



UNIVERSITY OF CALIFORNIA

Santa Barbara

Development of the MilliQan Bar Detector: a search for Millicharged Particles

The MilliQan Bar Detector is a newly installed subdetector of the Compact Muon Solenoid Experiment at the Large Hadron Collider. It will search for Millicharged particles, a hypothesized dark matter candidate characterized by an apparent electric charge on the order of $\approx e \cdot 10^{-3}$. The detector has completed commissioning and testing has begun on the full configuration, with early stages of data-taking and background analysis being performed.

A Dissertation submitted in partial satisfaction of the
requirements for the degree Bachelor of Science
in Physics

by
Samuel David Alcott

Supervised by:
David D. Stuart

June 2023

Preface

The purpose of this senior thesis project is to record my most meaningful contributions to the MilliQan Experiment. I have attempted to strike a balance between breadth and depth, emphasizing the aspects of the experiment I was more deeply involved with over those of which I am simply aware. However, I have also tried to ensure that almost every aspect of the experiment is mentioned and that necessary context is provided. It is my hope that the sections I have written in greater depth can serve as a valuable resource for members of the collaboration during detector interventions and problem solving, and that the less-detailed sections contribute a more holistic picture of the experiment and its purpose.

This thesis is written in sections primarily focused on construction, calibration, and analysis. I have been involved in each of these stages, so I have ordered these main sections roughly chronologically according to when I participated in each. I hope this helps to deliver a narrative following the detector's development. In order to provide full context, I will briefly summarize the timeline of my involvement in the experiment:

In March of 2021 I contacted Prof. David Stuart at UCSB asking about research opportunities, which he had instructed interested undergraduate students to do through his webpage. He very generously delivered a long and detailed pitch of MilliQan, and explained that the project was going to enter construction in summer and needed the hands of eager undergrads. I leapt at the opportunity, and began an independent research project with him over spring quarter of 2021. In this project I looked at older data from the MilliQan Demonstrator, and drew conclusions about what should be changed for the full-scale detector. These decisions had already been made at this point in time and the project was purely educational, but it was nice validation to learn that I came to the same conclusions as the collaboration.

I then joined the effort full-time at the start of summer, supported by the generous donors of the College of Creative Studies' SURF program. At this point the "raw" material for the bar detector had been obtained and was waiting in a lab in Broida Hall. This included large scintillator components, reflective and light-blocking material, and many large boxes of Photomultiplier Tubes left over from an older experiment and shipped from Ohio State University, MilliQan's founding institution. During this time, and the following academic school year, I worked with Katherine Larina, another undergraduate, to wrap the scintillator bars, and to test and calibrate roughly 120 Photomultiplier Tubes to select those ~ 72 best suited for use in the detector. We then mounted the Photomultiplier Tubes to the scintillator and prepared the components for shipment.

A few smaller projects were completed during this timeframe, including a dark-rate study on MilliQan bars, a test of light collection efficiency in detector panels, and a study of Cadmium-109 X-ray penetration depth in the detector among others, all of which are covered in this thesis. Katherine and I were trained one way or another during this time by the then-current members of the UCSB MilliQan team: David Stuart, Claudio Campagnari, Matthew Citron, Ryan Schmitz, Francesco Setti, Keegan Downham, and Joe Crowley.

MilliQan commissioning then began over summer of 2022, where the completed MilliQan components were shipped to Switzerland. A contingent of researchers from UCSB and other MilliQan institutions, including myself, joined this effort and traveled to CERN. I was there for 10 weeks, working to install, calibrate, and cable as much of the detector as possible, and also test the effect of the massive CMS magnetic field on our Photomultiplier Tubes. During this time I also assisted with some testing of the MilliQan slab detector.

After leaving Switzerland, the commissioning effort has continued in several phases and full data-taking just began in June of this year. As a result of this schedule there has not been a lot of data to analyze from the finished detector, but during my 4th year at UCSB I have used what is available to better understand the detector's backgrounds, and those efforts are included here.

Given this context, it should come as no surprise that my thesis is almost fully concerned with the MilliQan bar detector, with brief discussion of the slab detector. There is also a heavy emphasis on construction and commissioning, detail that I believe will be most useful to the collaboration and for which I am most likely to have specialized knowledge. My work on data analysis is not as technical or ground-breaking, but showcases exciting measurements and comparisons made with real detector data, and provides a satisfying conclusion to my time spent on the experiment. I have had the privilege to access many different stages of High Energy Experimental research through this novel project, and for that I am incredibly thankful.

Contents

1 Motivation	7
1.1 Dark Matter	7
1.2 The Standard Model	7
1.3 The Dark Photon	9
1.4 Millicharged Particles	11
1.5 mCP Searches	12
1.6 mCP Production	13
2 Detector Design	15
2.1 Detecting mCPs with Scintillator	15
2.2 Design of the Bar Detector	16
2.3 MilliQan Demonstrator	17
2.4 The Slab Detector	19
3 Experimental Setting	22
3.1 The Large Hadron Collider	22
3.2 The Compact Muon Solenoid	23
3.3 Observation and Drainage Gallery	24
3.3.a CMS Warehouse	24
4 The Bar Detector	26
4.1 Constraints at the LHC	27
4.2 Signal Criteria	27
4.3 Backgrounds	28
4.3.a Cosmic Muons	28
4.3.b Beam Muons	29
4.3.c Light Leaks	29
4.3.d Dark Rate	29
4.4 Bringing MilliQan to Life	29
4.4.a Engineering	30
4.4.b Transportation and Alignment	30
5 Construction	32
5.1 Scintillator	32
5.2 R878 Photomultiplier Tubes	33
5.3 MilliQan Amplification Base Boards	35
5.3.a Base Types: v3,v4,v5	37
5.3.b Base Testing and Troubleshooting	38

5.4	Bars	39
5.4.a	Bar Wrapping	39
5.4.b	Plastic Mounting and LEDs	43
5.4.c	PMT Mounting	44
5.4.d	LED Testing	45
5.4.e	Bar-Base Pairing	45
5.5	Panels	46
5.5.a	Panel Wrapping	46
5.6	Panel Mounting	49
5.7	Bar Assembly Units	50
5.7.a	Additional Mu-Metal Wrapping	52
5.8	Supermodules	53
5.9	Cabling	54
5.9.a	Cable Cutting	55
5.9.b	Cable Organization	56
5.9.c	Cabling Process	57
5.10	Power	61
5.10.a	High Voltage Supply	61
5.10.b	Low Voltage Supply	62
5.11	Supermodule Cage	63
5.12	Supermodule Insertion and Alignment	64
5.12.a	Installing Panels	67
5.13	Hodoscope	68
5.14	Voila	68
6	Calibration	69
6.1	Domino Ring Sampler (DRS4)	70
6.2	Gain	73
6.2.a	Single Photoelectron Pulses	74
6.2.b	Nominal Range	75
6.3	PMT Energy Response	75
6.3.a	Calibration with Cadmium-109	76
6.3.b	Height, Area, and Duration	78
6.4	Darkrate	79
6.5	Bar Procedure	79
6.5.a	Materials	79
6.5.b	Preparing the Bars	80
6.5.c	LV Light-Leak Safety Tests and Inspection	81
6.5.d	No Source/Darkrate Testing	82
6.5.e	Cd109 Source Testing	82
6.6	Panels	83
6.7	Units	83
6.8	Supermodule	84
6.8.a	Source Testing Procedure	84
6.9	Monitoring and Future Calibration	87

7	DAQ	90
7.1	Digitizers and Triggerboard	90
7.2	MilliQan Offline	91
7.3	Event Displays	93
7.4	Converting ROOT to CSV	94
8	Data Analysis	96
8.1	PMT Selection	96
8.2	Dark-rate instability study	97
8.3	Optical Gel Study	101
8.4	Light Collection in Bars	103
8.5	CMS B-field Study	104
8.6	Unusual Saturation Issue	108
8.6.a	Searching for Cause	108
8.6.b	Issue Tracking	109
8.7	Cosmic Shower Radiation Penetration	111
8.8	Direct Muon Hit Coincidence	113
8.8.a	Tagging Muons	114
8.8.b	Finding Coincident Hits	118
8.8.c	Comparison with Simulation	119
9	Conclusion	120
9.1	Acknowledgements	121

Chapter 1

Motivation

1.1 Dark Matter

In 1933, Swiss astronomer Fritz Zwicky encountered a gravitational anomaly when observing the Coma cluster, inferring the existence of an unseen substance that he called “dunkle materie”, or dark matter in English [5]. After almost a century of astronomical research, it has been determined that there is over five times more of this dark matter in the Universe than there is visible matter! This figure is striking, because it means that the summation of all human knowledge (save, perhaps, for General Relativity) may pertain only to a small fraction of the “stuff” in the universe! Dark matter is not just a rounding error that needs to be ironed out, it is one of the biggest questions in physics today. It is also one of the most elusive. Over 90 years since Zwicky’s discovery and we still do not know exactly what dark matter *is*—is it a uniform substance, or is it highly diverse? Is it made up of particles with wave functions like the matter we know, or is it something else entirely? Is it completely “dark”, or is there some way we can detect it? To answer these questions it is important to understand the basis on which most of these discussions are grounded, particle physics.



Figure 1.1: The Coma Cluster, as photographed in June of 2008 by the Hubble Space Telescope [15].

1.2 The Standard Model

We know that dark matter interacts via gravity, but gravitational attraction does not give much indication of what dark matter *is*, only that it is there. For that we must turn to the study of everything that isn’t gravity: particle physics. Particle physics ironically finds its origins in roughly the same time period as Fritz’ “dunkle materie” observation, with Paul Dirac marrying a well-matured quantum theory to Einstein’s special relativity from earlier in the century in order to explain electron behavior in 1928. Unlike the study of dark matter, however, particle physics was an extremely productive field, rife with discoveries, over the decades that followed. The culmination of all this effort is known as the “Standard Model”, and it is perhaps the most well tested physical theory in existence. A commonly

cited measure of the Standard Model's accuracy is its prediction of the magnetic moment of the electron, which agrees with theory to one part in one trillion (10^{12}) [8]. This is roughly equivalent to measuring the distance between Santa Barbara and Miami to the width of a strand of spider silk [23].

The Standard Model claims that all visible matter is made up of 12 fundamental particles called fermions, and that the way these particles behave is described entirely by 5 more fundamental particles called bosons, making 17 total unique particles in existence. That means that everything we can see and touch are just varying combinations of these 17 ingredients. As most learn in grade school, almost all matter we encounter is composed of atoms, and these atoms are composed of protons, neutrons, and electrons. Electrons are already a fundamental particle, denoted e in Figure 1.2, but protons and neutrons can be broken down one step further into a variety of fermions called "quarks". There are six varieties of quark, cutely nicknamed "up", "down", "charm", "strange", "top", and "bottom", but we only need the up and down quarks to build every atom in existence. Indeed, a proton is simply two "ups" and a "down" while a neutron is two "downs" and an "up". The electron belongs to a different family of fermions called the "leptons". Three of these leptons, including the electron, have charge -1 in units of the charge of a proton. These are the electron (e), muon (μ), and tau (τ). They each have different masses, and only the electron is stable while the others decay quickly after being produced. The other three leptons are called "neutrinos", and true to their name they have a neutral electric charge. They come in three flavors corresponding to the e , μ , and τ , and there are rules about how they interact with other particles based on this flavor. Finally, each of these particles has an anti-particle twin, which carries the same mass but opposite charge and annihilates upon contact with it to produce energy.

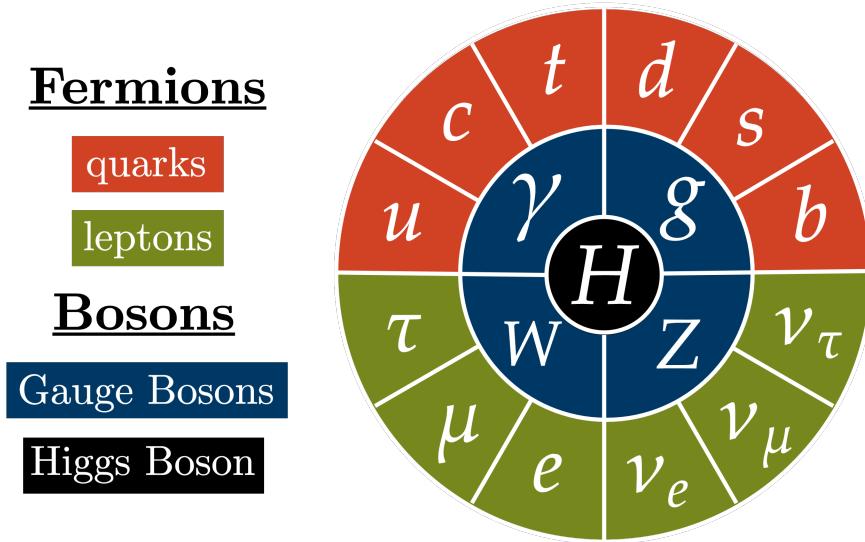


Figure 1.2: The particles of the Standard Model, in their common Greek-letter representation.

This discussion of charge and restricted interactions begs the question: where do these rules and ideas come from? What does it mean to have a charge? This is where the bosons come in. Four of these bosons: the photon (γ), the gluon (g), and the W and Z gauge bosons, are

called “force carriers”. Each of them interact with a subset of the fermions, as well as other bosons, with a strength known as charge. The easiest of these to understand is electric charge—this quantity determines how strongly a particle will interact with photons (such as visible light). The others are more complex in their interaction, but in short the gluon is responsible for holding quarks together and the gauge bosons are responsible for nuclear decay. The last boson, the Higgs boson, is not a force carrier. The Higgs is unique because it interacts with every massive particle in the Standard Model, including itself, with an interaction “charge” known as...mass. It is thus responsible for particles having mass, and how much mass each of them have. This also means that the Higgs boson would interact with dark matter, since it has mass, as a natural consequence of any attempt we make at incorporating dark matter into the standard model.

1.3 The Dark Photon

There are many ways to incorporate dark matter into the Standard Model, and to build so-called *Beyond the Standard Model* (BSM) theories that account for dark matter. One class of these models is called the “Dark Sector” (or Hidden Sector), and it has some natural appeal for particle physicists. In Dark Sector models, dark matter is not a uniform substance, but a diverse collection of particles much like in the Standard Model, and they are attracted to one another by *dark forces* mediated by dark bosons. One can see why particle physicists would be excited by the idea of a whole new “dark” Standard Model of particles to discover! This is not to say that the world of dark matter is a mirror copy of the Standard Model, but it is reasonable to assume that there would be some similarities.

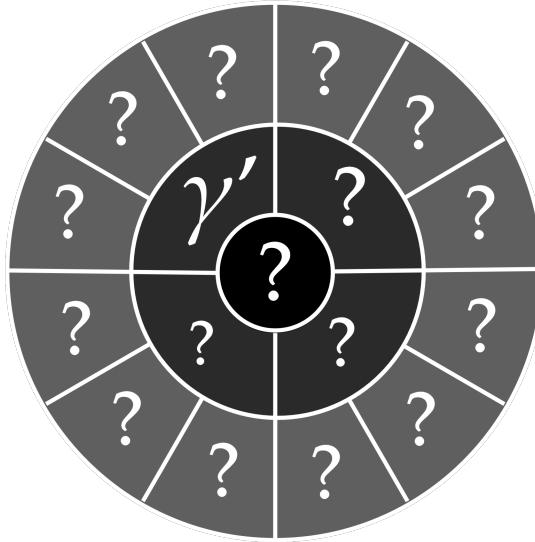


Figure 1.3: Could the Dark Sector look something like the Standard Model? Perhaps, perhaps not. But it is likely that it would at least share one feature in common: a unitary gauge group of degree 1, otherwise known as a photon. This “dark photon” would imply a dark QED (Quantum Electrodynamics).

The most natural assumption is that the math that describes the existence and interaction of these particles works the same as it does for visible matter, i.e. Quantum Field Theory and Group Theory. The latter of these provides an elegant way of understanding

the gauge bosons of the Standard Model (SM), where each generates transformations that preserve a particular *symmetry group*. The simplest of these is the unitary group, $U(1)$, which is mediated in the photon in the SM. The other groups in the Standard Model are $SU(2)$ and $SU(3)$, which are preserved by the “weak” ($W+Z$ bosons) and “strong” (g) forces. A consequence of this model is that it is possible for these groups to “mix”, allowing particles to interact with multiple bosons, i.e. multiple groups. $U(1) \times SU(2)$ is called the “electro-weak” force, and $U(1) \times SU(2) \times SU(3)$ represents the entirety of the Standard Model.

It is not known what group symmetries a dark sector would have, but the simplest extension would be its own version of the unitary group $U(1)$, which we will call $U'(1)$. Holdom [12] introduced this group, and found that its existence has calculable consequences. For this we return briefly to the former mathematical basis for particle physics: Quantum Field Theory (QFT). As with almost any theory in modern physics, the equations of motion for particles in QFT can be found by minimizing their Lagrangian, which is a scalar (coordinate invariant) quantity that holistically describes the physics and constraints of a system as a series of terms. QFT treats the fundamental particles of the SM as being disturbances in quantum fields that permeate space, so the terms of the SM Lagrangian pertain to each of these fields and the interactions between them. The SM Lagrangian (actually a Lagrangian *density*), \mathcal{L}_{SM} is very long, but an extremely abbreviated version of it can be found on coffee mugs and T-shirts at CERN, a premier particle physics research organization:

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\bar{\psi}\not{D}\psi + \psi_i y_{ij}\psi_j\phi + \text{h.c.} + |D_\mu\phi|^2 - V(\phi)$$

It is not important for our purposes to understand all of these terms (and the hundreds of terms that hide within), but they are nicely summarized to the layperson by Woithe et al. [13]: The first term represents the interactions that the force-carrier bosons have with one another, the second term represents the interactions that those force carriers have with the fermions, the third term represents how fermions interact with the Higgs field and thus get their mass, the fourth term (h.c.) are the hermitian conjugates of these terms (necessary for the algebra), the fifth term explains how the W and Z bosons interact with the Higgs field and thus get their mass, and the last term is the Higgs potential, and explains how the Higgs interacts with itself.

Now that we have the SM Lagrangian (density) under our belt, what happens to it when we add a whole new field under this dark $U'(1)$ symmetry? Holdom [12] found that this leads to a few more terms in the Lagrangian:

$$\mathcal{L} = \mathcal{L}_{SM} - \frac{1}{4}A'_{\mu\nu}A'^{\mu\nu} + i\bar{\psi}(\not{\partial} + ie'\not{A}' + iM_{mCP})\psi - \frac{\kappa}{2}A'_{\mu\nu}B^{\mu\nu}$$

The first term arises from the addition of this symmetry, and is a massless dark boson similar to γ that is comically called the “dark photon” [2]. It is represented by a field strength $A'_{\mu\nu}$ that derives from a vector potential A'_μ , and is defined:

$$A'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$$

The second term represents a dark Dirac Fermion, much like the electron, which has a dark charge e' (and from symmetry we may argue $e' \sim e$), and a mass M_{mCP} (the subscript will shortly become clear) [2]. The final term is the most interesting, and arises from the fact

that just as with the symmetry groups in the standard model, the new $U'(1)$ group can be mixed with the old $U(1)$ group to form $U(1) \times U'(1)$. This term therefore couples the dark photon field $A'_{\mu\nu}$ with the SM *Hypercharge* $B_{\mu\nu}$, which is similar to the standard photon field but obeys the symmetries of $SU(2)$ as well, that way it represents the electro-weak force. This term brings with it a coefficient κ , which from something called one-loop effects can be argued to be [3]:

$$\kappa \sim \frac{\alpha}{\pi} \approx 0.00232$$

Holdom [12] then performs a trick to remove this mixing term, by redefining the gauge boson's vector potential as:

$$A'_\mu \longrightarrow A'_\mu + \kappa B_\mu$$

which is perfectly allowed. Substituting this change and performing some math we end up getting a new but equivalent Lagrangian

$$\mathcal{L} = \mathcal{L}_{SM} - \frac{1}{4} A'_{\mu\nu} A'^{\mu\nu} + i\bar{\psi} (\not{\partial} + ie' \not{A}' - i\kappa e' \not{B} + iM_{mCP}) \psi$$

Which has major implications. By writing the Lagrangian in this form, we see that a new term has appeared in our dark-fermion expression, which looks a lot like the term that gave the dark fermion charge under the dark photon, only with a different coefficient and the SM hypercharge B instead of the dark photon. This means that the existence of a $U'(1)$ field that mixes with the standard $U(1)$ produces fermions (which I will call χ) that have mass and couple to the dark boson with a dark charge e' , but *also that couple to the SM hypercharge* with a charge $\kappa e'$! This means that the value of κ acts to suppress the strength of interaction between the SM and dark matter, meaning that these χ 's would look to us as if they have charge on the order of one-thousand times smaller than the electron [3]!

1.4 Millicharged Particles

We have seen that the direct consequence of introducing a dark sector with a dark photon, otherwise known as the “abelian gauge group” $U'(1)$, is that this dark sector will also be populated by dark fermions. We also saw that due to “mixing” between groups, these fermions will interact at a very feeble amount with the photon of the standard model, giving them a charge $Q \sim O(10^{-3}e)$. These are called “millicharged particles”, and detecting them may crack the lid on the dark matter mystery. Millicharged particles, or mCPs, are a particularly compelling candidate because their detection is fairly model-independent. Almost every dark sector theory incorporates a dark photon, so an experiment seeking to detect mCPs need not commit itself to any specific model in order to make a major discovery [2]. The question of the existence of millicharged and other fractionally-charged particles is also one that merits investigation, regardless of its connection to dark matter. The smallest known unit of charge is that of the d , s , and b quarks, with $Q = -1/3e$ (meanwhile the u , c , and t quarks have $Q = 2/3e$), where e is the magnitude of the proton/electron charge. These quarks are not stable on their own, however, so the electron might as well be thought of as having the smallest unit of charge, e . Bizarrely, however, there is no law that prohibits particles with smaller charges than this, and the question of whether charge is fundamentally quantized has yet to be answered [9] [7]. The discovery of an mCP would dethrone the electron as the smallest unit of charge, and may have major implications even beyond the identity of Dark Matter.

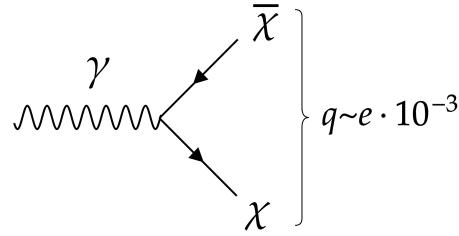


Figure 1.4: Here a very basic Feynman diagram depicts a photon (γ) pair-producing an mCP (χ) and its anti-particle ($\bar{\chi}$), with a coupling strength that gives them a charge of $q = \pm O(10^{-3}e)$.

1.5 mCP Searches

Millicharged particles are not themselves a novel idea, and as a result there are already many constraints on what their possible charge and mass might be. These constraints have arisen from two genres of studies: indirect observation and direct searches.

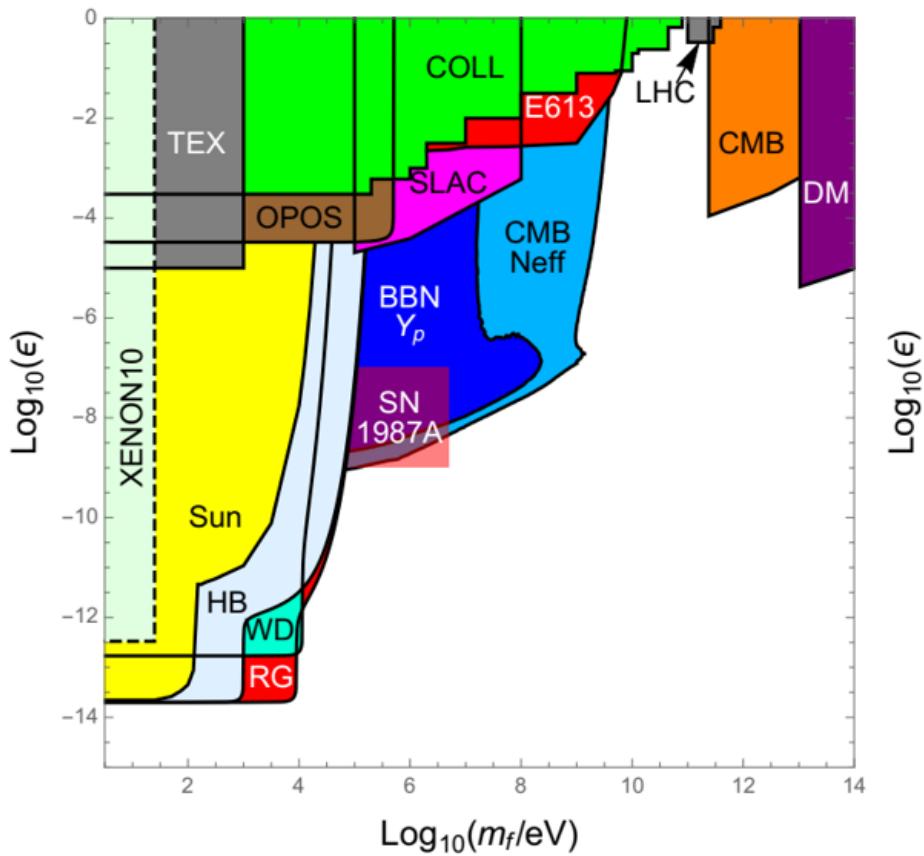


Figure 1.5: Graphed on the mass-charge phase space (where ϵ here is the ratio of the mCP charge to the electron charge), these are the current constraints on mCPs. Colored regions indicate charge-mass combinations that an mCP *can not* have based on direct and indirect searches, as well as theoretical limits.

In indirect observation studies, physicists determine what charges and masses an mCP *can't* have, otherwise it would have already been discovered or the physics we know would be significantly different. Some of these constraints come from collider studies, and others come from observations in high energy astrophysics, which can probe the existence of certain mCPs based on how many would have been produced in the early universe [9]. A direct search, on the other hand, is an experiment that involves a detector specifically built for finding mCPs. These are far more rare, with perhaps the most prominent being the mQ experiment at SLAC in the 1990s [11]. The combined efforts of indirect and direct searches have constrained the mass-charge phase space to that shown in Figure 1.5.

This graph is cramped, with many studies ruling out the same regions for mCP existence, but there remains a tantalizing corridor of unexplored phase space in the 100 MeV-100 GeV range that has not been fully ruled out even to the $Q = 10^{-1}e$ charge level. This mass range is especially promising, as it is well suited to the power of modern colliders.

1.6 mCP Production

We have seen that mCPs may interact with the SM photon, but this tells us little about how they could actually be created in an experiment. It turns out that mCPs may be pair-produced (as a particle and an antiparticle) anywhere that fermions are pair produced in standard particle interactions, *so long as it is kinematically allowed*. This means that the rate of these types of interactions in collisions at a standard particle collider (related to a quantity σ called the cross-section) will depend greatly on the mCP mass. It also turns out that the likelihood of one of these interactions to produce mCPs instead of another particle (related to something called the “branching ratio”) will depend quadratically on its charge, meaning that their production will be subdued even further. Figure 1.6 shows each of the production mechanisms for mCPs, graphing the cross-section times the branching ratio as a function of the mCP mass (denoted m_χ).

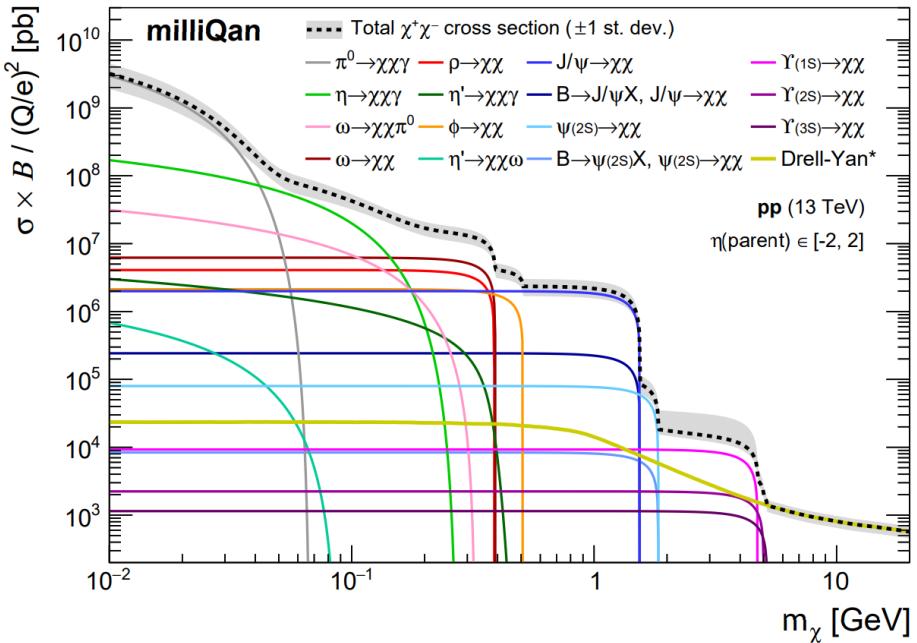


Figure 1.6: Cross sections for mCP production mechanisms for mCPs in 13 TeV proton-proton (pp) collisions.

There are many different production mechanisms summed up in the graph, but most of them are either Meson Decays, Dalitz Decays, or Drell-Yan scattering (note: a Meson is a quark-antiquark pair, and a hadron h is a bundle of three quarks, like the proton and neutron). Each of these staple processes are shown in Figure 1.7, in the form of *Feynman diagrams*. These diagrams come in a simple pictorial form, but precisely represent the imposing mathematics of perturbative quantum electrodynamics (QED). Imagine these diagrams, instead, as a graph of particle trajectories through time and space, with space on the horizontal axis and time flowing from the top to the bottom. If we consider again the compelling area of phase space in Figure 1.5, we see that this range of masses is dominated in particular by Drell-Yan processes, which result from collisions of hadrons. There are no historical examples of mCP direct detection experiments that operated at hadron colliders, motivating such a search even further.

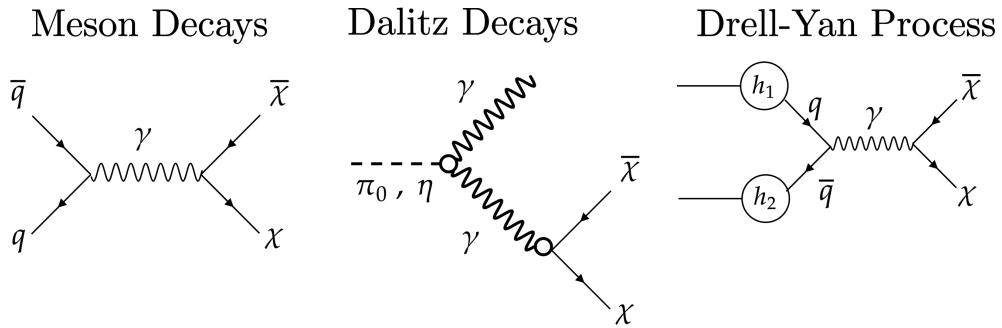


Figure 1.7: Three of the more relevant production modes, as Feynman diagrams. In Figure 1.6, we can see that Drell-Yan production begins to dominate in the GeV range.

Searching for mCPs is a convenient approach for dark matter direct-detection experiments, since the relevant physics relates purely to QED, which is the best studied area of the SM. This essentially means that mCP detectors need only be hyper-sensitive charged particle detectors—their novelty comes only from their design, as the necessary technology has already existed for decades. It is also clear that there is a unique opportunity for mCP discovery at this time in history, with very few experiments looking for mCPs at hadron colliders, particularly in the 100 MeV - 100 GeV mass range. This strongly motivates the creation of a new dedicated mCP detector, but also raises two initial questions: how will you make a detector sensitive to the minuscule charge of an mCP, and where will you put your detector in order to find them? These questions are the concern of the next two chapters.

Chapter 2

Detector Design

2.1 Detecting mCPs with Scintillator

We have seen that if we model dark matter as being part of a dark sector with an abelian gauge field $U'(1)$, we predict the existence of dark fermions called mCPs with apparent charge $O(10^{-3}e)$, making them not exactly *dark*, but extremely *faint*. This means that they may be detectable using traditional particle detector technology, which employs electromagnetic *scintillation* to detect energetic and charged particles. Scintillation is a process whereby certain materials, often crystals or plastics, produce light at a consistent visible or ultra-violet wavelength when charged particles or high-energy photons pass through them. There are a few ways this can happen, including Compton scattering and pair-production (at high enough energies), but the most common is ionizing radiation whereby a charged particle or energetic photon disturbs the structure of electrons in the scintillator as it passes through, and they produce scintillation light as they settle [31]. The more energy deposited, the more scintillation light gets produced. For charged particles, the rate of energy deposition per unit length is proportion to charge squared [31]. This means that scintillators could theoretically be activated by a dark fermion with an effective milli-charge, but it would need to pass through a much larger distance of scintillator in order to deposit the same amount of energy as an electron. This explains why mCPs have not been observed in currently operating detectors, which minimize scintillator thickness to better constrain particle trajectories, and it motivates a dedicated mCP detector at the forefront of high energy limits. Such a detector would have very thick scintillator, bringing the rate of ionization energy deposition from a hypothetical mCP to detectable levels. The goal of such a detector would naturally be to measure the charge Q and mass m_{mCP} of a *milli*-charged particle, not unlike the oil drop experiment performed by Robert Millikan to measure the charge of the electron [30] (I will refrain from explaining the wordplay further). In doing so, a detector would need to address a few key challenges:

- (a) **Charge Sensitivity.** Because mCPs have characteristically low charge, a detector must have very thick scintillator in order to detect them. This forces a balance between scintillator quality and quantity, because it would be prohibitively expensive to purchase adequately large chunks of the finest scintillator.
- (b) **Mass Sensitivity.** Arguments for $O(10^{-3}e)$ charge do not have a direct mass-analog for mCPs or dark-fermions. All there is to go off of is the known constraints on mCP mass from Figure 1.5, and the enticing 100 MeV-100 GeV window is a large range to

scan. It is thus a goal to maximize mass sensitivity as well as charge sensitivity. Unfortunately particles with higher masses are pair-produced less often in collisions [32], as are particles with very small charges [2]. Therefore for mCPs created using particle colliders, this means we want to maximize the energy and frequency of collisions, as well as the solid-angle acceptance of the detector around the collision point (to catch as many as possible). This also introduces a cost constraint, as it calls for not only a thick detector, but a wide one as well!

(c) **Overwhelming backgrounds.** Any detector sensitive to the feeble interaction of an mCP will be blinded by interactions with $O(e)$ charged particles and energetic photons. This means that a detector must find some way to “filter” out undesired particles produced in collisions without blocking mCPs, and that it must itself be heavily shielded from radiation, otherwise mCPs will be lost like whispers in a hurricane. There will inevitably be some disturbance of the detector by other particles even after these precautions, so the design of the detector must find a way to distinguish clearly between this *background* and the *signal* that would be produced by a real mCP.

2.2 Design of the Bar Detector

The MilliQan Bar Detector is an attempt to apply this detection philosophy. It targets a relatively low-charge and high-mass region of unprobed mCP phase-space (see Figure 1.5), achieved through its long titular bars. The detector is comprised of four co-linear layers, each containing a 4×4 grid of these scintillator bars with Photomultiplier Tubes (PMTs) mounted to the end. PMTs convert scintillation light to photoelectrons, which are multiplied to generate an electric signal (more on them later).

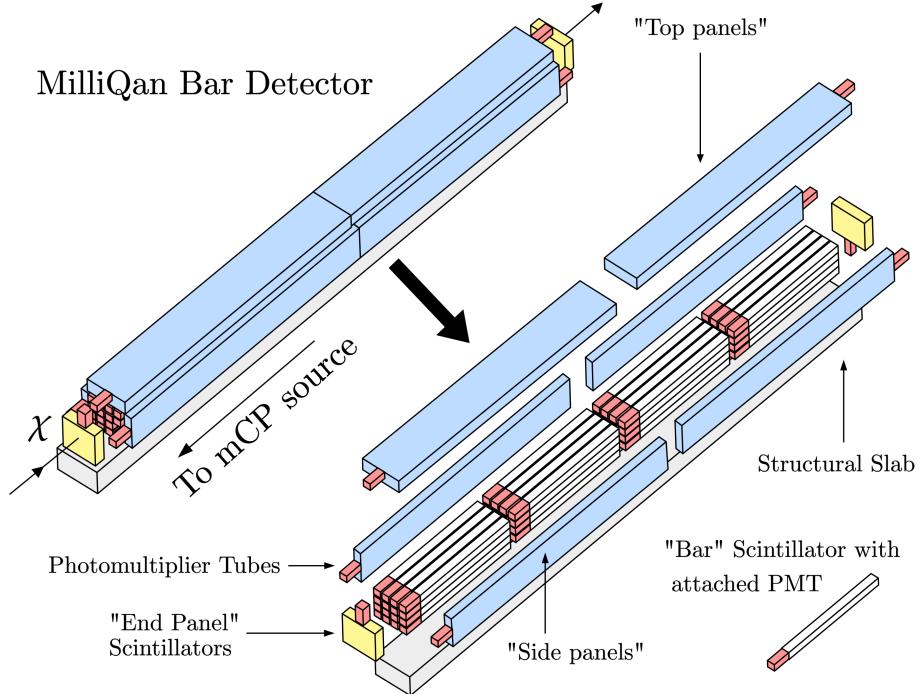


Figure 2.1: An exploded diagram of the full bar detector, at a conceptual level without any practical engineering considerations.

Each bar is $5\text{cm} \times 5\text{cm}$ wide, to match the diameter of our chosen cost-effective Hamamatsu R878 PMTs. They are grouped in four layers to suppress known mCP backgrounds, and they are 60cm long to maximize charge sensitivity while allowing all layers to fit at an angle of 43.1° in our detector's experimental site. These bars make up the bulk of the detector and provide all of its mCP sensitivity. Surrounding these inner layers of bars are four $25\text{cm} \times 170\text{cm}$ scintillator panels (side panels) in groups of two on the left and right flanks of the detector, two $30\text{cm} \times 170\text{cm}$ scintillator panels (top panels) on the top of the detector, and two $25\text{cm} \times 25\text{cm}$ slabs of scintillator (end panels) on the front and back along the axis of the detector.

All of these panels serve as both active and passive shielding for the detector: active in the sense that they will be used as veto criteria in the final signal trigger, excluding events in which they are activated, and passive in the sense that they will naturally absorb a significant amount of incoming radiation by virtue of their thickness. The end panels have an additional purpose, being to tag high-energy non-mCP particles originating from the same source as the mCPs. This assists in alignment for the detector and prevents unusual false-positive signal events that occur when a non-mCP passes perfectly through the gaps between bars (to be discussed in Section 4.3.b).

The bar detector is fine-tuned to the experimental setting in which we are building it, which I will elaborate greatly on in the next two chapters, but the general philosophy behind its $4 \times 4 \times 4$ design is one of statistical background suppression. The detector is aligned such that mCPs created at the interaction point of a collider will travel straight through it, and in doing so will deposit a very small amount of energy in each layer of the detector, hitting four bars in a row. By *only* considering events in which four bars in a row are triggered, with time ordering and no hits in any other bar or panel, we are able to statistically exclude radiation (which usually deposits a lot of energy and effects multiple channels within a single layer) and random firing of our PMTs that could produce a false positive. There is nothing optimal about the 4×4 cross-section of the detector, this is merely a constraint of the budget and size of its experimental site. Ideally a MilliQan-like detector would be as wide as possible, and one successor called the FORward MicrOcharge SeArch (FORMOSA) is already planned with a 20×20 bar cross-section [36].

2.3 MilliQan Demonstrator

Before the bar detector, there was a prototype detector called the MilliQan demonstrator. This demonstrator was constructed and collected data from 37.5 fb^{-1} worth of collisions (these are units of integrated luminosity) in 2017 and 2018 [3]. As the name suggests, the demonstrator was a proof of concept for the bar detector, being smaller in scale and experimenting with a number of design elements in order to decide which would be best for the full detector. Apart from its smaller size, I will highlight the key differences between the demonstrator and bar detector, and what discoveries lead to those changes.

- (a) The demonstrator bars were slightly longer than the bars eventually used in the bar detector (80cm vs 60cm). This is because both detectors were designed to fully occupy the space available along the detector axis, and the demonstrator only had 3 layers instead of 4. This change was made because of the critical discovery that 3 layers were not sufficient for statistically suppressing backgrounds in the demonstrator run,

and that in fact more sensitivity could be achieved with 4 layers of bars at 3/4 the original thickness.

- (b) The demonstrator panels were significantly thinner than those in the final bar detector. This was due to funding restrictions, and they were not as effective as a result.
- (c) Rather than having just a front and back panel like the bar detector, the demonstrator had four intermittent scintillator slabs and two blocks of lead. The purpose of the slabs was similar as in the bar detector, with the additional inner-layer slabs providing more spatial and timing information on high-energy through-going particles, and the lead slabs to block secondary particles produced in one layer from hitting adjacent layers [3].
- (d) The demonstrator has a *hodoscope* at each end, which themselves are composed of a pair of perpendicularly crossed bunches of 2x8 small scintillator strips. This creates a much more precise cartesian grid for tracking high energy through-going particles, and thus can be used as a tool assisting in precision alignment and validation of alignment for the detector. There are plans for a hodoscope to be added to the bar detector as well, but implementation of those hodoscopes has not begun.

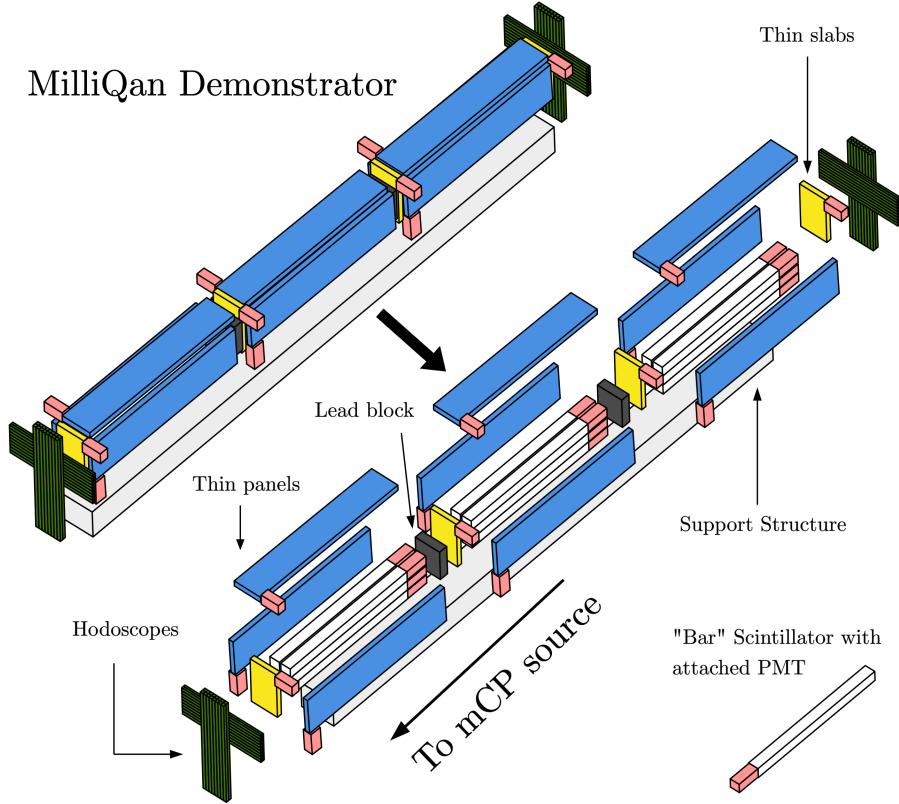


Figure 2.2: In the same style as Figure 2.1, the MilliQan demonstrator detector at a conceptual level. All scintillator parts (including bars) were different shapes for the demonstrator, as indicated by slightly different colors in the case of the panels.

Overall the demonstrator was a major success for its scale and timeline, resulting in a paper and probing mCP phase space up to and in some regions even beyond world-leading sensitivity. Those results can be compared with the Run 3 detectors in Figure 2.6. Another benefit of the experience installing the demonstrator is that it solved a lot of logistical issues in advance for the bar detector, such as determining how to safely route power and ethernet. The installation of the demonstrator also motivated the installation of a support and alignment structure that was eventually recommissioned for use in the full bar detector, though this was not always the plan. Originally the bar detector was going to be made up of a 9×9 array of bars, with significantly larger scintillator shielding, but this plan was abandoned in favor of a smaller bar detector and additional sister detector called the “slab detector” [4].

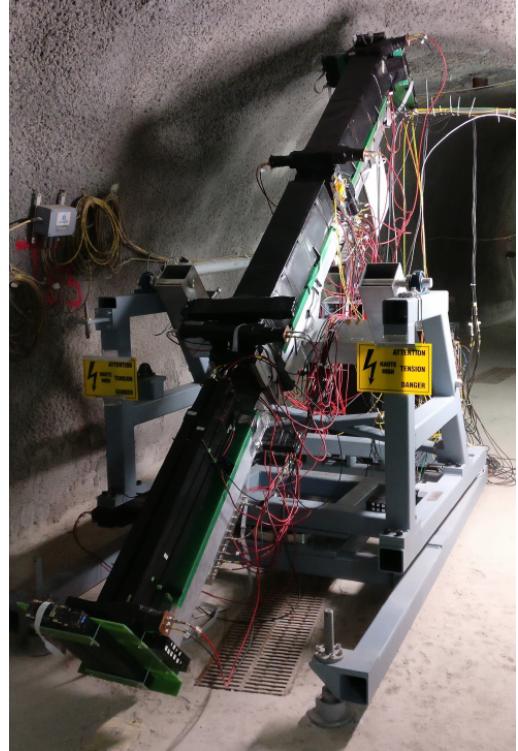


Figure 2.3: An actual image of the MilliQan demonstrator after its construction, with the front-most panels not attached.

2.4 The Slab Detector

Concurrently with the development of the bar detector, some MilliQan collaborators such as Ohio State University, University of Nebraska Lincoln, and New York University have been simultaneously developing and constructing the MilliQan *slab* detector, which employs the same technology but targets a different region of mCP phase space. Rather than searching for very low-charge mCPs by using thick scintillator oriented to maximize ionization by charged particles, the slab detector trades charge sensitivity for mass sensitivity by using its available scintillator volume to maximize width (and thus solid-angle acceptance) instead of thickness.

The slab detector is named for its “slabs”, which are $40\text{cm} \times 60\text{cm} \times 5\text{cm}$ rectangular scintillator prisms with four mounted PMTs each. The detector is a diagonal set of four shelves, each carrying four interlayered groups of 3 of these slabs, for 12 slabs to a layer and 48 slabs total. This arrangement maximizes the accessibility of the slabs, which sit horizontally on their shelves and can be easily inserted or removed [4].

It is worth noting that the slab detector does not have the shielding that the bar detector has, nor does it have dedicated scintillators for tagging high-energy through-going particles such as a hodoscope, beyond the slabs themselves. Doing so would be difficult to reconcile with the otherwise desirable design of the slab detector, and as a result these features were excluded.

The slab detector was developed very much so in conjunction and communication with the bar detector, though it was conceived of after the bar detector was already in development (during the drafting of [4]), so much of the technology and methods are carried over between detectors. Much of the construction and calibration work that I contributed to on the bar detector was used to inform work on the slab detector, and I myself worked with a few students at New York University to devise a calibration procedure. Beyond that, however, my involvement in the installation of the slab detector was minimal and my expertise almost exclusively relates to the bar detector. As a result I will not be discussing the slab detector in significant detail throughout this writing.

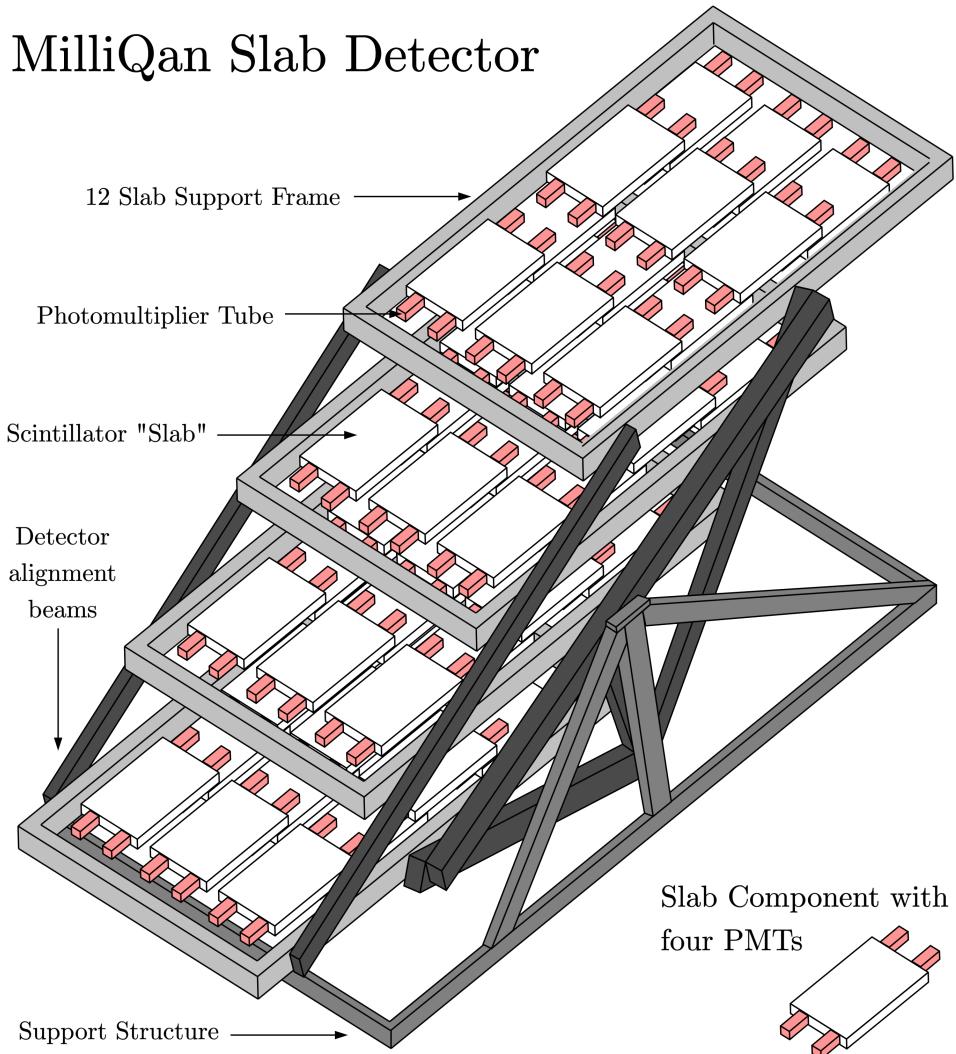


Figure 2.4: Conceptual design of the slab detector, which contains 48 slabs arranged in four layers of 12 each, overlapping to avoid gaps. In this earlier model the shelves pivoted and could switch between a data-taking position and an intervention position, but in the final detector these shelves are rigid.

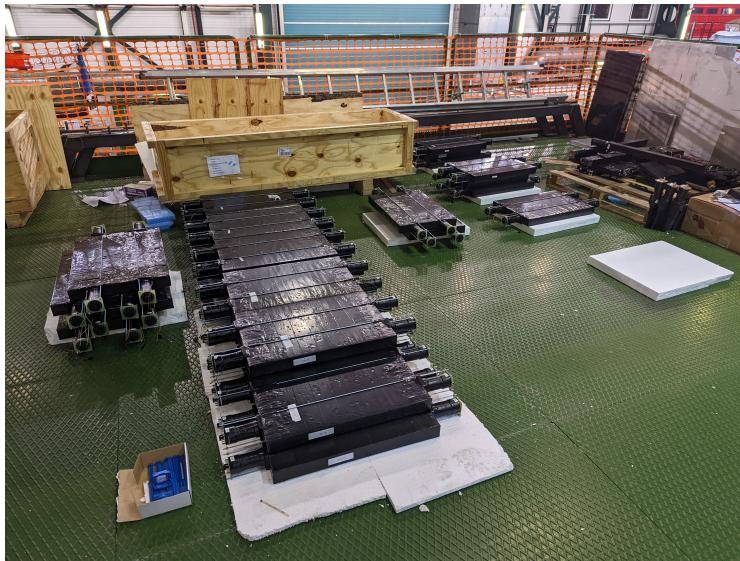


Figure 2.5: 24 slab components stacked atop Styrofoam following their assembly, awaiting installation in the final slab detector.

The slab detector is now undergoing its own commissioning, but it is on a different timeline than the bar detector and there is no exact date known for when it will begin taking data. The collaboration hopes to complete its installation and begin data taking in 2023. Once it is installed and data taking is complete, the slab detector will be sensitive to lower charges in higher-mass regions (10 GeV) scale, whereas the the bar detector will excel at low charge sensitivity in the (100 MeV) scale.

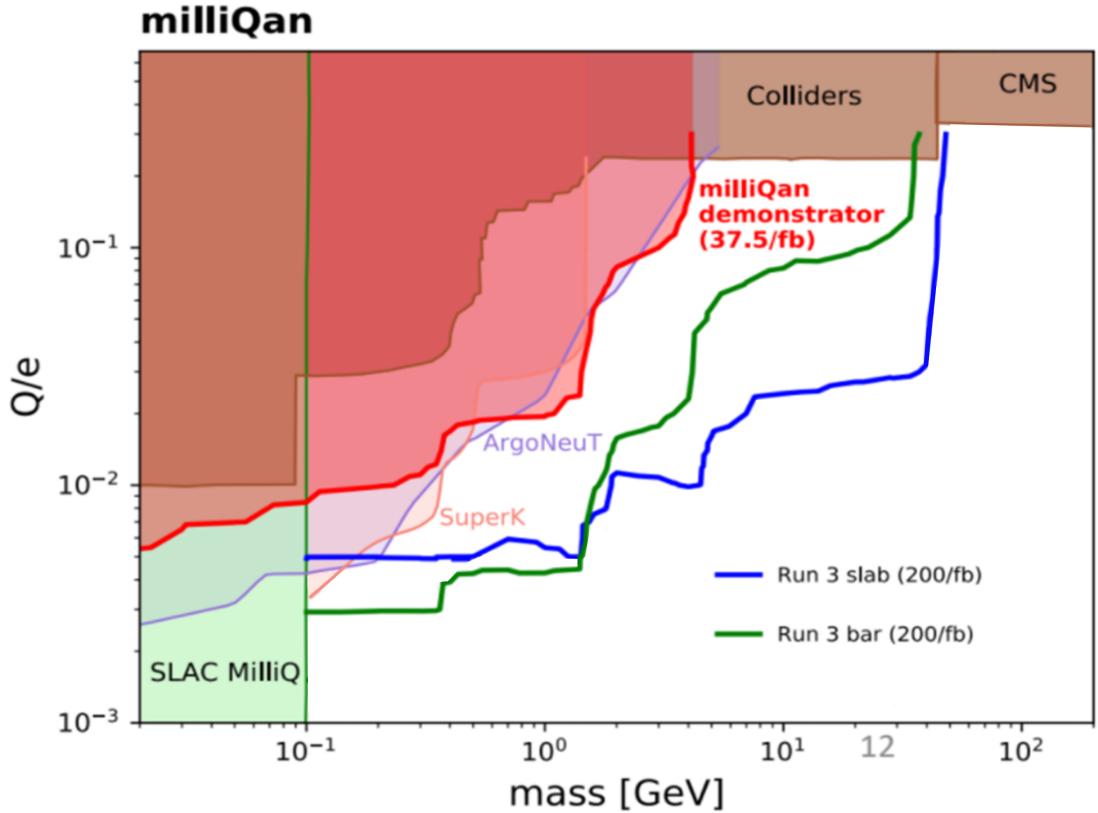


Figure 2.6: Sensitivity projections in phase space of mCP charge and mass for the demonstrator, as well as the bar and slab detectors. The demonstrator was already world-leading in charge sensitivity in some areas with mass above ~ 150 MeV, and the bar and slab detectors will be significantly more charge-sensitive than current limits in the 100 MeV-10 GeV mass range. The slab detector, due to its wider solid-angle acceptance but thinner IP-aligned scintillator excels in higher masses but also higher charges, whereas the bar detector excels in lower masses at lower charges.

Chapter 3

Experimental Setting

3.1 The Large Hadron Collider

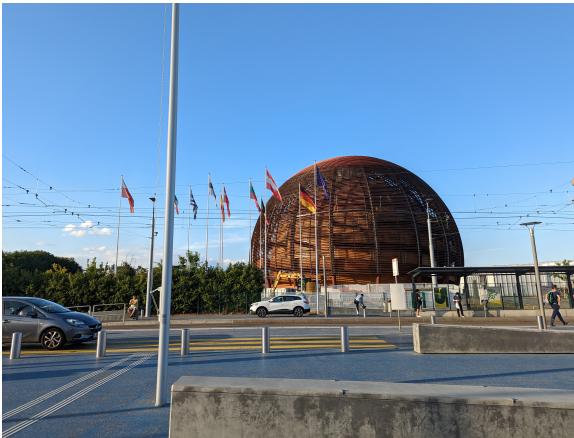


Figure 3.1: CERN’s iconic museum, *Globe de la science et de l’innovation*, marks the entrance to its main campus in Meyrin. A hundred yards to the left of the photo lies the French-Swiss border.



Figure 3.2: Between the CERN hostels on main campus is a 2m tall statue of the Hindu god Shiva, gifted by India as a celebration of its long-term participation in CERN [22].

The Large Hadron Collider (LHC) is a particle collider found at the European Organization for Nuclear Research (CERN). Measuring 27 kilometers in circumference, it is the largest circular particle collider in the world [16]. First powered on in September of 2008, the LHC has gone through a series of regularly scheduled “Runs”—where collisions take place and data is collected—and “Long Shutdowns”—periods with no collisions where maintenance and upgrades are performed on the LHC and its experiments [18]. The first Run of the LHC saw the famous discovery of the Higgs Boson, announced in July of 2012. Long shutdown 1 (LS1) took place in 2013 and 2014, followed by Run 2 of the LHC between 2015-2018 [14]. During Run 2 the energy of collisions was upgraded from 7-8 TeV in Run 1 all the way up to 13 TeV, close to the 14 TeV maximum collision energy of the collider [17]. The luminosity, a quantity related to the rate of collisions, was also increased between these runs. As of LHC Run 3, which began in July 2022 and will go until December 2025, the LHC will record proton-proton collisions at an astounding center-of-mass energy of 13.6 TeV. The record-high energy of LHC proton-proton collisions (good for Drell-Yan) is optimal for the mass sensitivity of any mCP searches, as is its record-high luminosity.

3.2 The Compact Muon Solenoid

The Compact Muon Solenoid (CMS) is one of the two large general-purpose experiments at the LHC, directly across the LHC ring from its sister experiment ATLAS. As its name suggests, CMS is considerably smaller than its counterpart (21×15 × 15 meters in dimensions), due to different design compromises being made by each team [19]. CMS operates with a 3.8-Tesla magnetic field sourced by a massive superconducting solenoid within. This magnetic field is confined by a massive steel yoke, and as a result the experiment weighs roughly 14,000 tonnes (about twice the weight of ATLAS) [19][20]. CMS is housed in the experimental cavern UXC-55 and operates approximately 100m underground, containing an interaction point (IP) known as “LHC Point 5” (P5) for proton-proton (pp) and heavy-ion (Pb-Pb) collisions. CMS hosts a massive and diverse experimental community, with CERN reporting “about 5,500 physicists, engineers, technicians, students, and support staff from 241 institutes in 54 countries” as of May 2022 [19].



Figure 3.3: CMS-S resides in Cessy, France, beside fields of wheat and sunflowers. Behind it, the snowy peak of *Mont Blanc*.

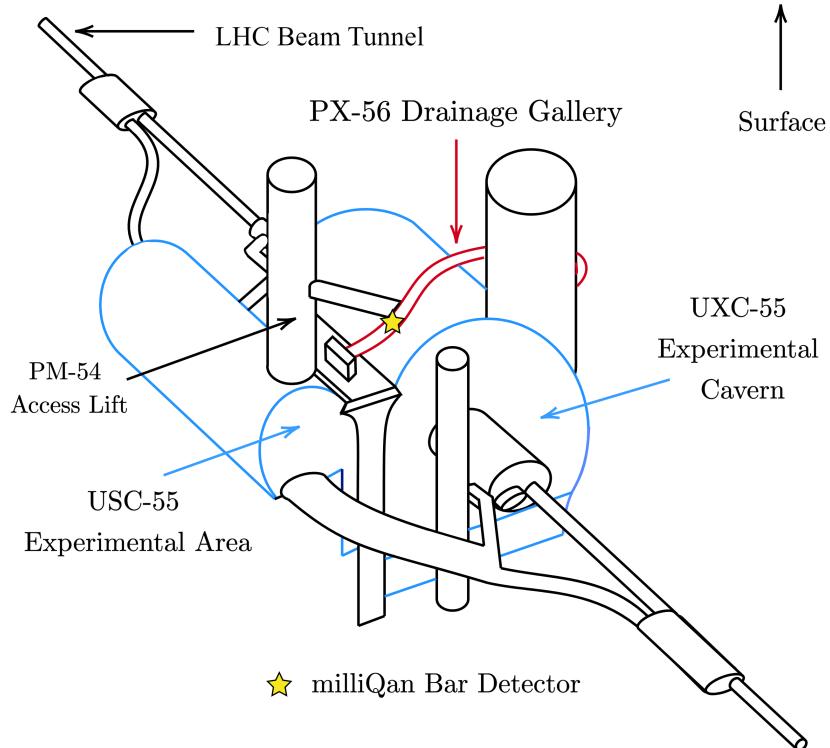


Figure 3.4: A diagram of the CMS experimental cavern and its connected passages, with the drainage gallery shown in red. The yellow star denotes the position of the detectors. The beam tunnel extends far beyond the diagram, forming the 27km circumference of the LHC.

3.3 Observation and Drainage Gallery

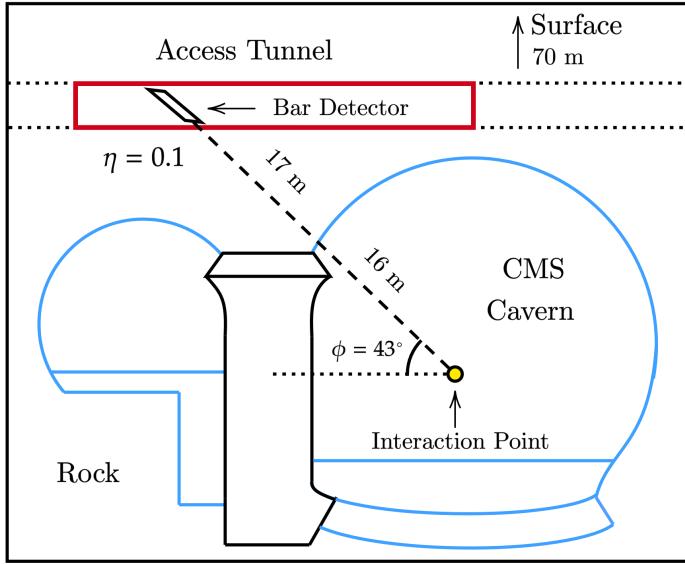


Figure 3.5: A cross-sectional view of the CMS cavern in relation to the bar detector. The detector has a pseudorapidity η of 0.1, which is a CMS coordinate indicating angle relative to the beam axis.



Figure 3.6: Katherine Larina and myself sat at a work table in the drainage gallery, positioned just behind the bar detector.

The UXC-55 experimental cavern is supported by the PM-54 access lift, the USC-55 Experimental Area, and the PX-56 Observation and Drainage Gallery, a tunnel connecting UXC-55 and USC-55, situated primarily above the CMS cavern. The gallery is well shielded from most known LHC radiation, save for muons generated in collisions (“beam” muons), and its depth underground significantly reduces cosmic ray flux. Due to this convenient background suppression, and its accessibility during beam-on periods of the LHC, the PX-56 gallery was chosen as the home for the MilliQan detectors.

The MilliQan Bar detector is positioned in this gallery at an angle of $\sim 43.1^\circ$ and distance of $\sim 33\text{m}$ (17m of which is stone) from the CMS IP. To access the gallery one must enter the CMS warehouse and pass through radiation and retinal security at the PM-54 access lift, taking the large elevator down to the USC-55 Experimental area on level 2. From here, one must ascend 5 flights of stairs on the other side of the shaft, accessible through a fenced platform next to the shaft. After 5 flights there is a narrow caged catwalk leading to a heavy door sealing off the gallery. On the inside of this door one can activate the ceiling lights in the gallery, revealing a tunnel. Following this tunnel one will find self-rescue safety mask kits and a T-intersection. Turning left leads toward the main UXC-55 cavern, and is currently blocked by an interlocking door to ensure radiation safety during beam-on periods. Turning right reveals the MilliQan bar detector, angled down toward the floor so as to align with the CMS IP.

3.3.a CMS Warehouse

The CMS detector is now located fully underground, but it was built in a warehouse on the surface (CMS-S) located in the small agricultural and residential town of Cessy, France.



Figure 3.7: An interior view of CMS-S, as seen from the MilliQan workstation on a 2nd-floor metal platform.

Prior to being lowered, each major component of CMS was built here on the surface before being lowered by crane into the cavern [19]. Though much smaller than CMS, this was much the case for the MilliQan bar detector, which had many of its components assembled and tested on a level-2 platform of CMS-S before being sent to the gallery (albeit via different methods, and down PM-54 instead of the main shaft). CMS-S remains as a fully functional facility for engineering, data-monitoring, and secure access to the underground. A large tarp at the end of the warehouse depicts the front face of the CMS detector at true scale, to commemorate its construction and awe visitors.

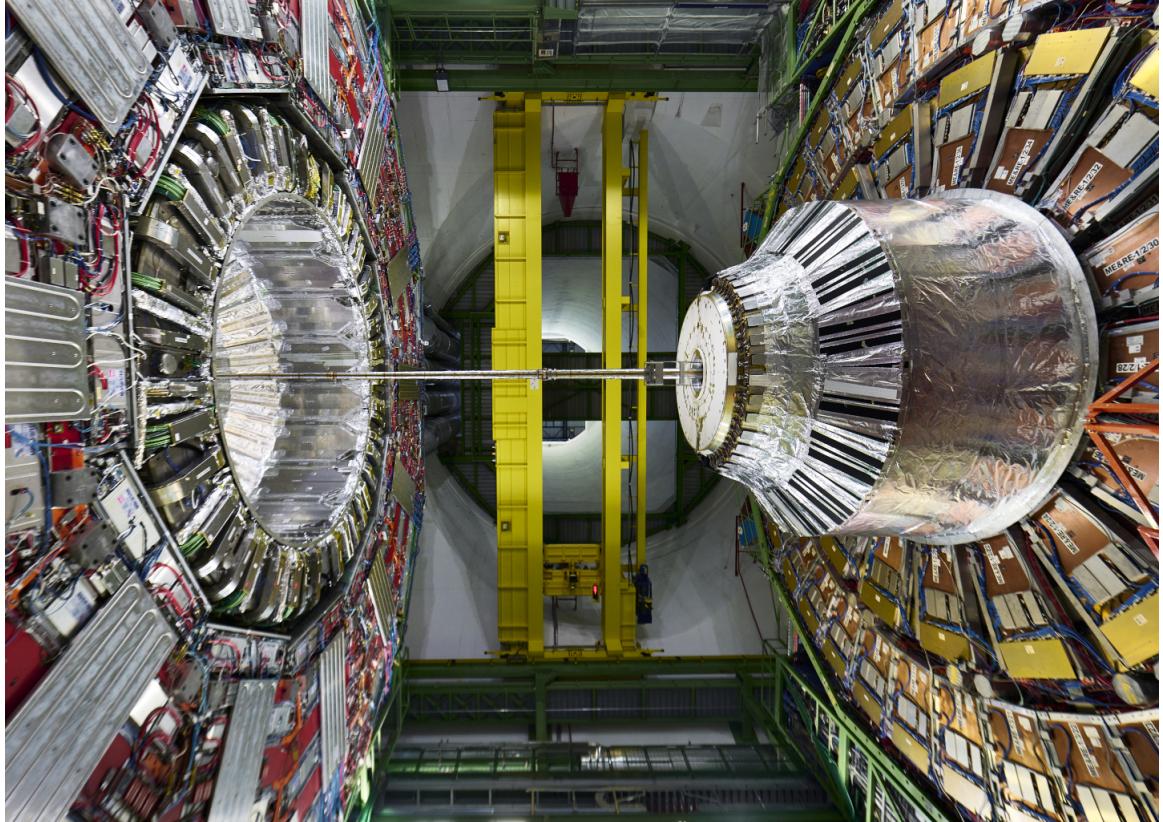


Figure 3.8: The CMS experiment. Photo courtesy of [19].

Chapter 4

The Bar Detector

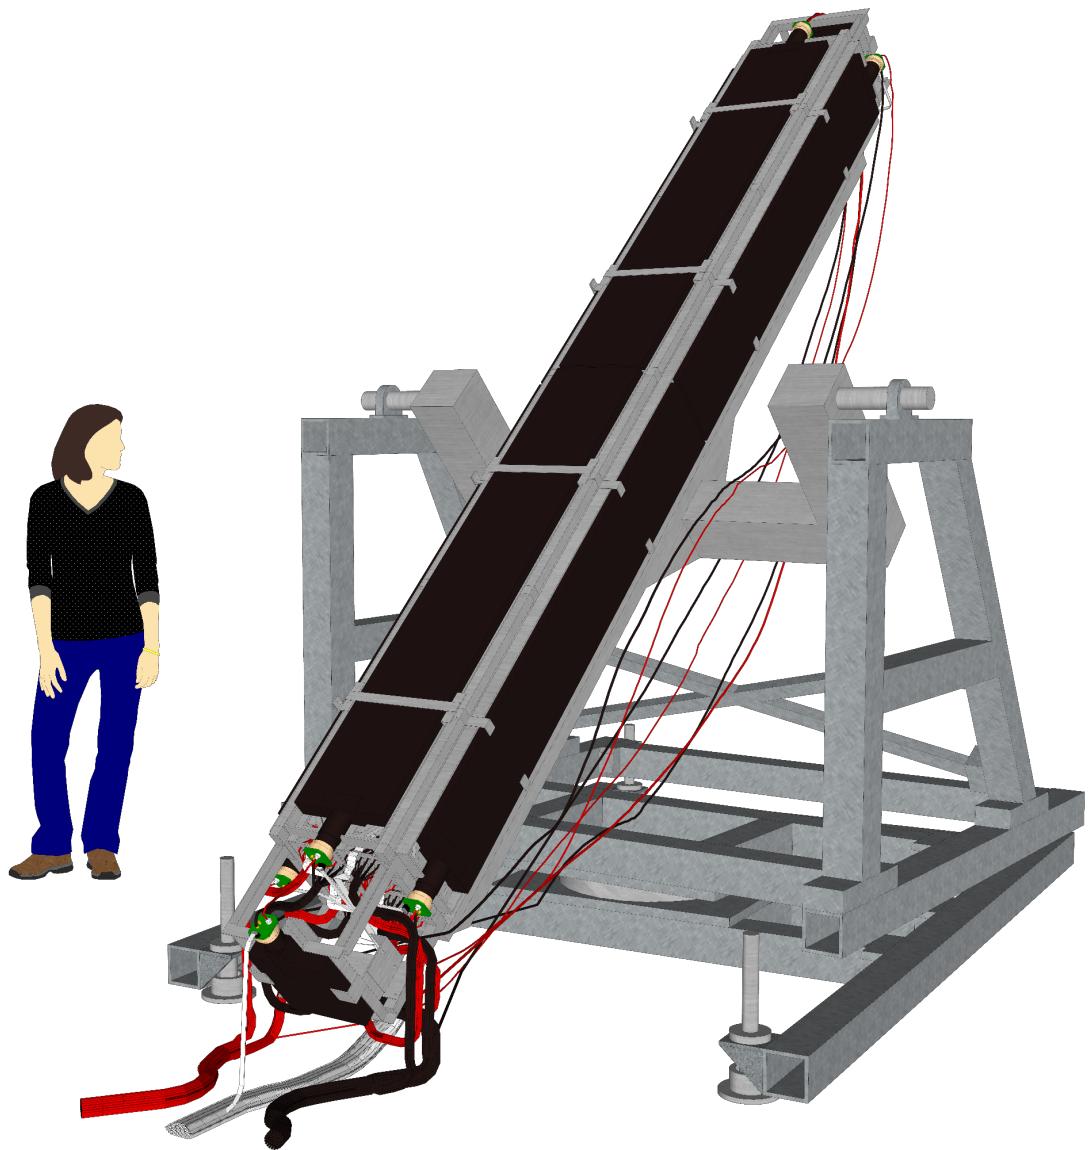


Figure 4.1: A 3D model of the bar detector, designed and rendered using SketchUp.

4.1 Constraints at the LHC

The LHC is unparalleled for data collection rate and mass sensitivity, but like any experimental setting it introduces unique constraints and backgrounds to the MilliQan experiment. The most immediate consequence is the detector's size, as it must fit within the PX-56 Gallery. The gallery is less than 3m in height, meaning that to achieve 4 layers the completed bars had to each be about $\sim 80\text{cm}$ long, accounting for scintillator, PMTs, and all other detector components.

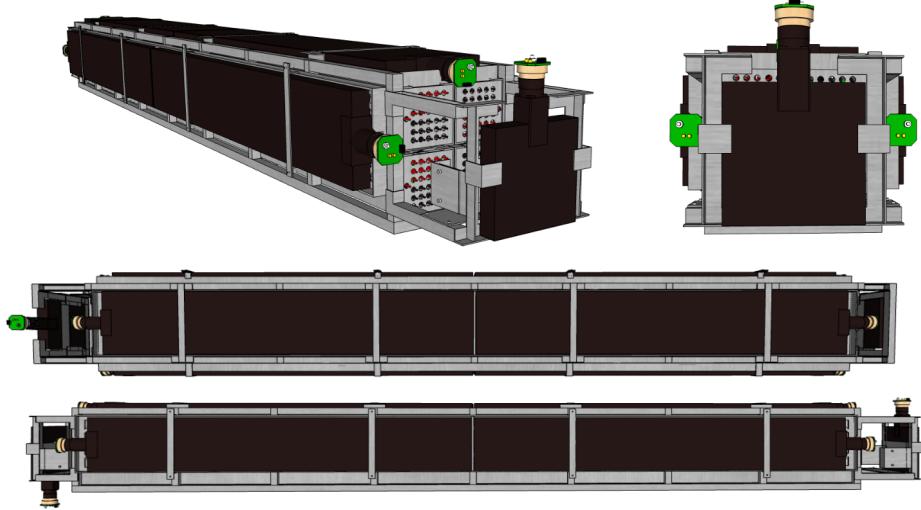


Figure 4.2: More angles of the bar detector, detached from its support and alignment structure. The direction of the PMT mounts on its front and back panels are opposite to allow for it to maximally fill its space in the tunnel, avoiding PMT contact with the floor and ceiling (respectively).

4.2 Signal Criteria

An mCP generated at the CMS IP with its momentum directed toward the MilliQan detector will travel in roughly a straight line, with magnetic field deflection reduced due to the particle's very small charge. Likewise, its small charge will allow it to pass through the 17m of rock separating the main cavern from the gallery unabsorbed. It is expected that it will then pass through the detector, hitting exactly four colinear bars, parallel to the detector's axis. It will leave a small deposit of energy in each of these bars, on the order of a single photoelectron (SPE) in PMTs for $Q/e \sim 3 \cdot 10^{-3}$, with deposits of increased size for signal regions with larger charge. For charges on the lower end of the spectrum, no signal is expected in the front or back panel of the detector, as they are not sufficiently thick in the direction of the IP to be sensitive to scintillation from ionization. To maximally suppress backgrounds, any event not meeting these specific criteria will be vetoed. Specifically, any event with more than one hit in each layer, any hit in the side panels of the detector, hits in each layer that are not coincident in time, events with hits that are too small/big or drastically varying in size, and in low-charge signal regions, events with hits in the front or back panels. Hits that pass all of these criteria will be counted and measured against measurements of background, with a significant deviation suggesting the existence of mCPs.

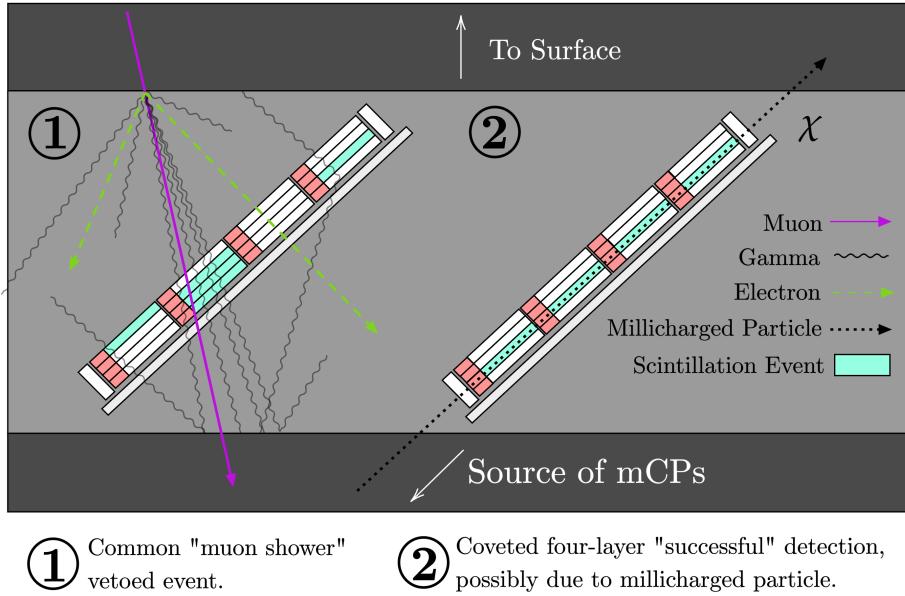


Figure 4.3: A diagram contrasting a common rejected event in which the detector is showered by radiation resulting from a cosmic ray with the ideal signal event which indicates detection of an mCP.

4.3 Backgrounds

There are a number of backgrounds that can produce mCP-like signals, and studying them is necessary to understand what an mCP discovery would look like, and to accurately model the detector's sensitivity.

4.3.a Cosmic Muons

When particles from space (primarily protons) make their way to Earth, they collide with nuclei in the upper atmosphere and produce showers of particles, especially pions. Charged pions then decay into muons and neutrinos. Muons have the same charge as an electron but are not easily absorbed by matter, allowing them to penetrate deep underground and interact with charge-sensitive detector materials like scintillator [14]. It is possible for these muons to pass through the drainage gallery, or for them to collide with the rock near the walls and ceiling of the gallery, producing gamma-rays, electrons, and other radiation that will all potentially interact with the detector and produce signals of a variety of sizes. Because the detector is sensitive to $O(10^{-3}e)$ charged particles, the effect of $O(e)$ charged particles or gamma rays is usually very large. That said, it is possible for low energy gammas or Compton scattering of other particles to produce an SPE-level signal reminiscent of an mCP. It is unlikely for such a signal to be produced in precisely the correct geometry and timing to replicate an mCP without other channels picking up radiation and thus providing veto, but this background does become a concern when it coincides with SPE pulses caused by light leaks or dark rate.

4.3.b Beam Muons

The position of the bar detector was chosen specifically so that most radiation from LHC collisions at P5 is absorbed in the rock surrounding CMS before it can reach the detector, save for feebly-interacting particles like mCPs, but this does not block all radiation. High energy muons produced in collisions, in particular, are able to penetrate the 17m of rock and pass through the detector. Because they will follow roughly the same trajectory as an mCP might (directly away from the IP, save for a curvature in their path due to CMS's powerful magnetic field), they will travel directly through all layers of the detector. Thankfully these muons will deposit so much energy in the detector that there is little risk of them replicating an mCP signal, and they are sufficiently uncommon so as not to affect a significant number of events.

The only risk of a beam muon producing an mCP-like signal could occur if the muon passed perfectly through a gap between bars in the detector, not causing scintillation directly but producing low-energy photons or electrons through scatters in the material that happen to all go in the same direction, hitting each of 4 bars all in a row and producing an SPE-signal in each. This background is nullified by the front and end panels, which fully cover the inter-bar gaps and can be used easily to veto events caused by situations like these: an mCP traveling through their shallow thickness would not produce any hits, but a muon would.

4.3.c Light Leaks

Because MilliQan PMTs must be sensitive to single-photon scintillation events in the bars, it is essential that low-energy visible light is not permitted to enter the bars. Any tear or gap in the protective wrapping surrounding the detector's scintillator introduces a large flood of photons contaminating and eclipsing any signal, while potentially damaging the PMTs at their operational 1450V power. Thankfully light leaks have a very recognizable effect on data, with an increase of SPE trigger rate at several orders of magnitude.

4.3.d Dark Rate

Not all backgrounds are due to the LHC environment. Dark rate, or dark current, is a background inherent to all PMTs that is caused by internal PMT activity and not light from the scintillator. This can be thermionic electron emissions originating from the photocathode, small current leaks between dynodes, radiation passing through the PMT directly, or just plain electronic noise [27]. Dark Rate is a *major* and unpreventable background for the MilliQan experiment, due to the expected infrequency of mCP detections in our target region and the relatively high activity of our chosen PMTs (an older model) at room temperature, though it can usually be distinguished from light leaks or low-energy activity. Nonetheless, significant effort has been dedicated by me and many others to understand and minimize PMT darkrate for MilliQan.

4.4 Bringing MilliQan to Life

A lot of work is required to translate between a design philosophy and an actual working detector, but thankfully the research professors and CERN engineers of the MilliQan

collaboration were hard at work overcoming these hurdles long before I came on board.

4.4.a Engineering

There are a few ingredients to a functional detector that are absent from the design discussed up until this point. The 64 bars and 8 panels of the detector require a support structure to hold them in place, each one of the PMTs needs a power, readout, and amplification system, and each of those systems will require cabling that must reach every component individually without impeding on the design of the detector. Additionally, it must be possible to access or remove each individual component of the detector at a minimal inconvenience and disruption to the rest of the detector.

For this reason the bar detector has been designed to be modular, arranged in tiers by a factor of four. The 64 bars are divided into small groups of four, which are together bundled with a 3d-printed plastic frame to form a *unit*. Four of these units are then secured end-to-end in an aluminum frame and become a *supermodule*. Finally, four of these supermodules (each containing 16 bars) are inserted into each quadrant of a larger aluminum *cage*, which also has attachment points for each of the detector's external panels.

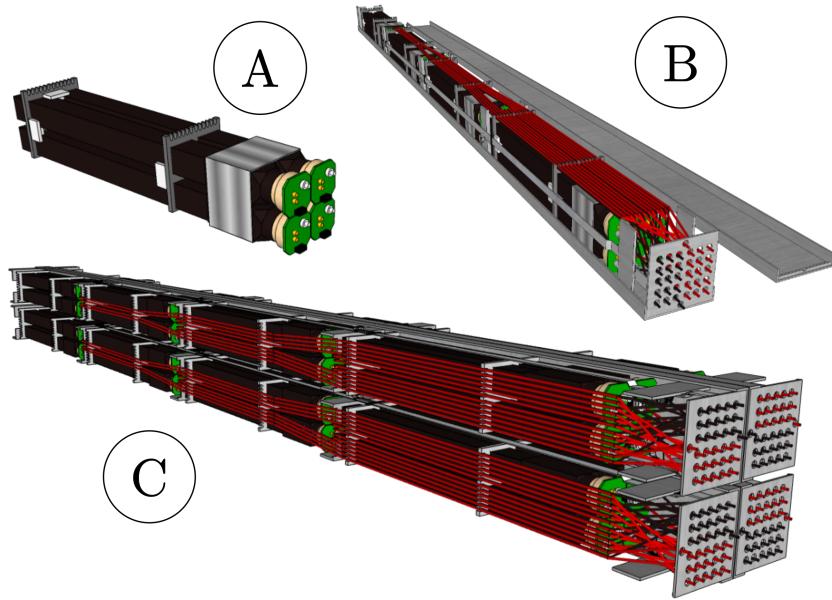


Figure 4.4: Organizational structures of the detector by tiers of fours. A: Units, B: Supermodules, and C: full detector, with supermodule lids and cage excluded to reveal cabling.

4.4.b Transportation and Alignment

Another challenge is getting the detector, which is by no means small or inexpensive, into the PX-56 drainage gallery. Most of the demonstrator had been assembled on the surface and lowered all together by crane down the main shaft to CMS/UXC-55 during a gap in Run 2 collisions, as shown in Figure 4.5, but this was not possible for the bar

detector due to the timing of Run 3. Instead, most of the bar detector components were tested separately on the surface and then transported underground individually.



Figure 4.5: A crane in CMS-S being used to transport a large box containing the demonstrator detector, readying it to be lowered down the main shaft.

Smaller components, such as the cabling elevator and up the staircase by hand. Larger components of the bar detector, such as the support cage and supermodule frames (Chapter 5), were lowered by crane down the PM-54 access lift.



Figure 4.6: The CERN survey team working with MilliQan to align the detector's support structure toward the CMS IP using special CMS-coordinate marking devices.

and detector units, were carried down the elevator and up the staircase by hand. Larger components of the bar detector, such as the support cage and supermodule frames (Chapter 5), were lowered by crane down the PM-54 access lift.



Figure 4.7: A creative experiment in transporting supermodules to the Gallery without a crane: a supermodule lid is balanced diagonally to fit in the radiation scanner required for all materials entering and exiting the CMS underground.

Once on-level with the entrance to the Gallery, a small door in the caged catwalk suspended over the PM-54 shaft was opened to allow these components to be fed through by crane to receivers positioned there on the catwalk. The MilliQan collaboration thanks the CMS Transport Team for performing this sensitive maneuver. The bar detector is tightly secured to a free-rotating support and alignment structure, the same used by the demonstrator. This structure was inserted into the gallery through the large entrance that joins the main CMS cavern and lift, where it had been lowered by crane in preparation for the MilliQan Demonstrator and Run 2 of the LHC. Precision alignment of this structure was performed in 2017 by the CMS Survey team, who used specialized devices to locate the structure in coordinates specific to CMS, and secure it to point directly toward the IP. It now may only rotate freely in one direction in order to access the bar detector, and there are stops in place to easily rotate it back toward the IP without advanced alignment each time. One final re-alignment will likely be performed again by the survey team once the bar detector is ready to take consistent ongoing data.

Chapter 5

Construction



Figure 5.1: (Left to right) Myself, Ryan Schmitz, and Katherine Larina in the PM-54 elevator lobby, waiting for our first trip underground.



Figure 5.2: (Left to right) Mukesh Ghmire, Ryan Schmitz, Dariush Imani, Katherine Larina. A supermodule lays cabled for the first time at the P5 work station.

5.1 Scintillator

Before detailing the construction of the detector that I have discussed thusfar, it is important that we familiarize ourselves with its specific components, reiterating some information in greater detail.

Scintillator is material that produces light in a consistent and convenient spectrum when activated by an energetic or charged particle, either through ionization, Compton scattering, or pair-production. Scintillator comes in a variety of forms, with most being solid,

and ours is EJ-200 plastic scintillator. This scintillator was chosen to optimize cost, and because the light it produces is between 400 and 500 nm, peaking at 425 nm [35]. This is a visible light-purple, and our PMTs are well optimized to detect it. When a particle, such as an X-ray, passes through the scintillator, it deposits a lot of energy by displacing and exciting electrons in the plastic. As the electrons settle back into stable configurations, they emit the desirable visible photons, with more energy leading to more scintillated photons. Our scintillator, like most used in experiment, is wrapped with internally reflective material such that a maximum number of low-angle photons are fully-internally-reflected and funneled toward a collection device at the end called a Photomultiplier Tube.

As discussed in Chapters 2 and 4, the bar detector utilizes four distinct shapes of scintillator. There are sixty-four scintillator bars with the dimensions $60\text{cm} \times 5\text{cm} \times 5\text{cm}$, which make up the bulk of the detector and serve as its signal region. There are four side panels with the dimensions $170\text{cm} \times 25\text{cm} \times 5\text{cm}$, which shield the left and right edges of the detector from radiation and serve as an active veto for cosmic ray showers. There are also 2 top panels with the dimensions $170\text{cm} \times 30\text{cm} \times 5\text{cm}$, which serve the same purpose for the top of the detector, but are slightly wider to fully cover the internal bars of the detector, which are angled toward the IP in such a way that they take up more horizontal space than if they were packed tightly. Finally, there are two end panels, which complete the hermeticity of MilliQan's shielding, provide an active veto for any high energy radiation (particularly muons) originating from LHC collisions, and serve as a low-resolution hodoscope for validating detector alignment using those same beam muons.



Figure 5.3: Bar scintillator, wrapped in plastic. Styrofoam spacers prevent scratching.

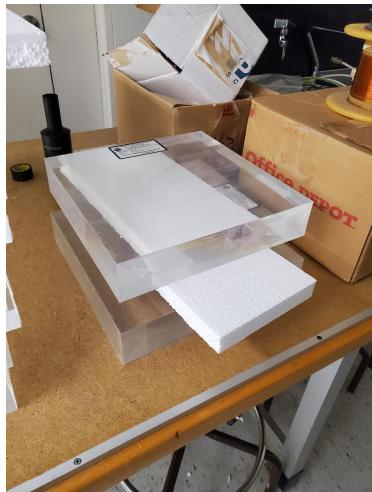


Figure 5.4: Scintillator for the two end panels, also wrapped and separated.



Figure 5.5: Scintillator for the six side panels, the two top panels wider than the others.

5.2 R878 Photomultiplier Tubes

Photomultiplier Tubes, or PMTs, are highly sensitive vacuum-tube-style light detectors, capable of producing a characteristic signal for even a single freed photoelectron. They function through a process of photoelectric cascade initiated by a single photoelectron freed by a low-energy photon.

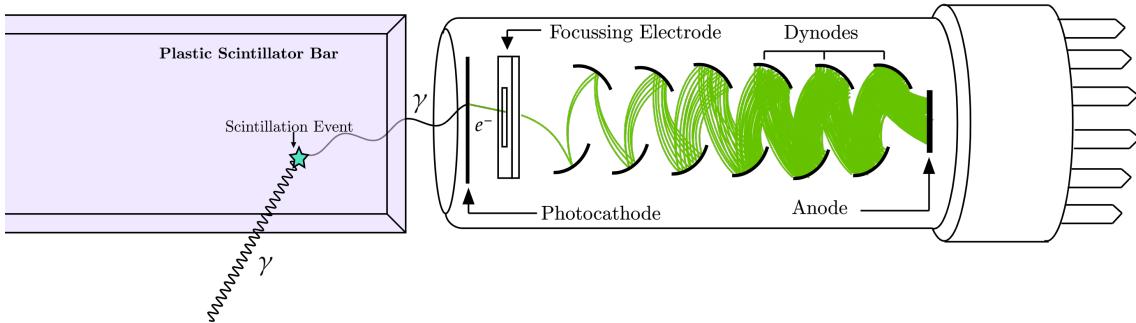


Figure 5.6: Simplified cross-sectional diagram of a PMT. A high-energy photon γ , perhaps an X-ray, interacts with the scintillator and produces a visible (low-energy) scintillated photon.

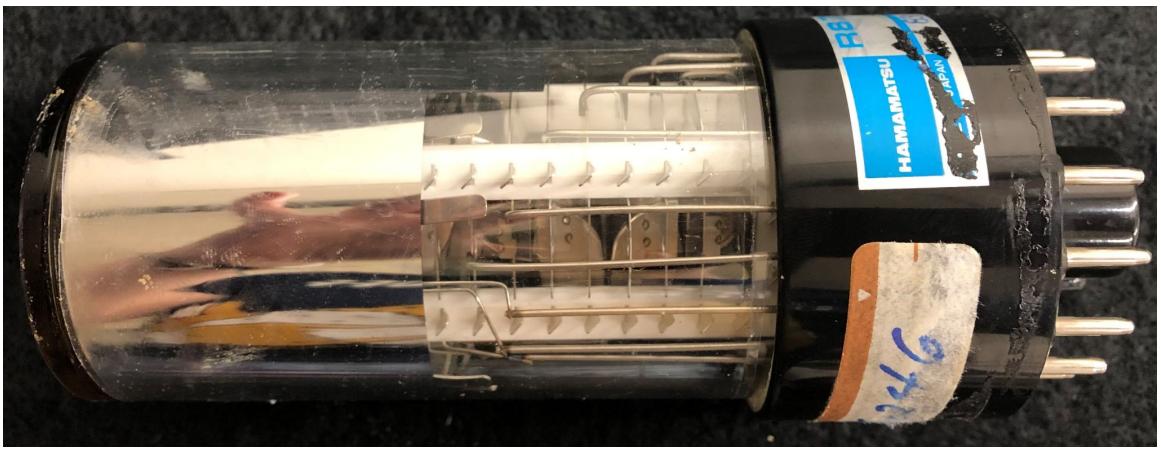


Figure 5.7: One of MilliQan’s many R878 PMTs. In the detector these PMTs are individually wrapped with mu-metal, which blocks magnetic fields, but it is removed here to show the visible dynodes.

A PMT is a glass vacuum tube containing a *photocathode*, a *focussing electrode*, a series of curved metal plates called *dynodes* (in our case 10 [28]), and an *anode* at the end. A PMT is powered by a high-voltage signal (in our case 1450V), which is placed across the photocathode and anode, with each dynode coupled by a parallel resistor to ground such that they act as a kind of *voltage staircase* from one end of the PMT to the other.

When a ~ 425 nm scintillated photon reaches the end of a scintillator bar, it meets the glass interface of a PMT, which is positioned to rest in contact with the scintillator in order to maximize surface area contact (some photons are unavoidably lost, since the square face of the scintillator extends in its corners past the circular interface of the PMT). This photon is absorbed by a photocathode, freeing a photoelectron via the photoelectric effect. This photoelectron is guided by an electric field through the focusing electrode, where it is guided to the first dynode. This photoelectron has been accelerated and impacts with enough energy to free additional photoelectrons, typically multiplying by a factor of ~ 10 at 1450V. These freed electrons are then guided by the potential difference between the first two dynodes such that they slam into the second, freeing even more electrons. This process continues between each dynode, multiplying electrons exponentially until a measurable current arrives at the anode as signal. This signal comes as a short pulse with relatively

consistent properties, but the height and area of this pulse will vary statistically with the number of photoelectrons freed at each dynode. This signal often needs to be amplified before it is read out as data [27].

MilliQan utilized the R878 PMT because it is cost-effective and well matched to our scintillator, with quantum efficiency maximized for photons between 300 and 650 nm, with a peak in quantum efficiency at 420 nm. This minimizes loss, since the scintillator produces light with wavelength peaking at 425 nm as discussed previously.

The R878 was not the only model of PMT under consideration for use in the detector. R878s were used originally in the demonstrator detector because the collaboration already possessed many of them leftover from other experiments, but 4 of the 32 channels instead used Electron Tube 9814B PMTs, and another 2 used Hamamatsu R7725 [3]. These PMTs brought greater timing precision and stronger resistance to background effects, but perhaps surprisingly displayed larger variability in the size of single photo-electron (SPE) pulses produced. The advantage in timing precision was not worth the additional cost and inconvenience, so the R878s were chosen for use in the full bar detector as well.

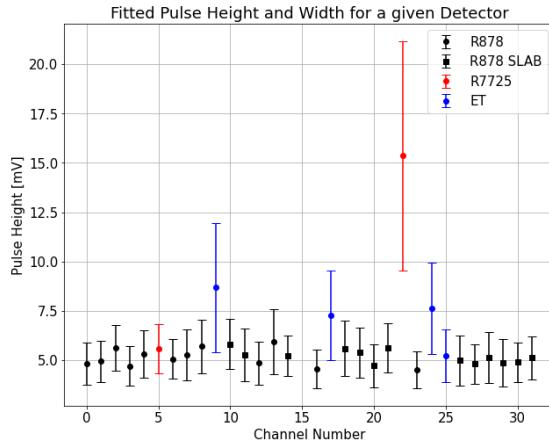


Figure 5.8: The height of SPE pulses in each type of PMT from the demonstrator, with error bars being used to indicate the standard deviation in the pulse height distribution of each PMT, rather than the error. Black points indicate the R878s that would eventually be used, whereas the red and blue were alternatives tested that would have been much more costly.

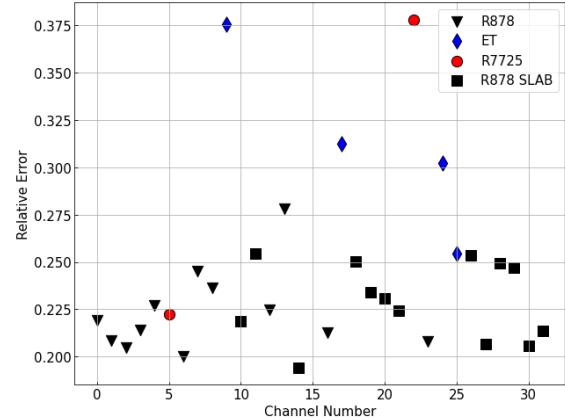


Figure 5.9: A measure of the standard deviation of each pulse height distribution, divided by the center (gain) of that distribution. Gives a better indication of how consistent SPE response of each PMT is. It is more clear here that the R7725 and ET PMTs have more variability in pulse height than the R878s that were finally chosen.

5.3 MilliQan Amplification Base Boards

To operate a PMT, you need a safe and controlled way of supplying high voltage (HV). For MilliQan this means a plastic “base” with a port on one side for a Safe High Voltage (SHV) cable, and 14 holes on the other side, across 10 of which this voltage is divided. The 14 prongs of the PMT (10 linked to dynodes) are inserted into these holes for operation.

You also need a way to output the signal received at the anode of the PMT, and likely a way to amplify and control the gain of this signal. The MilliQan bases receive this output signal through the 11th prong [27][28], and it is amplified by an operationan amplifier (op-amp) on a Printed Circuit Board (PCB) before being output through a “SubMiniature version A” (SMA) signal cable port. There are actually two of these ports, with one at a slightly attenuated gain in order to produce a signal of equal shape but different magnitude for bench testing. The MilliQan bases contain a number of other convenient tricks and tools, including a port for supplying low-voltage pulses to a a Light Emitting Diode (LED), which is an additional feature included in the bars in order to calibrate and trigger on controlled SPE pulses, and will be discussed in greater detail later.

This low voltage is controlled to avoid damaging the LED or PMT, and is configured to both a Digital Potentiometer (DigiPot) and tool-operated physical potentiometer to control the brightness and signal of that LED. There is a switch on the base to alternate between these two potentiometers, with the former being ideal for standard operations and the latter used primarily for bench-testing.

To provide all of these features and a few more, the MilliQan bases require a low-voltage (LV) supply of their own that is distinct from the HV supplied to the PMTs, spanning 10V total. The reason that this must be supplied separately and not divided out from the HV given to the PMT (1450V) is because HV supplies operating in that range can not guarantee precision down to the 10V level, and the components requiring LV could be severely damaged if their supply voltage wavered by a single volt or more. This fact is critical, because it means that now each individual PMT must be supplied with a signal, HV, and LV cable, making 216 cables total.

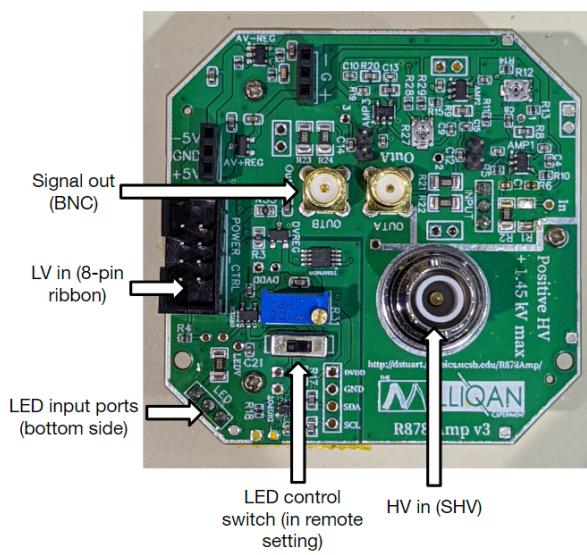


Figure 5.10: A v3 MilliQan amplification base board. Arrows indicate the low-voltage input, the SMA signal output, the LED control switch and ports, and the SHV high-voltage input. These boards are mounted on a plastic base that connects with the end of a PMT.

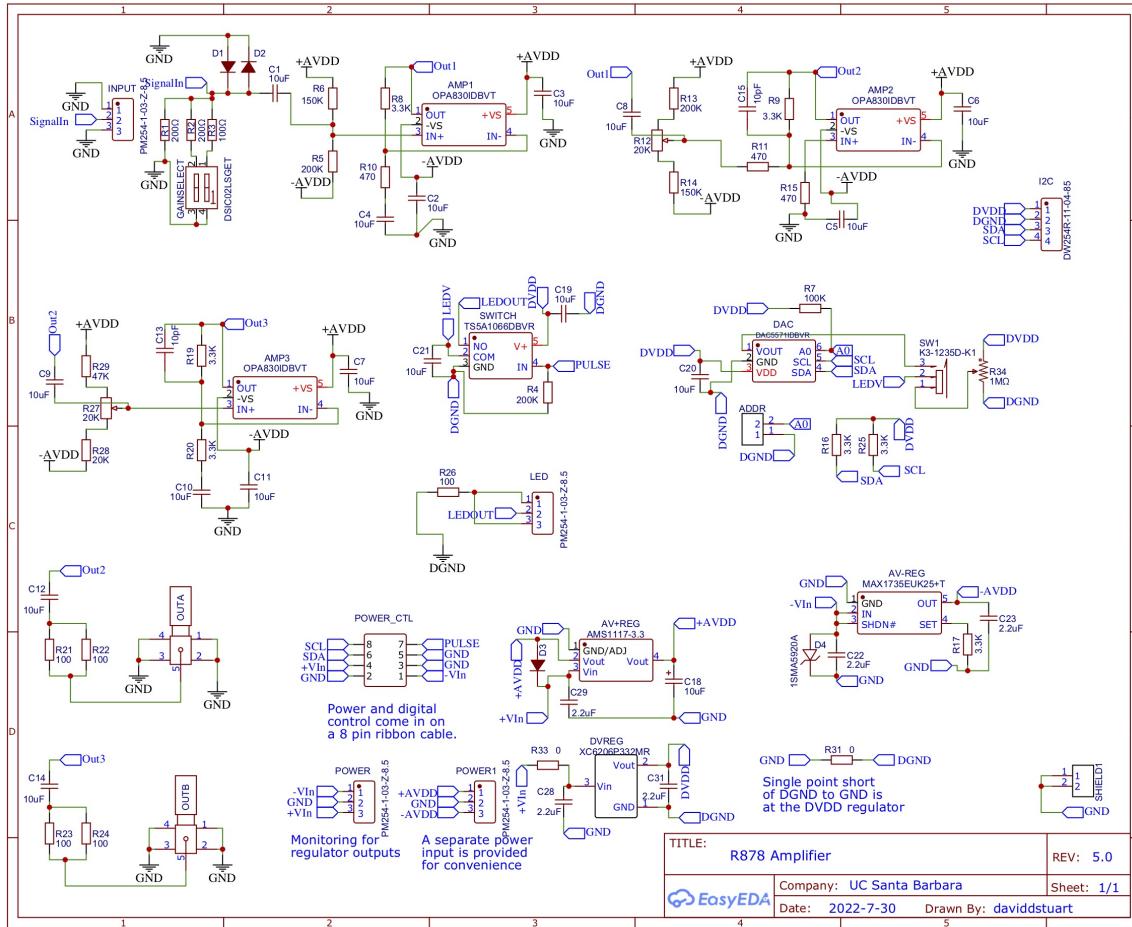


Figure 5.11: A functional circuit diagram for the improved v5 base boards, designed in EasyEDA and implemented by David Stuart. Major addition is the gain control switches in the top left corner.

5.3.a Base Types: v3,v4,v5

The MilliQan bases went through many iterations since the demonstrator, which had no PCB functions. Since construction on the bar detector, there have been 5 distinct versions of amplifier base. The first two were used almost exclusively for early PMT calibration and selection at UCSB, and had lower gain than later iterations. The v3 bases were produced en masse for the first wave of bar detector commissioning at CERN, and were more stable than earlier models.

During this period the v4 and v5 bases were designed by Prof. David Stuart at UCSB, who had created every model thusfar. The v4 bases were never implemented, but v5s became the mainstay and were used to replace all earlier models during the second main thrust of bar detector commissioning in spring of 2023. The v5 bases showed a number of improvements to stability and function, such as increasing the ease of attaching SMA cables to the board, improving the reliability of voltage regulators, and adding a 2-bit hand tune-able pair of gain switches that could change the signal gain of the base, spanning over roughly a factor of 4 between the highest and lowest settings.



Figure 5.12: A photo of the v5 amplification board, complete with the visible addition of gain control switches (red arrow) and other minor changes. These gain control switches allow easy tuning of gain during installation.

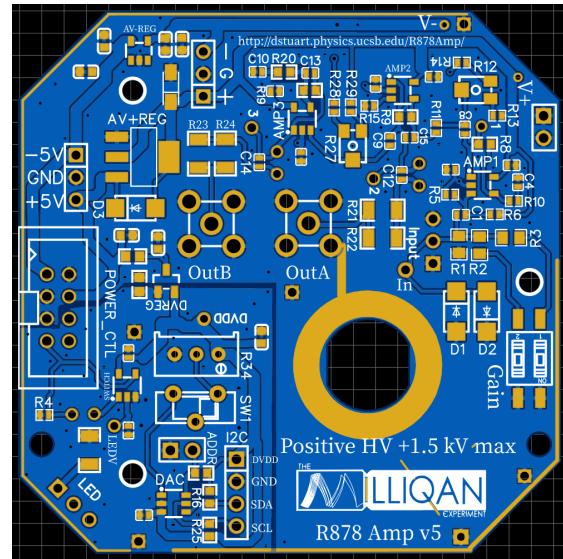


Figure 5.13: A simulated view of the v5 board, designed by David Stuart in EasyEDA. Holes indicate the position of through-hole soldered components added at UCSB by Zachary Zerman after the boards are delivered.

5.3.b Base Testing and Troubleshooting

Issues often arise when working with PMTs, so it is important to be able to isolate whether the issue is a product of the base and not something else. There are a lot of components that *could* be faulty on a base, but the most frequently observed source of high-impact issues were the voltage supply and regulators. If the base is not receiving the correct voltages of +5, 0, and -5, and the regulator is not setting these voltages to the desired values, the signal from that PMT may appear unusual or completely dead. A good way to check this quickly during operation is by using a voltmeter. By probing the four soldered ports indicated on the PCB in Figure 5.14, you should observe both extremes of the supply voltage, and the regulated outputs of +3.3V and -2.3V.



Figure 5.14: Convenient probing locations on the base board and their expected voltages, used to determine whether it is functioning correctly.

5.4 Bars

As has been emphasized, the titular bars of the bar detector are its most critical components. As a result, special care is put into their assembly to assure the highest quality materials available and minimal human error.

5.4.a Bar Wrapping

To prevent light leaks and improve the light collection efficiency of the bars, they are wrapped in a layer of Tyvek (1073D) and black electrical tape, with a square strip of highly reflective 3M DF2000MA specular film covering the end of the bar opposite the PMT.

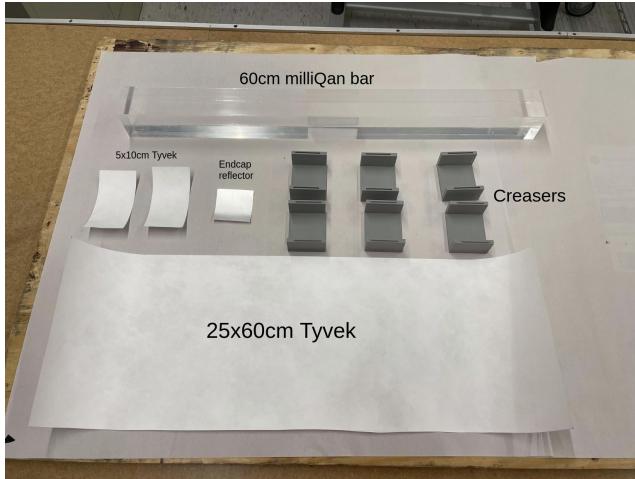


Figure 5.15: Step 1: Materials used to wrap a MilliQan bar. Black tape, gloves, and low-particle wipes not pictured.



Figure 5.16: Katherine busy wrapping a bar with tape during her first summer on the project at UCSB.

The Tyvek has a reflectivity of 0.97 [35], helping to contain additional photons that are not internally reflected at the scintillator's interface. This in combination with the reflective cap turns the bars into light-guides, where scintillated photons produced in the bar are reflected toward the PMT. The process of wrapping a bar follows these steps:

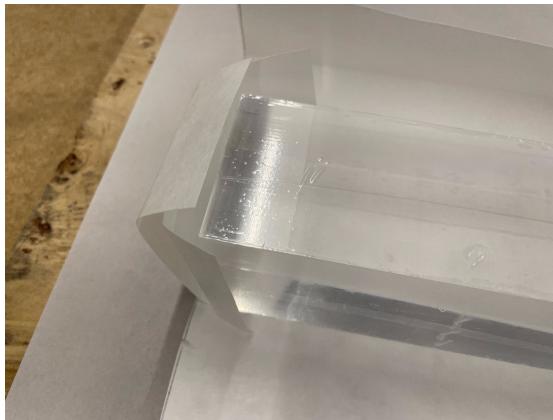


Figure 5.17: Step 1: Desired shape of end-cap.

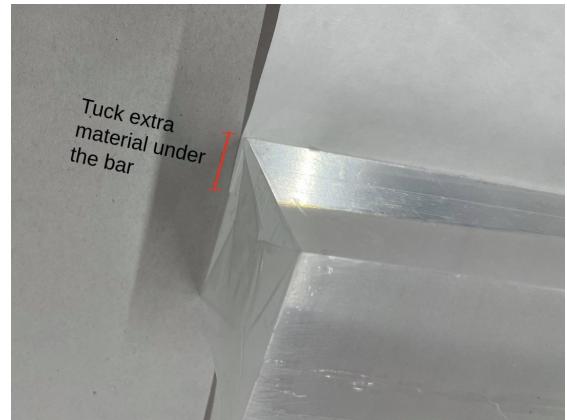


Figure 5.18: Step 2: Tucked scintillator.



Figure 5.19: Step 2: Bar resting on creasers, ready to have plastic unwrapped.



Figure 5.20: Step 3: Creasers are slid to ensure tight first fold of Tyvek over bar.

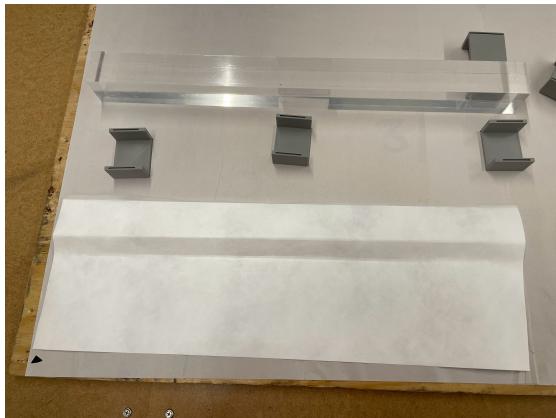


Figure 5.21: Step 4: Tyvek removed, allowing for remaining plastic to be removed from bar.



Figure 5.22: Step 5: Tyvek folded fully around bar, creased at each fold.



Figure 5.23: Step 6: Tyvek strips and square of reflective material, laid in the shape of an end cap.



Figure 5.24: Step 6: Inserting the first strip of the end cap, not yet having tucked in one of the two sides.



Figure 5.25: Step 6: Full end cap inserted and secured.



Figure 5.26: Step 6: Tape used to loosely secure Tyvek wrapping.

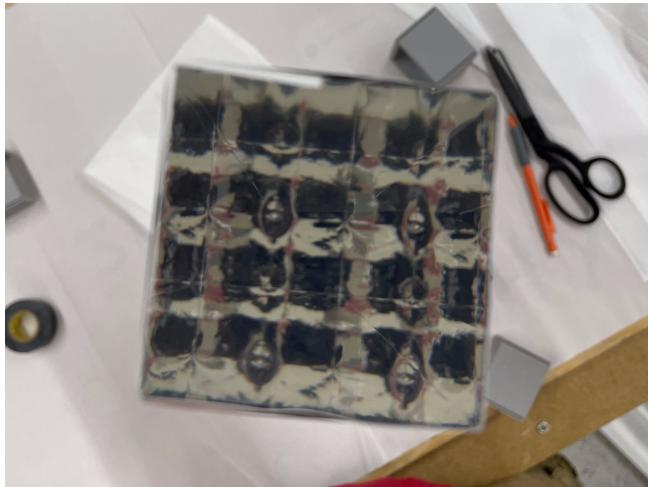


Figure 5.27: Step 6: Kaleidoscopic interior of bar, viewed from exposed end with Tyvek around sides and reflective cap at opposite end.



Figure 5.28: My face through a bar without an end cap, demonstrates its properties as a light-guide.



Figure 5.29: Step 7: End of bar fully wrapped in tape. This is continued over entire bar.



Figure 5.30: Step 8: Bar double-wrapped in tape, with the exposed end covered by cloth.

1. First prepare the necessary materials, as shown in Figure 5.15. Wear nitrile gloves and have sterile *Kipwipes* or another low-dust cloth at hand.

Next, take each of the two $5\text{cm} \times 10\text{cm}$ strips of Tyvek, and use the end of the plastic-wrapped bar to fold their edges such that the strip becomes a $5\text{cm} \times 5\text{cm}$ square of Tyvek with a $2.5\text{cm} \times 5\text{cm}$ flap on either side.

2. Next, tuck a small ($\sim 1\text{cm}$) lengthwise section of the Tyvek underneath the bar, and fold the Tyvek tightly over the edge of the bar so that you form the first crease.

Next, lift the bar from the sheet of Tyvek and slot it into three of the plastic creasers so that it is supported evenly, and so that the crease of its plastic wrapping is above the creasers. With gloves, gently peel this layer of plastic off of the bar, leaving the other half of the plastic wrapping in tact. Alternatively, this unwrapping stage can wait until after the next step.

3. Now, place the creased portion of the Tyvek over the top of the bar scintillator, and use the two remaining creasers to gently form a second crease in the Tyvek. These creasers may need to be lifted over each of the bottom creases in order to avoid contact with the other creaser and bunching of the Tyvek. Adjust all creasers as needed to form a complete and even crease.

4. At this stage, remove the Tyvek from the bar and lift the bar from the creasers. Wearing gloves, pick the bare end of the bar up with a cloth using one hand so that the bar sits vertically on its end. Use your other hand to slowly peel away the rest of the plastic, being sure not to let the bar fall. Once the plastic is fully removed, set the bar back down into the cupped portion of the creased Tyvek sheet.

5. Once the bare scintillator bar is resting in the Tyvek, begin folding from the end with the short tucked-crease so that the bar can be gently rolled without making contact with the table. After each turn, form a new crease in the Tyvek so that it is tightly wrapped around the bar. Once complete, there will be a few cm of overlap between the first crease and the final section of Tyvek. Apply black electrical tape along this overlap, tightly binding the wrapped Tyvek.

6. Turn now to the $5\text{cm} \times 10\text{cm}$ winged strips of Tyvek and the $5\text{cm} \times 5\text{cm}$ square of reflective material. Remove the plastic on the non-reflective side of the reflective square, revealing the adhesive side. Adhere the square to the center of one of the Tyvek strips, forming an *end cap*. Next, remove the plastic film on the *reflective* side of the reflective material, revealing its full reflectivity. Slide this end cap into one end of the wrapped bar so that the Tyvek flaps slide *under* the Tyvek body of the bar, holding the cap securely to the end so that it is fully in contact with the bar.

Next, insert the Tyvek strip without the reflective material over the same end of the bar, that way all four edges of the bar are fully covered. Securely tape the end-cap to the rest of the Tyvek on the bar. You should now be left with A Tyvek-wrapped bar with one end securely covered. When you look down one end of the bar, you can observe its lightguide properties, which will appear as a kaleidoscope pattern. This will allow you to check whether the end cap is aligned properly.

7. Beginning from the end of the bar, use black electrical tape to tightly wrap every surface of the Tyvek, revealing none visible in the end. Pay particular attention to

the edges and corners, as these are the most common locations of light leaks in the bars. The width of the roll of tape is not critical, but for greatest efficiency it is best to use a wide ($>5\text{cm}$) roll.

8. Once a full wrapping has been complete, complete a second wrapping of tape so that the bar is double-layered. Ensure that there are no bumps or protrusions in the tape, as this will interfere with installation later. The bar can be no wider than 5.4cm by the end of the wrapping process. Now the bar wrapping is complete. There should still be one un-covered end with completely bare scintillator. Cover this end with an opaque cloth and secure the cloth so that no light can enter and so that no dust can gather on the scintillator. Store the bar in a dark place, particularly away from natural light sources. This is critical, as UV radiation from the sun can reduce the light yield of the scintillator, reducing the sensitivity of the detector.

5.4.b Plastic Mounting and LEDs

Once a bar is wrapped, it must have a PMT mounted to it. This mounting is achieved by a 3d-printed grip assembled from four parts, with a top and bottom half each assembled from a print with a rectangular notch for the wrapped scintillator and a print with a cylindrical trough for the PMT. These prints slot into one another and are glued securely with epoxy. One half of the mount has a pinhole that leaves an opening between one corner of the scintillator and a cylindrical hole on the exterior of the final mount. This allows the insertion of an LED that can be used for calibration and to provide a random trigger, and which is secured in a light-tight manner using epoxy and black electrical tape.

To construct these mounts, a few basic steps should be followed for consistency. First, 3d print the four parts of the mount: each half of the PMT-cover side (shown in grey in Figure 5.31) and each half of the interface-cover side (shown in white). Half of these parts should have protruding plastic edges and the other half should instead have a hole for them to notch into. One of these bilateral halves will also contain pinholes on their inside for an LED. There will also be notches and holes that allow the PMT-covers to connect with the interface covers.

Once the parts are obtained, eject a small amount of epoxy onto a smooth surface like cardboard and mix using a cotton bud. Then take the cotton bud and paint the mixed epoxy glue onto the curved flat interface of each of the PMT-cover parts. Add a little bit of epoxy to the plastic protrusions as well. On the side with the pinhole, be sure not to apply epoxy over the pinhole or let too much build up in its vicinity.

Quickly but carefully take the interface-cover sides and bind them with their corresponding PMT-cover sides, making sure that excess glue spills over and that there are no gaps: the two pieces should be aligned with no angle between them. Wipe away excess glue and avoid contact with hands.

Keeping each half separate, stand them vertically and insert a pre-prepared LED with soldered and shrink-wrapped extended leads so that the bulb of the LED rests inside the mount and the leads stick out of the top. It is important that these are color coded consistently, i.e. white or red = positive and black = ground. In some cases other color

combinations were used when certain lead colors were not available. Once the LED is inside, fill the opening with epoxy in order to block any holes for light to reach the PMT through (the pinhole goes directly up to the scintillator's interface).



Figure 5.31: Bar mounts glued with epoxy and propped up to dry. The leads of LEDs sprout from the top, positioned so the glue sealing their insertion hole is able to dry.

Secure these parts upright so that they stay together in the correct configuration until the epoxy dries. The epoxy will be fairly solid after about 20 minutes, but leave them for an hour or two for best results.

After they have dried, make sure that they are sturdy and have retained their shape. Do not stress them too hard, as the thinner plastic parts can be fragile and they are not generally constructed to withstand strong bending. If one half has been slightly misaligned at the juncture, the two halves will not be able to mount the PMT and bar as desired.

5.4.c PMT Mounting

To mount a PMT to a scintillator bar, the bar is first carefully lowered into the half of the plastic mount with the LED hole. It is best then to secure this half of the mount to the bar using tape, so that the mount remains rigidly attached when the bar is fully rotated, but being careful not to obstruct the scintillator or any part of the mount needed for joining with the other half of the mount. The PMT is then gently rested in the opposite end of the mount, so that its interface is gently resting in contact with the scintillator. The PMT must be oriented so that its amplifier alignment ridge is pointed upwards, otherwise the LED control ports on the power amplification board will not be aligned with the leads of the LED. This alignment can be checked by attaching a power amplification board to the desired PMT prior to mounting. The other half of the mount is then lowered onto the bottom half, so that the male and female plastic connectors are fully joined and there is no air-gap between halves. Being careful to maintain the physical connection of the PMT interface to the scintillator bar, the entire mount is then secured with tape and wrapped with 1-3 layers, with additional tape usually needed at the regions joining the mount to the bar and to the PMT, where light leaks are more critical and likely to occur.



Figure 5.32: Stages of mounting a PMT. Right: Fully wrapped bar, with exposed scintillator. Middle: Two glued halves of plastic mount, one of them holding a PMT against the bar and the other lying adjacent, ready to be closed over the top. Left: Fully wrapped bar with mount and PMT attached. Leads of LED stick out next to the PMT.

5.4.d LED Testing

There may come a time when the LEDs need to be tested for functionality (or polarity, if a non-standard convention was adopted for cable colors...) *after* they have been fully sealed in a bar, at which point it would be frivolous to unwrap the bar without being certain it was necessary. A good solution in this scenario is to hold a 1-2V voltage difference across the leads of the LED and see whether it draws any power from your voltage supply. Voltage this low will not reverse-bias the LED, and should be sufficient to continuously power it with direct current (DC) in a way that does not leave lasting activity in the scintillator or PMT.

5.4.e Bar-Base Pairing

In non-v5 bases, the gain can not be altered without a soldering iron and surface-mount resistor, and the gain varies between some versions of bases. In the first iteration of bar detector construction, many v2 and v3 bases needed to be used, meaning there were different amplification options available. This would usually be bad, because the size of an SPE pulse from the PMTs should ideally be normalized across the detector, but in this case the PMTs *also* had a decent spread in inherent gain. Because of this, high-gain bases were matched with low-gain bars and vice versa, ensuring that no bar had an SPE pulse height peak below 20mV, and that few had one above 50mV.

This is not an issue now, as all bars have been upgraded with v5 bases, but nonetheless it has been documented which bar corresponds with which PMT and which base. That way if an issue arises with a bar, known behaviors of its PMTs and base can be used to identify the source of the problem.

5.5 Panels

The panels play a less central-role in the bar detector, but they are important nonetheless and there are some critical nuances to their construction.

5.5.a Panel Wrapping



Figure 5.33: Katherine pulling back the protective plastic wrapping from a panel scintillator, Ryan in background.



Figure 5.34: Panels also act as a light-guide, reflecting my face once again. Reflective caps are not necessary, since it is unlikely for light to make it across the panel to the PMT.

Panel wrapping is very similar to bar wrapping, except that the panels are much larger and much heavier. As a result of human error, these features make panels much more prone to light-leaks. In anticipation of this, the wrapping of panels is slightly thicker, with an additional layer of black paper between the reflective Tyvek and the black electrical tape. The wrapping procedure is also more involved, requiring 3-4 people. This procedure is most easily visualized with the smallest panel, but I will explain the distinctions for the larger variants as well:



Figure 5.35: Step 1: End-panel scintillator, still wrapped in plastic. Other materials not pictured.

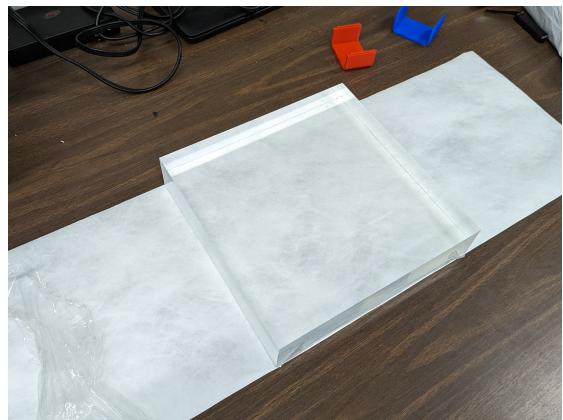


Figure 5.36: Step 2: Unwrapped scintillator placed on appropriately-sized rectangle of Tyvek.



Figure 5.37: Step 2: Tyvek wrapped fully around scintillator.



Figure 5.38: Step 3: Tyvek end caps added and secured with strips of tape.



Figure 5.39: Step 3: another view of panel with completed Tyvek.



Figure 5.40: Step 4: Wrapping process begins again with black paper.



Figure 5.41: Step 4: Black paper being wrapped around Tyvek.

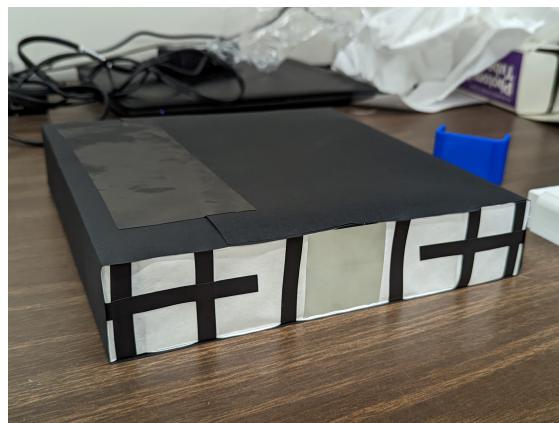


Figure 5.42: Step 4: Black paper secured, ready for next stage.



Figure 5.43: Step 5: Fully wrapped end panel, with scintillator-half of plastic PMT mount attached and wrapped.



Figure 5.44: Step 6: inside of wrapped panel, as seen through square hole left open for PMT mounting.

1. As with the bar, gather your materials. You will need to cut squares of tyvek *and* black paper to the size of the panel that you are wrapping. It should be long enough to perfectly fit the longest edge of the panel, and wide enough to wrap across all faces of the panel with $\sim 5\text{cm}$ of overlap. In the case of the end panels, this means $(25\text{cm}) \times (\sim 70\text{cm})$ rectangles of Tyvek and paper. For the two remaining uncovered faces of the panel, you will require a pair of end-caps. These will be slightly different from those used on the bars: they will be made of 3 parts instead of 2: one larger piece that covers the face, and two short $10\text{cm} \times 5\text{cm}$ strips that fold over the other edges. They are layered with both Tyvek and black paper, with no reflective material. These will cover *both* of the open faces, since the panels have external PMT mounts slid over a $5\text{cm} \times 5\text{cm}$ square at the center of *one of these ends*, and the rest of that side must still be covered. This also means that a $5\text{cm} \times 5\text{cm}$ square hole must be cut into the center of one of the larger endcap rectangles, to provide access for the PMT.
2. Like the bar scintillator, panel scintillator comes tightly sealed with plastic. Unwrap this plastic carefully with gloves and Kimwipes, and lay it flat on the largest section of Tyvek, approximately in the center. Making sure its matching sides are perfectly parallel, fold the Tyvek tightly around the panel using creasers, until it overlaps. Apply black electrical tape to hold the wrapping in place.
3. Once the Tyvek is secured around the bulk of the panel, you will need to attach the end sections. Like with bar wrapping, these should slide snugly under the previous wrapping. The order of adding each strip of the cap does not matter, so long as all edges are covered. Secure these end-caps tightly with strips of tape.
4. Now you will need to repeat the previous steps with black construction paper. In the case of the end panels this is not difficult, but for the heavier side panels this will be a team effort. At least two people will need to lift the larger panels, allowing at least one other person to set down the large section of black paper. Alternatively you can lay down the black paper and Tyvek together at the start, but the larger panels will need to be lifted by multiple people either way.
5. Once all Tyvek and Black paper are applied, it comes time to wrap the entire panel in tape (the wider the roll the better). This tape should be applied in *at the very*

least three layers for the panels, because the detector is more adjustable to their size than it is for the bars, and because the panels are extremely prone to light-leaks due to their size and the difficulty of handling and attaching them gently to the detector. Extra tape should be used around the edges and corners.

Note: in the case of the end panels, do not add multiple layers of tape before attaching its plastic mount (discussed in the next section), as this will interfere with its fit. At this stage in the wrapping process, these mounts should be added and wrapped with tape along with the rest of the end panel.

- At this point you are done wrapping the panel. There should be a $5\text{cm} \times 5\text{cm}$ square of exposed scintillator in the center of one of the short-ends, where the PMT will be mounted. Cover this square with a folded Kimwipe and tape it securely. In the case of the end panels, stuff a Kimwipe into the circular portal and tape it inside.

5.6 Panel Mounting

There are two types of panel mounts, one for the end panels and one for the side panels. The end panel mounts are special 3d-printed parts that come in two halves and can be attached/glued to a PMT-cover part just like was done for the bars. These are used because it is useful to have LED pulsing in the end panels, but not really consequential in the side panels since their primary purpose is shielding. The only difference for the end panels versus the bars is that both halves of the scintillator-end of the mount should be taped to the panel together before the PMT-end is inserted. At a later time, the PMT should be inserted into this hole along with the other two mounted parts, and glue should be applied.

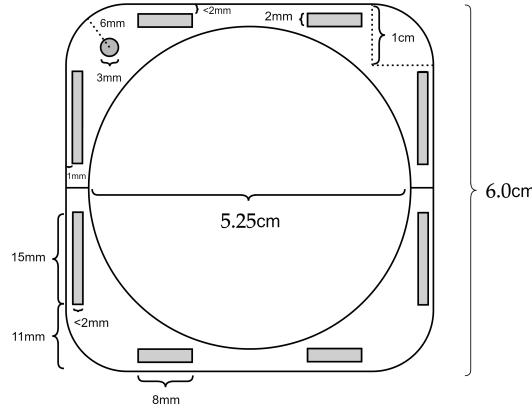


Figure 5.45: Measurements on the front face of the end panel PMT mounts. Rectangular holes are for merging with the other half of a PMT mount, and the small circular hole is for LED light.

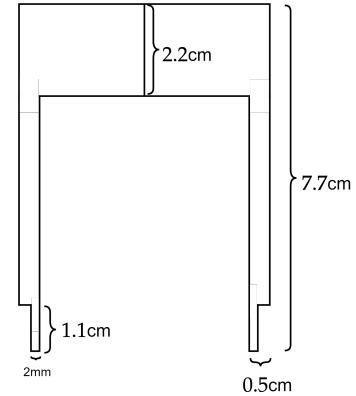


Figure 5.46: Side view of the end panel PMT mount. It is shaped in two halves to close around the end of a 5cm thick scintillator.

The mounts used in the side panels are meanwhile composed of only a single flexible plastic component that is slid easily over the square interface window on the end of a panel. Once

on, these mounts are secured and covered with black tape so as to be light-tight, and a PMT is slid gently into the cylindrical port until it interfaces with the scintillator. Additional tape is then added between the mount and the PMT, completing the hermiticity of the wrapping. These are nice because their simple installation allows for faster PMT repair or replacement in-situ.

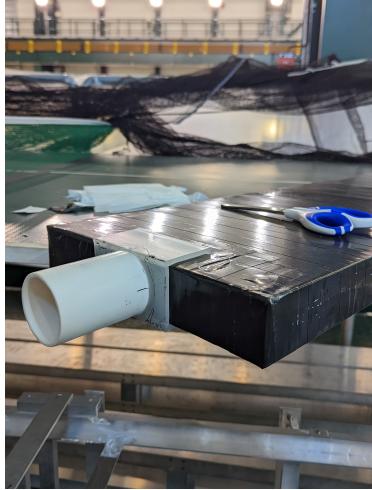


Figure 5.47: Side panel PMT mount, slid easily over the end of a panel in the center over the interface.



Figure 5.48: A PMT can be easily slid into these mounts after they have been attached, fitting snugly.

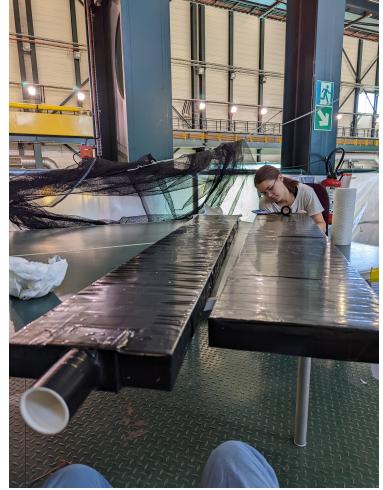


Figure 5.49: As with any other mount, these mounts are fully wrapped with black tape against the scintillator.

5.7 Bar Assembly Units

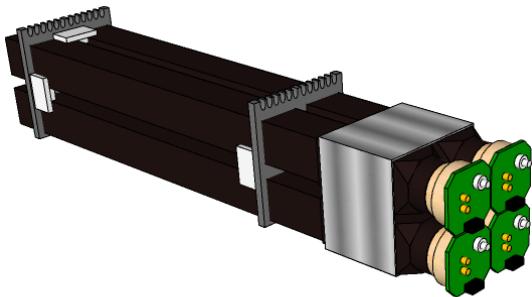


Figure 5.50: Four bars secured tightly into unit frames, held with friction by the frames themselves and by eight white, adjustable clamps. The end of the unit is wrapped in additional mu-metal.

As discussed previously, the bar detector utilizes a rule of fours in its organization. Four bars constitute what is known as one “unit”, four units are contained in one “supermodule”, and four supermodules make up the quadrants of the full detector. Each component in a given stage of this assembly process is self contained, simplifying the transport, organization, and repair of the bars. This modular organizational strategy also encourages different stages of calibration, to be discussed in the next chapter. The rule of 4 does not apply to the panels, which are assembled and attached externally to the detector separately after the other construction is complete.

Units, the lowest level of this system, are assembled using two 3D-printed plastic braces which each contain four square 5×5 cm holes in a 2×2 grid. One side of each brace also has 12 semi-circular notches with a diameter of 5mm. They also feature an extruded section in the shape of a cross at the intersection of the four square holes, and a cylindrical hole running through the center of this cross.

These braces serve two purposes: enforcing even separation and alignment between bars, and providing flush organization for the high-voltage and signal cables needed for each bar in the detector. In assembly, the braces must have metal fixtures attached to them, not completely symmetric between them. Both braces are fitted with a cross-shaped metal component with four threaded holes, one at each end of the cross, and another threaded hole down its central axis. These components are placed on and aligned with the cross-shaped extrusion. For one of the braces this cross is secured using a simple bolt that feeds through the central hole of the brace and is screwed tightly into the cross component on the other end. On the other brace, a longer threaded metal rod is fed through in the same manner, and secures the connection between the brace and the metal cross using a metal nut spun down to hug the brace.

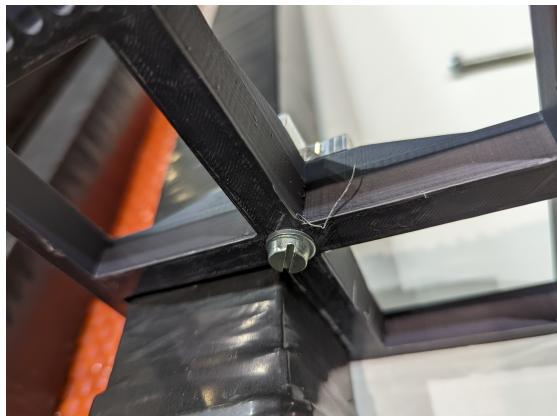


Figure 5.52: The “front” unit frame attached closest to the PMTs, secured with short bolt. One bar is inserted.

Next, two of these braces are slid over a grid of four bars, with care taken to align the bars consistently so that the LV port on their base faces away from the notches. This can be easily ensured by attaching a base to each of the bars prior to insertion. It is critical that the extruded portions of each brace face one another on the inside portion of the completed unit, and that the end with the longer metal rod is on the end of the unit opposite to the PMTs. It is also critical that there is a gap of a specific size between these braces, and that the brace with the longer rod is a specific distance from the end of the bar. This is achieved using an assembly table and metal spacer specifically designed to enforce those



Figure 5.51: A collection of unit braces, newly printed in a CERN workshop by Robert Loos.

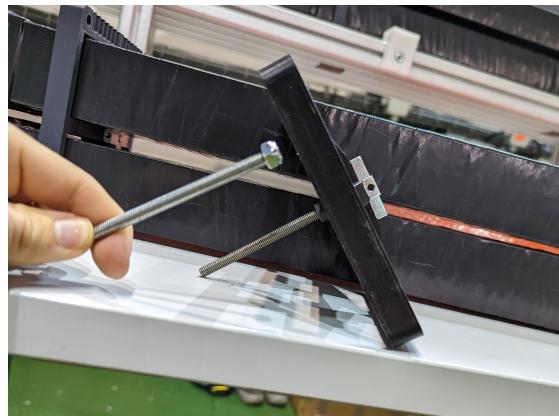


Figure 5.53: The “back” unit frame, with a much longer bolt in the opposite direction used to secure the unit to the supermodule frame.

specific distances (shown supporting unit in Figure 6.1).

Once a unit of bars has its braces, it is necessary to secure the braces so that they do not shift. This is done using eight 3d-printed plastic T-shaped components with holes down their central axis and grooves to allow flush insertion of small metal bolts. Sure enough, each of these plastic parts is combined with a partially-threaded metal bolt, which is then tightened into each of the four threaded holes in the metal cross-shaped components of each brace. This provides significant friction between the brace and the bars, preventing any sliding while not damaging or altering any of the bars. Once this step is complete, the unit has been successfully assembled.

5.7.a Additional Mu-Metal Wrapping

During the commissioning of the bar detector in the summer of 2022, studies were conducted to determine the effect that the CMS B-field had on the MilliQan bars. This effect had been predicted and tested by David Stuart using a Helmholtz coil, determining that the most efficient orientation of the MilliQan bars relative to the CMS magnetic field was with the LV port on the bases facing the Earth and the PMT facing away from the CMS IP, an orientation used to inform construction (interestingly the construction ended up flipping the orientation of the bars, but this 180° did not have a large effect on its performance in the magnetic field vs the planned orientation).

Testing against the true CMS B-field, it was apparent that there was a substantial relationship between PMT orientation and a few of its relevant parameters in the magnetic field (more detail contained in Chapter 8), and that even at the most efficient position there was some potential loss or variability in these parameters. As a result, additional mu-metal was added in a $25 \times 12\text{cm}$ sheath around each unit. The most efficient size and position of this sheath were determined through experiment (more in Chapter 8). These sheaths were cut from a roll of raw mu-metal and electric tape was applied around their edges to increase the safety of assembly and prevent any holes from being torn on the wrapping of the bars.

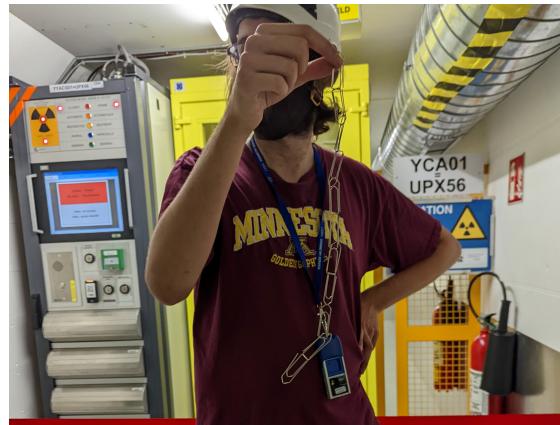


Figure 5.54: Here Ryan demonstrates the strength of the CMS magnetic field in its surrounding tunnels by holding a line of paperclips in the air outside the secure entrance to UPX56. They curve with the direction of the field lines.



Figure 5.55: Two hex wrenches are magnetically held together after being spontaneously magnetized while they were being used to construct a table in the Drainage Gallery.



Figure 5.56: Strips of additional mu-metal lay in a pile with their sharp edges covered by black tape.

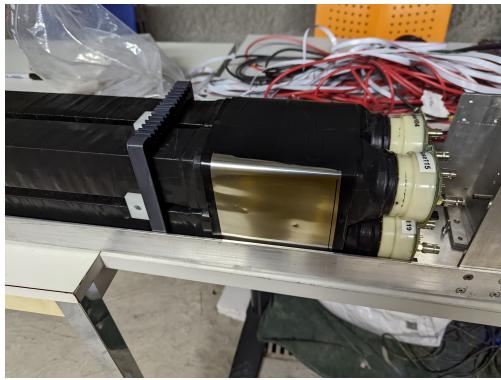


Figure 5.57: One unit is wrapped with a strip of the mu-metal shown in Figure 5.56.

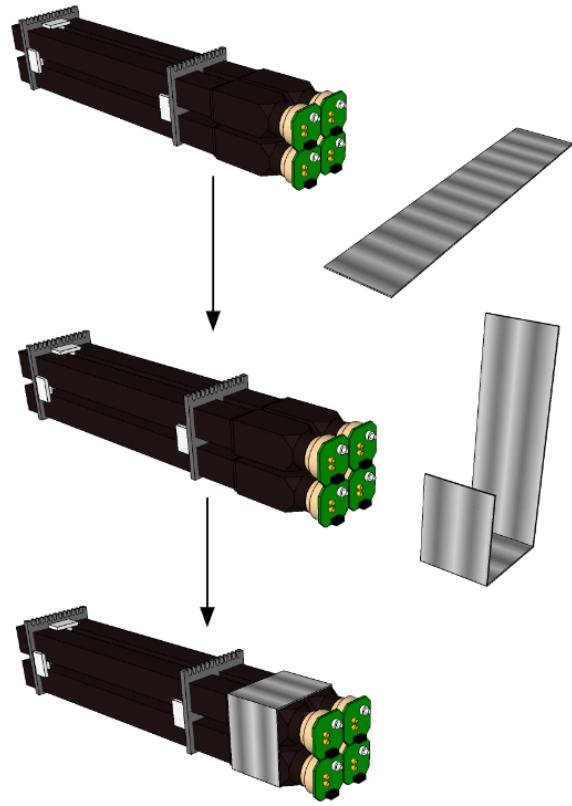


Figure 5.58: The simple wrapping process for additional mu-metal. It should be tight and have some overlap at the edges.

5.8 Supermodules

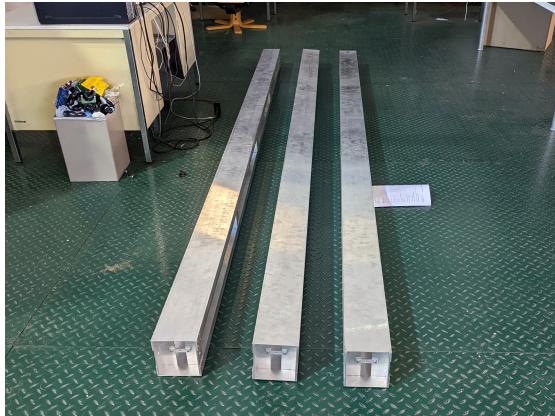


Figure 5.59: Three of four aluminum supermodule frames, made of two practically symmetrical lids and unit spacers connecting them, lie on the floor at the CMS-S work station.



Figure 5.60: A supermodule with all units cabled and inserted lies across two tables in the Gallery. Ribbon cables are routed and lightly taped to the flush sides of the supermodule.

Once units are assembled, they must next be assembled into supermodules. One supermod-

ule contains four units, inserted end-to-end into a metal frame composed of two long metal lids held together by metal brackets secured by screws, and containing four spacers to which the long bolt that protrudes from the back of a unit brace is attached. At the very front of the supermodule, thirty-two holes provide cable egress for the thirty-two 5mm diameter power and signal cables used for those four units. The purpose of the supermodules is to provide additional rigidity and structure to the detector, as well as managing and protecting the cabling of the detector.

The following steps are necessary for assembly:

1. The supermodule is laid flat, and the front is noted thanks to four small holes drilled in diagonals at the far front end of the lid.
2. At the front end, an L-shaped metal component with four matching holes is secured with screws, the position of those screws ensuring only one correct orientation.
3. At this stage the 32-hole grating is secured on the outside end of the front of the panel, using two screws attaching to the L-shaped component.
4. Along the sides of the detector are evenly spaced positions containing four small screw holes. At each of these positions, two L-shaped metal spacers are secured so that both of their inner-corners face toward the front of the detector so as to hold the back ends of a unit's bars.
5. At this stage units are cabled and inserted one by one from the front to the back, and after each is inserted it is secured to its brace using a metal binding component that spans the back of each pair of spacers. This metal component has three holes: one for the long bolt at the end of a unit and two for securing to either spacer. The long bolt is then secured with a metal nut so that the back end of the bars fit snugly into their spacers.
6. Once all units are inserted and cabled, the top lid of the supermodule is lowered over it, with care taken to ensure its front end is aligned with that of the bottom. It must have its own front L-shaped component installed, and after covering it must be secured symmetrically using all of the same screws in all of the same positions.

In these steps I glossed over the process of cabling, but that is extensive in and of itself and will require its own dedicated guide in the next section.

5.9 Cabling

When inserting and routing all of the cables contained in a supermodule, a strict organizational system is obeyed to ensure proper cable management. Each bar takes power through a 5mm diameter SHV power cable, and outputs signal through a 5mm BNC cable. Each base amp board has an SHV port and two SMA ports, meaning that the BNC cable can only be coupled to each board through a BNC to SMA connector attached to each board beforehand. Each board also takes LV power through an 8-pin ribbon cable port, which is fed through a corresponding 8-pin ribbon cable to the LV power supply.

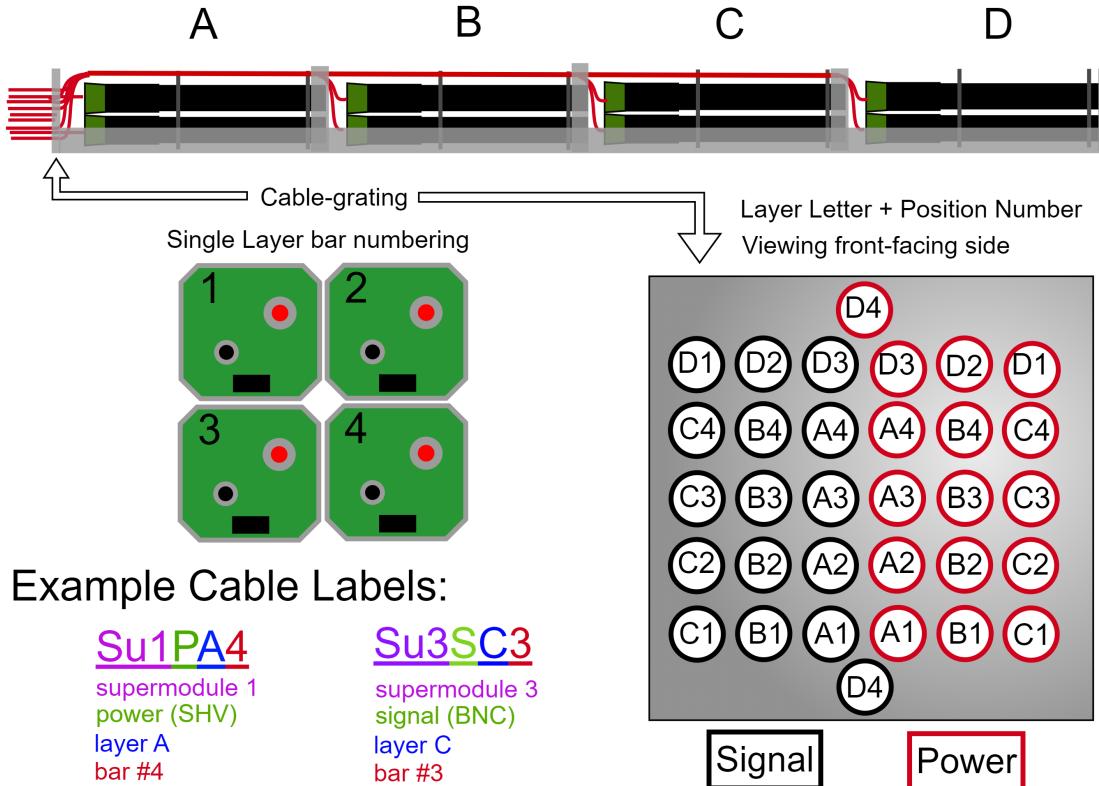


Figure 5.61: A comprehensive labeling guide for the supermodules. This organization system brings redundancy, so that if any cable labels fall off it is still possible to determine the routing of any cable by its positioning.

5.9.a Cable Cutting

Most of the cables for MilliQan were cut by the team at the University of Chicago or on-site by Katherine, myself, and others at CERN. The SHV and BNC cables were cut and attached with their respective connectors in what I can only assume is the standard procedure, but there is an exception to that rule with the ribbon cables. Ribbon cables (in this case 8-pin) have one colored wire to signify the orientation of the cable, as internally the voltages carried by each of its wires are asymmetrical. In order to receive the same voltages at one end as you supplied at the other, you make sure that the colored cable is in the same orientation relative to the ribbon connector at each end. This is not the case for MilliQan.

MilliQan's LV cables are all powered by a large PCB called a power distribution board, and all of the connectors on this board are inverted (due to a layout issue). As a result, all ribbon cables must have their orientation flipped between each end, as shown in Figure 5.62. Correct voltages may be verified with a potentiometer.



Figure 5.62: Ribbon cables in MilliQan are inverted from end-to-end, due to an inversion in the pin mapping on the power distribution board.

5.9.b Cable Organization

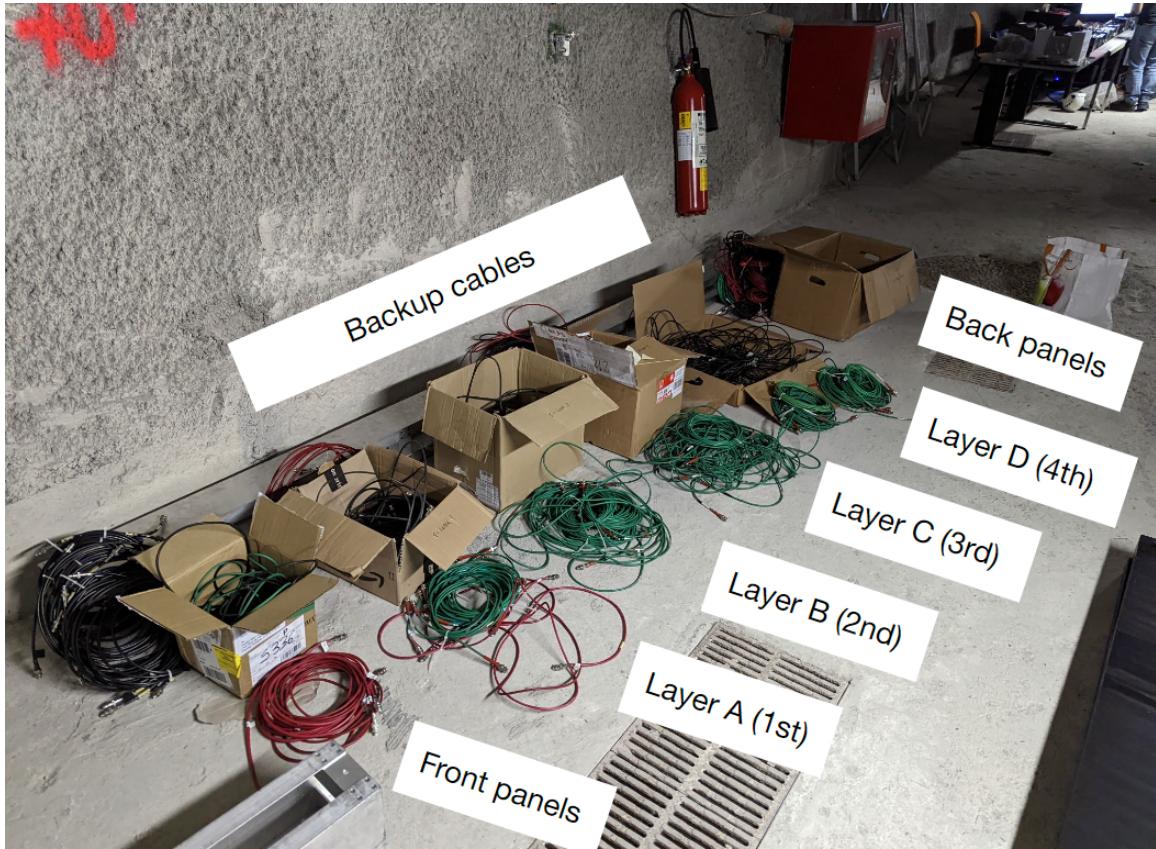


Figure 5.63: Cables were cut and labeled by length by the team at the University of Chicago, though length measurements had to be approximate for the detector, with enough leeway for rotation and repositioning. As a result, cables were organized by their appropriate position in the detector while supermodules were being cabled.

The core array of bars in the bar detector requires 128 5mm diameter cables (64 SHV and 64 BNC), as well as 64 ribbon cables. It is essential that these cables are organized, protected, and that they do not conflict with the desired geometry of the detector. Every 5mm cable is named and labeled according to which supermodule its bar belongs to, what type of cable it is, which unit in each supermodule (or which “layer”) its bar belongs to, and the position of its bar within that layer. The specific labeling criteria are shown in Figure 5.61.

The goal in routing cables for the supermodule is minimizing overlap, and keeping things consistent and symmetrical so that it is easy to tell at a glance which cable routes to which bar. These 5mm cables are run from each base on each PMT, up through the gap between each unit in the supermodule, and then through the plastic notches of all of the units in front of it. The SHV and BNC cable for a given PMT are run together through the same notches, and stay together until they reach the front of the supermodule. The notches that a given one of these cable pairs is routed through is based on their position within a unit, and the unit’s position within the supermodule. The four cable pairs in one unit are always fed through the inner four notches of the unit in front of it, with the bottom two bars feeding into the inner two notches and the outer two bars feeding into opposite outer notches beside them. When those cables reach the following unit, they are routed

diagonally and symmetrically to further notches, allowing space for the cables of the next unit to be routed in the same manner. The result looks like this:

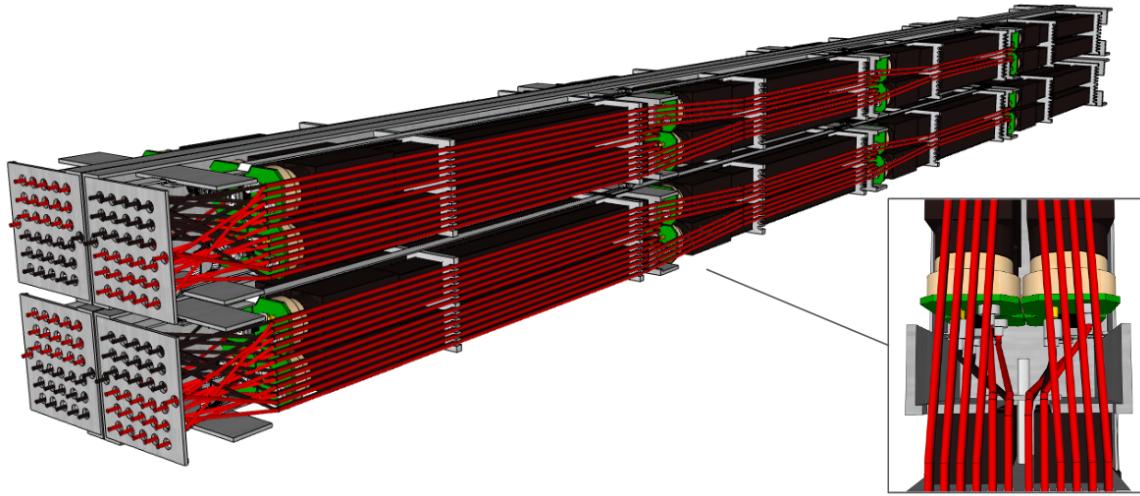


Figure 5.64: The supermodules in the detector are rotated such that the cables are all on the *outside* of the detector, minimizing any gaps between bars. The cables themselves are organized for easy insertion and visual clarity across the detector.

In the case of the front-most bars, their cables are routed directly out of the supermodule and require no notches. Once all of the cables reach the front of the detector, the SHV are split from the BNC and fed through a specific one of 32 holes assigned based on the position of its bar, easily identified from its label. Signal and power cables are separated to allow organized routing between the detector and their separate power/readout destinations, storing the inevitable messy criss-crossing of power/signal cable separation in the extra contained space at the front of the supermodule as opposed to the exterior of the detector.

5.9.c Cabling Process

The process of cabling a full supermodule involves the following steps:

1. Bars with tested, functional LEDs are matched with appropriate, tested bases such that gain differences are compensated for. The wires of the bar's pulsing LED are then inserted into the appropriate ports on the base, with the positive-voltage LED lead being inserted into the center port and the GND wire going into either of the other ports.
2. Additional mu-metal is added according to the process in Section 5.7.a if it has not been already.
3. The first, front-most unit is propped-up on the front of the supermodule so that all of its bases are accessible, and then it has all of its appropriately-lengthed-and-labeled cables inserted into those bases.

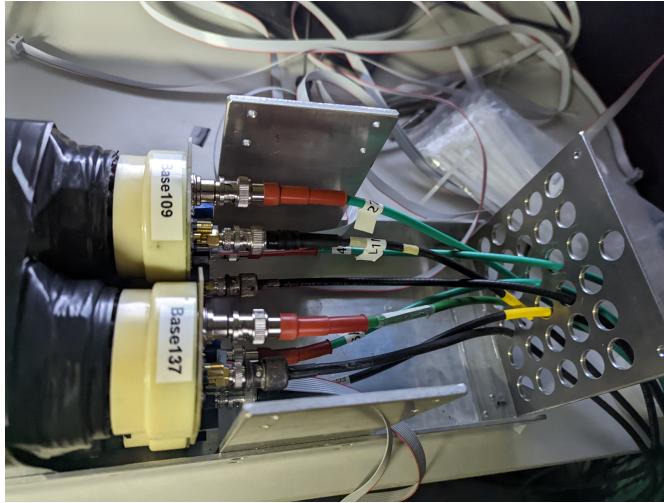


Figure 5.65: At the front of the detector there is additional space for the many cables to be routed to each of their appropriate holes, as shown in Figure 5.61.

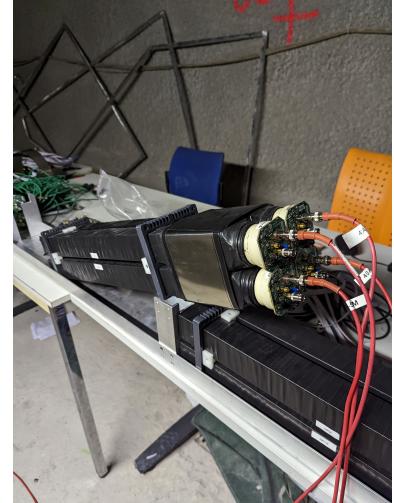


Figure 5.66: A unit is propped in order to have its SHV cables tested.

4. The HV and LV cables of this first supermodule are briefly inserted into their power supplies to ensure fit and functionality, then they are removed. Each ribbon cable should draw $12\text{mA}(\pm < 1)$ of current. SHV current may vary, meanwhile, so differences do not warrant concern within an order of magnitude. Once disconnected, run these cables through the appropriate port in Figure 5.61.
5. Before sliding the unit into position, the small aluminum plate used for securing each unit to its spacer bracket must be removed if it has not already, or it will obstruct clean installation. The unit is then lowered and pushed gently into position, ensuring that the cables are not tugged so as to pull the bases out of position. Once the unit is in, the plate is returned and secured with a screw.

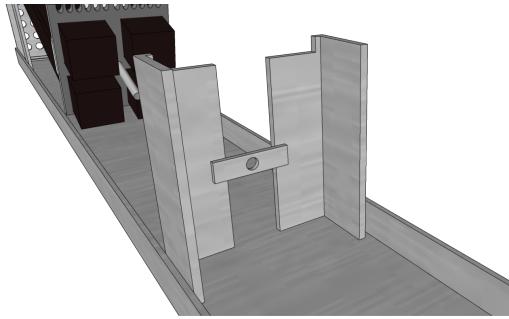


Figure 5.67: A small aluminum rectangle with a hole in its center is used to secure each unit against its spacers. The hole on the inside is not threaded, and is meant to hold the rear bolt of the unit secured after insertion using a threaded nut.

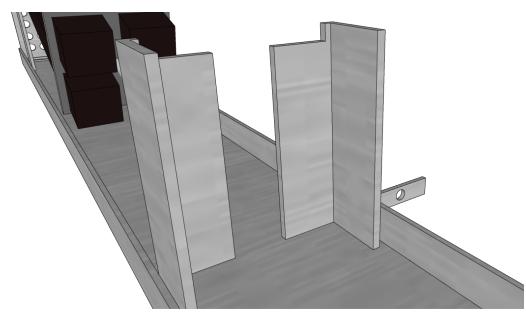


Figure 5.68: This spacer must be removed for the unit to be gently lowered into its position. Afterward it is slipped over the bolt at the back of the unit, and secured to the frame with screws. The unit is secured with a nut.

6. Repeat this process with each unit until all four are inserted. Route cables according to the rules established in Figure 5.61. The LV cables should be connected to each base and lain cleanly on the table at either side of the supermodule.



Figure 5.69: David Vannerom monitors the HV power draw of each cable while units are being inserted.



Figure 5.70: David Vannerom checks each unit in a fully cabled supermodule to ensure no LEDs or bases have come loose.

7. In each section of the supermodule, bundle cables tightly and bound them securely with zip ties to avoid them splaying apart or overlapping one another.

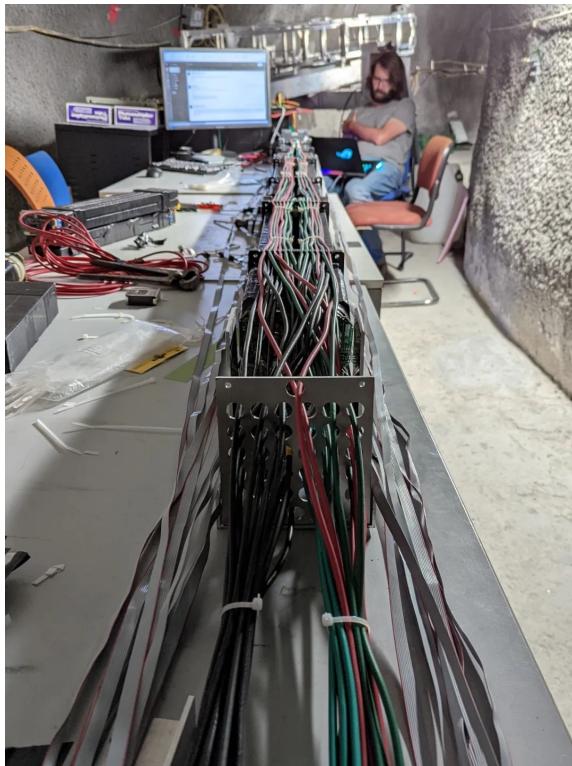


Figure 5.71: Before ribbon cables are secured to the side of the detector, they sprawl on either side and are easily pulled out by accident.



Figure 5.72: LV cables are properly taped to the sides of this supermodule, holding them in place.



Figure 5.73: Each SHV cable is labeled with the code according to Figure 5.61.

8. Once every unit is in and every cable is tested and routed, and you are confident that no LED has slipped out of its port, it is time to put the lid on. Be wary that the lid takes 40 screws, which are very time-consuming to remove if anything needs to be changed. Before screwing the lid on fully, run all of the ribbon cable so that

it lies flush with the sides of the detector and does not extrude anywhere. If any ribbon cable extends past the aluminum frame, it will be caught on something during insertion and may cause ribbon cable to be pulled out of a bar. Many of these steps would need to be un-done and repeated if this occurred. In order to keep the ribbon cable flush, duct tape was used on the first two supermodules.

Some notes on closing the lid: the aluminum is quite malleable and not all of the holes on the lid will align with the holes on the body perfectly at first. If a hole is so misaligned that a screw can not be inserted, find the closest alternative and insert a screw there instead. This will usually move the first hole closer to alignment. Expect this process to be difficult and frustrating.

9. Take the loose LV ribbon cables and run them along the exposed edges of the supermodule, along the exterior of the spacers but flush with the aluminum lids. Secure them tightly to the edges using black tape or plastic sheets. The goal here is to prevent these ribbon cables from being tugged out of their base during the supermodule insertion process, as their connection to the base is much less secure than the SHV or BNC cables.
10. Zip-tie the ribbon cables at the front end of the detector, keeping them separate from the power and signal cables.

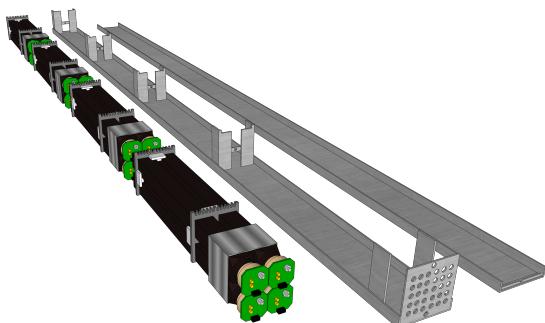


Figure 5.74: Step 1&2: bars are prepared for insertion

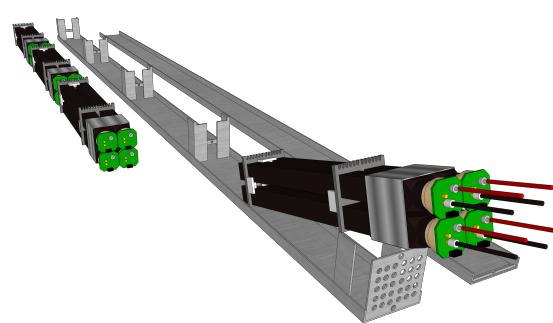


Figure 5.75: Step 3&4: First unit cabled and tested.

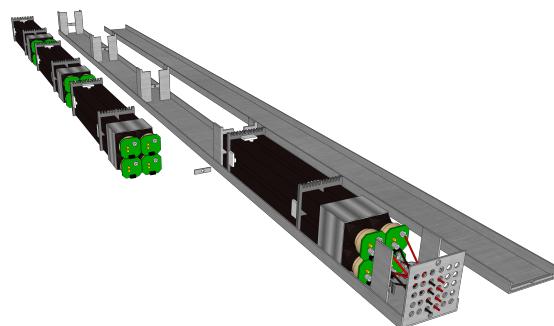


Figure 5.76: Step 5: First unit inserted and secured.

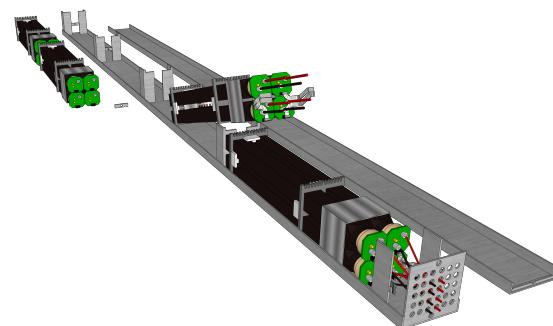


Figure 5.77: Step 6: Process repeated for second unit.

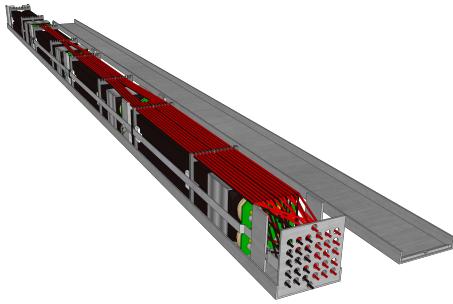


Figure 5.78: Step 6&7: All units cabled and labeled.

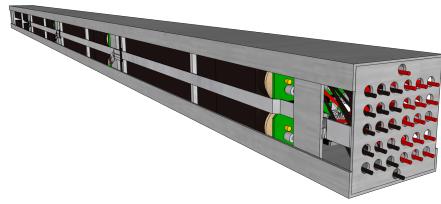


Figure 5.79: Steps 8&9: lid covered and LV taped. Ready for insertion

5.10 Power

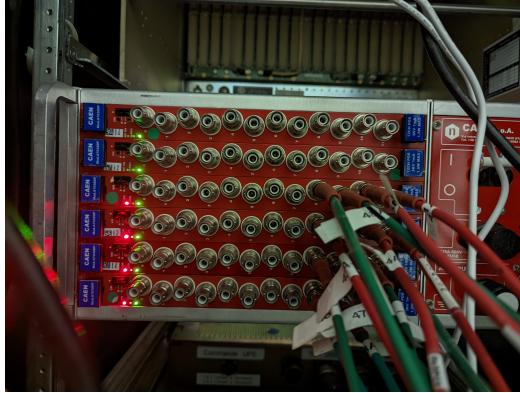


Figure 5.80: An image of the CAEN HV power supply, with ports for Su3 in use.

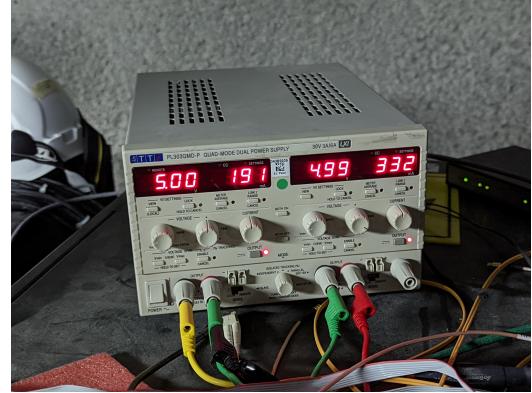


Figure 5.81: The LV supply uses two channels to provide a range of -5 to 5 V.

5.10.a High Voltage Supply

HV to the detector is supplied by six Mod.A7030DP CAEN boards stacked in a control rack that takes power from the wall. Each board has 12 SHV ports, making 72 total. This makes it just barely sufficient for operation of the bar detector, and there is a standard convention used to map each port to each channel:

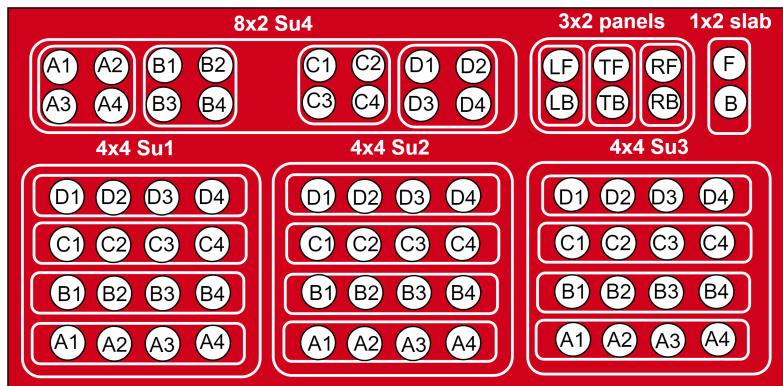


Figure 5.82: Until splitter boxes are installed, the bar detector occupies every channel in the HV crate, organized for safety and convenience.

This works for now, but it is not the long-term plan for detector power. In fact, it is planned that the slab detector will share this CAEN supply with the bar detector. This creates an obvious issue: the slab detector will require 196 HV channels, and there are exactly 0 available! The solution to this is being worked out by the team at the University of Chicago, who are designing and building large splitter boxes to fan out HV from the CAEN supply to the bar detector and to each layer of the slab detector.

5.10.b Low Voltage Supply

One might expect that LV would be easier to supply than HV, but at least in the case of MilliQan they would be incorrect. The supply itself is a PL303QMD-P Quad-mode dual power supply from Thurlby Thandar Instruments, which is nothing unusual. In order to supply ground, -5V, and +5V, both channels are used, with banana cables establishing the negative port of one channel and the positive port of the other channel as being GND.

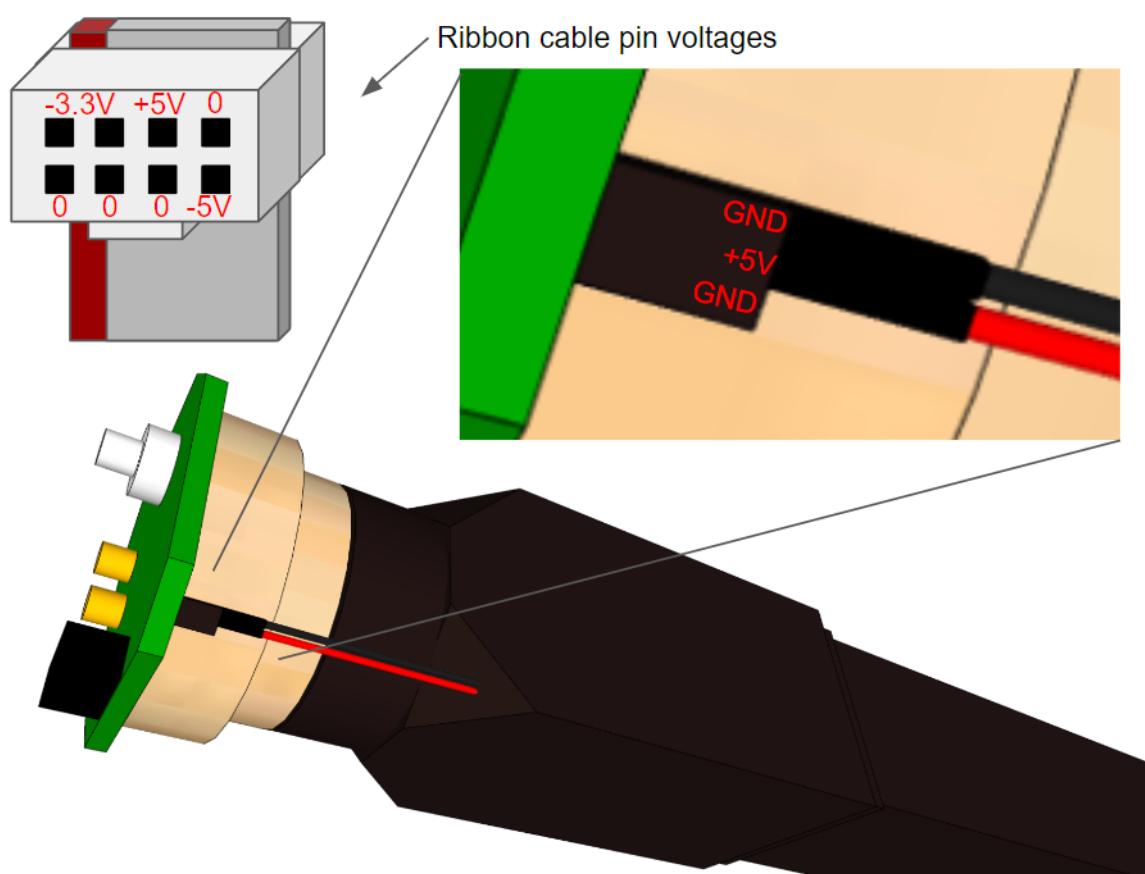


Figure 5.83: Four unique voltages run through variations of the 8 wires in each ribbon cable. Some power the base, some control the LEDs.



Figure 5.84: An LV power distribution board, with an RPi used for controlling the pulses of LEDs in each bar.

Where things get complicated is that this supply provides DC voltage to the *Power Distribution Board*, which fans that power out to 80 channels and converts the three inputs to the 8 needed for each ribbon cable and port. The ribbon cables themselves carry two lines for I2C communication, one port at +5V, another at -5V, and four of them at GND. These serve the purposes of providing power to the base boards, which as we have seen control the amplification of PMT signal and pulsing of the LED. The LED leads themselves are inserted into the base, of course, into the appropriate ports as shown in Figure 5.83.

The LV is challenging to work with because there are many things that can go wrong at each stage of the process, and MilliQan uses primarily custom equipment to achieve this.

There is another source of stress when it comes to LV: it controls the LED sitting inside every bar. If this LED was to accidentally be powered for a long period of time, it could effectively destroy a PMT and require that a bar be removed and un-mounted. Thankfully there are protections built into the bases and power distribution board to prevent this, but it exemplifies the perhaps surprising risk of low voltage in an experiment that uses such high voltage overall.

5.11 Supermodule Cage

After the supermodules is the next stage of construction scale: the supermodule cage. The supermodule cage securely supports and angles each of the four supermodules. This is achieved using internal metal struts that cross and separate each supermodule into its own square, a bolting system at the back for the supermodules to be secured to, and plastic wedges customized to each supermodule in order to angle it ever so slightly from the straight-pointing position, thus increasing the total cross-sectional area of the detector relative to the CMS-IP and eliminating to first-order the effect of mCPs and muons passing through the detector in radial instead of parallel lines.

On the outer faces of the cage there are connectors for each type of panel. Each end panel is mounted so that its PMT interferes with the rotation of the detector minimally: the back end panel is attached so that its PMT is pointing toward the ground (not hitting the ceiling), and the front end panel has its PMT facing upward (not hitting the floor). As will be detailed further in the next section, the front section used for end-panel attachment must



Figure 5.85: Robert Loos (left) is seen through gaps in the interior of the MilliQan cage as it is attached to the support structure for the first time.

be removed to insert supermodules, and this is fairly easy to do. There is also additional space at the front of the detector between the end-panel and the face of the supermodules: this is used to route the 216 cables from each supermodule out of the detector and to their respective supplies/outputs.

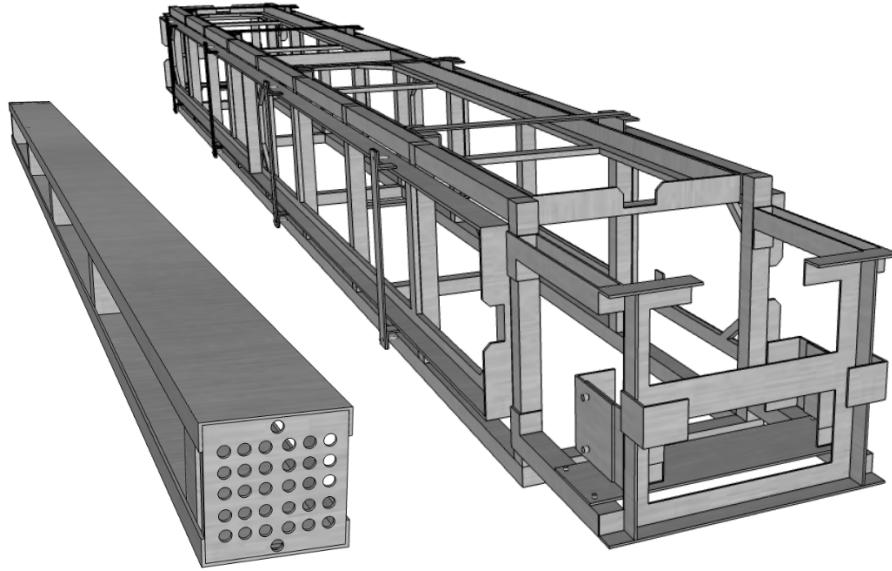


Figure 5.86: A 3D model of the aluminum cage and one of the four supermodules (left) inserted inside it.

The supermodule cage was designed by Mouhammad Kamra and Robert Loos, two engineers who helped extensively in the practical execution of MilliQan. The cage went through a few iterations, but ultimately the rectangular model with sliding supermodule insertion was most convenient to construct and operate.



Figure 5.87: The interior of the MilliQan cage after its construction, without attachment points for panels yet added to the outside.



Figure 5.88: The full supermodule with all panel attachment ports, spacers, and parts used for angling the inserted supermodules.

5.12 Supermodule Insertion and Alignment

The supermodules are very heavy, and at this point they contain a large amount of fragile and expensive equipment. For this reason supermodule insertion is a 3+ person

task, and one that follows a specific procedure for a given completed supermodule:

1. Prep the detector for the insertion of the finished supermodule. First note the power draw in the LV channels, and write it down. If these values change, it indicates that a bar has been connected/disconnected during insertion. Next turn off HV power and carefully raise the detector to its horizontal position. This will require the removal of some lead blocks from the support frame and the alignment of a support bar on the front end to hold it up. Lastly, remove the necessary side panels and front end panels, as well as the front-end panel aluminum mounts.
2. Before inserting the supermodule, remember the two small white wedges used for alignment: they will need to be removed from that supermodule's slot in the cage. These wedges secure the supermodule with friction and align it to the IP. Be very careful not to mix them up when setting them down, as they are not identical.
After inserting the supermodule, these wedges will be wedged back into their prior position, and screwed back in.
3. Recall that the orientation of the supermodule for cabling is different from its orientation in the detector. Depending on whether you are dealing with the left or right half of the detector, the supermodules will need to be rotated 90 degrees counter-clockwise or clockwise (looking up the axis of the supermodule from the front). In the final orientation, all of the routed cables will be stored on the outside left and right of the perimeter of the detector, in order to minimize gaps in the center. Rotating the supermodule must be done very carefully by at least 3 people, with at least one person holding the bundle of cables so that none of them catch or tangle.
4. Now it is finally time to slide in the supermodule. One person should be holding the bundle of cables at the end of the supermodule, another should be supporting the back end, and a third should be supporting the front/middle and feeding it in while it is pushed. Ideally this would be four people, with two people supporting the supermodule's full weight while the extra person moves and make sure it slides in smoothly. Once the supermodules are in, there is a plate and some large bolts to bolt it in place.
5. Once the supermodule is in, all of its (labeled) cables will need to be routed to their respective ports on the LV distribution board and the CAEN SHV supply. Until the large HV splitter boxes arrive from UChicago, the cables for the bar detector are being routed as shown in Figure 5.61, following the same cable naming conventions as previously. Be careful to familiarize yourself with the channel naming conventions of the HV power supply itself, as they are unintuitive.
6. Now all that is left is to return the panels and cage parts that were previously removed, leaving the detector in the original state it was but now with a new supermodule. The detector may be carefully lowered back to 43.1 degrees, and the power may be switched on.

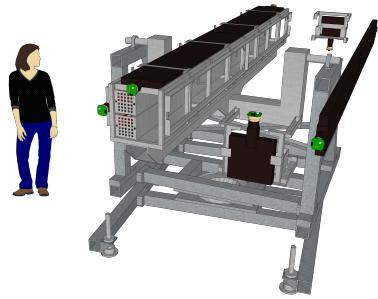


Figure 5.89: Step 1: End and some side panels (in this case right) are removed for insertion.

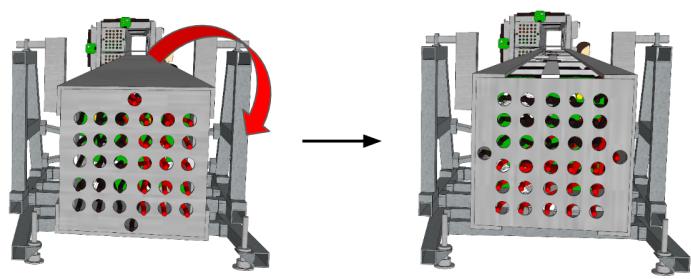


Figure 5.90: Step 3: This right-sided supermodule is rotated before insertion. Left-sided supermodules are rotated the other way.

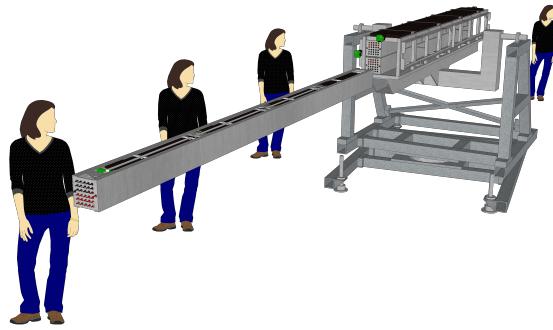


Figure 5.91: Step 4: A spotter and team of 3 prepare the supermodule for insertion.

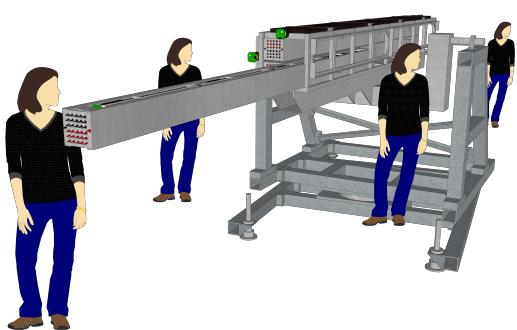


Figure 5.92: Step 4: The supermodule is slid into the cage.

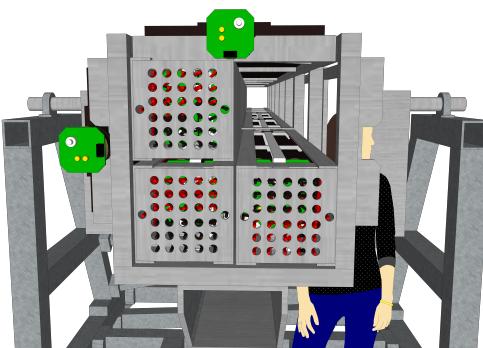


Figure 5.93: Step 6: Panels are attached again, and supermodule orientation verified.

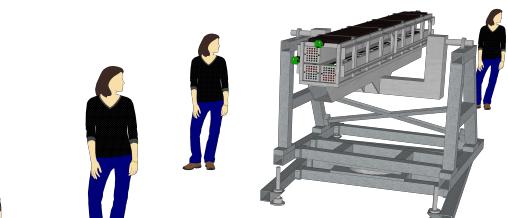


Figure 5.94: Step 6: The supermodule is fully inserted.



Figure 5.95: (Left-right) Katherine, Keegan Downham, and Hualin Mei stand ready as Rob secures the supermodule to the cage. Ryan supervises.

Note on power: the CAEN HV supply displays currents for each of its channels, so it is easy to check that they are all working. The LV power supply only displays global current, however, so some math will need to be done. Each v3-v5 base consumes 12mA of current, with a select few v2s in the past consuming 23-24mA and a few unusual ones drawing much more. The positive channel (right) of the LV power supply will always run at higher current than the negative channel (left) because portions of the board are powered. At 0 ribbon cables the negative channel draws 0mA, so use that one to measure.

5.12.a Installing Panels

Installing the panels is much more straightforward than installing a supermodule. Each panel has a box frame that they can be slotted into, and then secured using removable metal bars tightened with screws to hold them in place. There are three things to be mindful of when inserting the panels, however:

1. The edges of the cage are very sharp, and the panels fit very snugly, so avoid tears. Light leaks in the panels are not disastrous, but they will pollute the entire detector's data.
2. The slots in the side of the cage are larger than the panels, but the panels themselves are meant to be fitted as close as they can be to the center of the detector, that way they cover the bars maximally with their shielding.
3. The top-back panel is very close to the ceiling, with its cables just barely fitting. It was so close, in fact, that a portion of the ceiling needed to be lightly chipped away for it to fit.



Figure 5.96: All four supermodules inserted in the detector, visible due to removed panels.

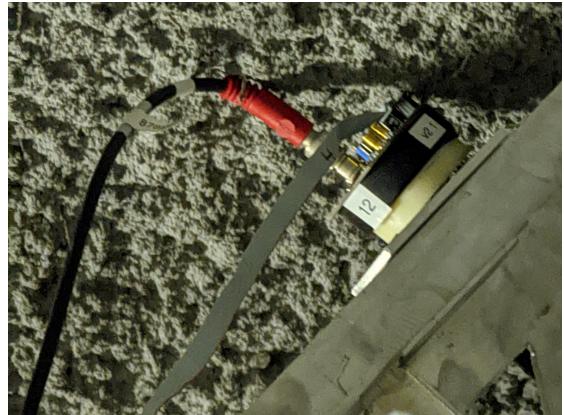


Figure 5.97: The backmost top panel brushes the ceiling of the Gallery when aligned.

5.13 Hodoscope

As mentioned in Chapter 3, there is a plan for MilliQan-b to use hodoscopes, but these have not yet been constructed. The current plan is to use a bundle of square scintillator strips in two paired layers that cross perpendicularly, that way high-energy beam muons leave a signal in the hodoscope that gives much more precise x-y data in an event. This will allow for better modeling, simulation validation, and background veto for beam muons.

5.14 Voila

Following these steps, you would have completed the construction of a bar detector!

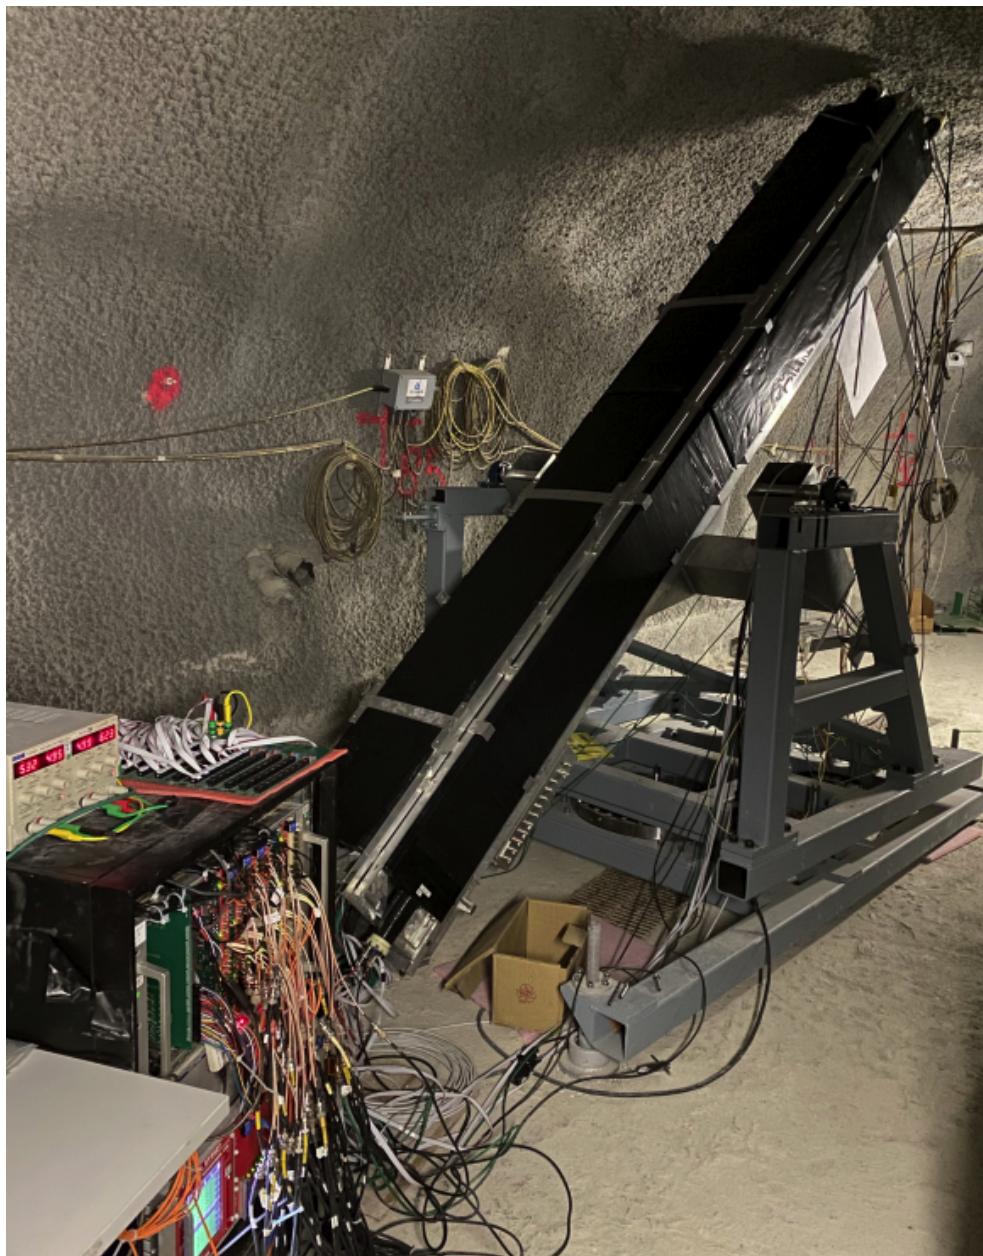


Figure 5.98: The fully assembled detector is powered and in position, taking data.

Chapter 6

Calibration

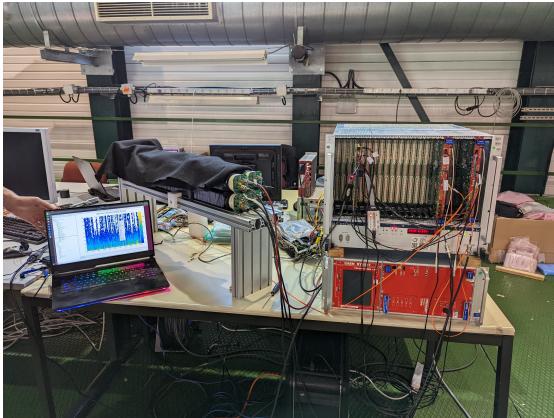


Figure 6.1: A newly-constructed four-bar unit is tested at CMS-S, draped in black cloth to cover any light-leaks created during assembly.



Figure 6.2: (Left to right) Katherine, myself, and Ryan look at data from the first ever full supermodule test, at a table in CMS-S.

In commissioning the bar detector it was not sufficient to merely assemble the relevant materials, but also to perform extensive and intensive calibration on them. This meant testing and monitoring the components—particularly the PMTs, to which the detector owes most of its sensitivity, to make sure that their behavior was consistent, well documented, and optimized for sensitivity to mCPs. To understand that task, it is important to highlight exactly what qualities are desired in a PMT.

In Chapter 4 I discussed an unfortunate feature of PMTs: dark rate. It is no surprise then that MilliQan seeks to minimize dark rate in its PMTs. There is another issue, light leaks, that we considered in Chapter 5. This is distinct from dark rate, but interferes with data in a similar way, so I will consider these together as making up the dark rate of a bar.

In the previous chapter we also considered the intrinsic gain of a PMT when discussing the MilliQan base boards. This gain varies appreciably between our R878 PMTs, and because gain factors are multiplied, it is absolutely consequential to the overall gain observed in a bar. This gain needs to be known and consistent so that thresholds can be set in the detector for what size of pulse corresponds to an SPE. In this sense there is not an effort to minimize or maximize gain, but rather to constrain it to a convenient range.

Lastly, and perhaps most critically, is what I will call the energy response. When a particle passes through scintillator and deposits energy, that energy is converted to scintillated photons, which are in turn converted into photoelectrons, which are multiplied into an electronic signal. If we want to know how much energy a particle deposited (a question critical to understanding the physics of an observation), we need to know the approximate conversion rates involved. And if we wish to produce at least one photoelectron (and thus any signal at all) for as small an energy deposit as possible (such as from an mCP with very low charge), we want to maximize this overall conversion rate. Thankfully there are methods and tools to measure each of these effects at each stage in construction.

6.1 Domino Ring Sampler (DRS4)

The Domino Ring Sampler (DRS) is a form of wave-form digitizer, or digital oscilloscope. The DRS4 receives up to four signal inputs through SMA ports on one end of its smooth metal casing, as well as optional trigger inputs and outputs that can be daisy-chained between multiple DRS's using micro-coaxial connector (MCX) cables. The DRS is then plugged into a computer, to which it sends an output that can be displayed on specialized DRS software.

The DRS is *not* the primary means of collecting data for MilliQan, as it is limited in trigger logic and channels (and quite expensive for just those four channels). Its advantage, however, is that it provides a fast and easy way to probe channels of the detector without configuring all of the cables and logic of the detector's DAQ, and it displays data (including real-time generation of histograms) *without* needing to record anything. It is also extremely portable.

This means that the DRS is a great tool for detector calibration, since it can have channels swapped out easily, it allows you to see if a PMT is functioning correctly at a mere glance, and it can be transported between labs easily. For these reasons the DRS was used for the majority of data-taking during construction and commissioning, and continues to be used for small-scale testing to this day.

Setting up the DRS

The DRS is a general-purpose waveform digitizer, so it has a lot of features not needed for work on MilliQan. There are a few steps that must be taken to configure it to the specific needs of bar testing:

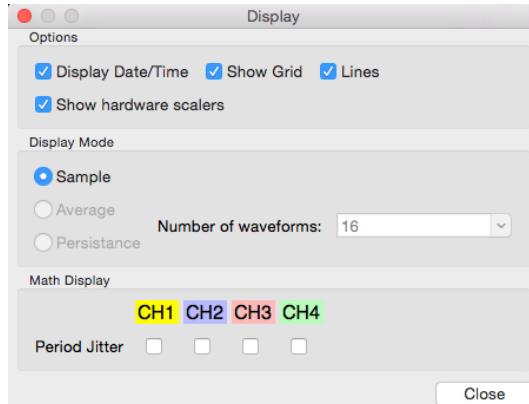


Figure 6.3: Display screen for the DRS-4. “Display hardware scalars” is the 4th option.

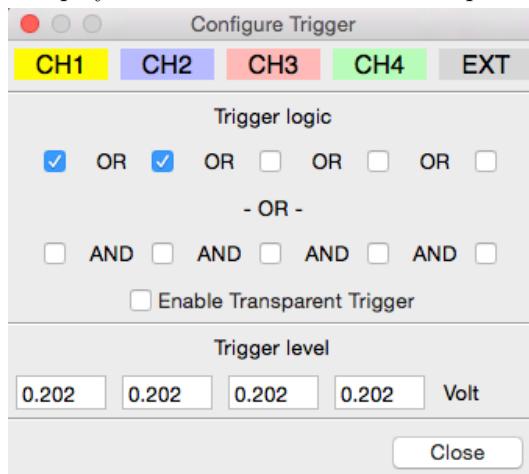


Figure 6.4: “Trigger” menu. Trigger logic (AND/OR) can be selected between multiple channels, or a single channel can be triggered on. Trigger threshold can be customized between each channel separately.

1. Make sure the DRS Oscilloscope (DRSosc) software is installed on the computer connected to the DRS4, and that any necessary installation has been completed. The DRS software can be started in Linux command line as `./drsosc`, or it can be double-clicked in a GUI. After opening the DRSosc, open its settings and select the option that enables settings to be saved between sessions.
2. On the main window, up under the “Trigger” section, select the type “Normal” instead of “Auto”, and make sure the trigger is occurring on the rising edge of a pulse (small square icon). Ignore the slider for now.
3. Open the CFG Menu, shown in Figure 6.4. This section contains the trigger logic and thresholds for each of the 4 DRS channels. For standard PMT testing, we trigger on a single channel at 15mV or 20mV, if 15mV is not above noise. You can check this by increasing the trigger threshold marginally and watching to see if the hardware scalars decrease significantly.

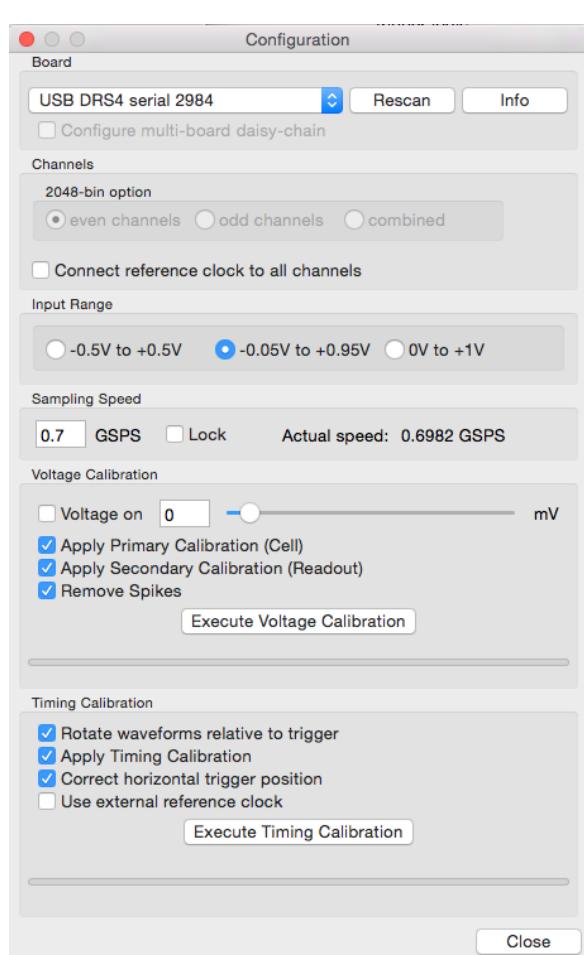


Figure 6.5: ”CFG” menu. Input range should be “-0.05V to +0.95V”, sampling speed should be 0.7 GSPS. Changing these will require a voltage and timing calibration.

4. On the main window under the “Horizontal” section, press the arrow buttons until the middle display reads “200 ns/div”. This is purely a display setting, and sets the time-width of the horizontal axis. Ignore the slider for now.
5. Under the “vertical” section, click which of the 4 channels you wish to display (and save!). Click the arrows until each of the channels reads “50 mV”. This is also a display feature, setting the mV/division on the vertical axis. 50 mV is useful, because it means a Single Photoelectron (SPE) pulse will be roughly one-box high on the grid. Ignore the sliders for now.
6. Under “Config” (shown in Figure 6.5), set the Input Range to “-0.05V to +0.95V”, and set the sampling speed to 0.7 GSPS, with no Lock. Every lower box should be checked except for “Voltage on” and “Use external reference clock”. When first using the DRS, the button “Execute Voltage Calibration” may need to be pushed. This will also need to be done when switching between Input Ranges.
7. Under “Display” (Figure 6.3), check the box labeled “Show Hardware Scalers”. This will allow real-time viewing of the activity rate of a channel in the upper-left corner. With proper triggers, this serves as a good proxy for PMT darkrate.
8. **(Wait until testing)** Once a PMT is powered on and is sending a signal to the DRS, there should be a series of digital colored waveforms corresponding to events in each connected channel, appearing like that in Figure 6.7. If these are not displayed over the full display window, power off the LV power supply and power it on again. If the issue persists, restart the DRS program.
9. **(Wait until testing)** Once the waveforms are displaying at all, the sliders from earlier need to be adjusted so that it can be seen properly. Under the “trigger” section, move the *horizontal* slider somewhere between “Min” and “Delay”, so that the rising edge of the pulses being triggered on come >300 ns into an event. Moving this slider will move a yellow tag labeled “T” at the top of the display window. This tag needs to be 2 gridlines to the right of the left-hand side of the waveform.
10. **(Wait until testing)** Under the “Horizontal” section, move the slider until the waveform fits fully within the center of the display window.
11. **(Wait until testing)** Under the “vertical” section, move each of the sliders so that their corresponding waveforms are juxtaposed in useful positions on the display. It is often useful to spread them out so that they can all be viewed at once, as we see in Figure 6.8.
12. **(Wait until testing)** When ready to record data for a PMT, press the “Save” icon on the main window. Name your file appropriately, and select the desired data format. We have been using “.dat” files, which save much faster than “.xml”. After this step, the program will ask you how many events you wish to save. Enter that amount, and the data-taking will begin. Progress can be monitored at the bottom right-hand corner of the display window.

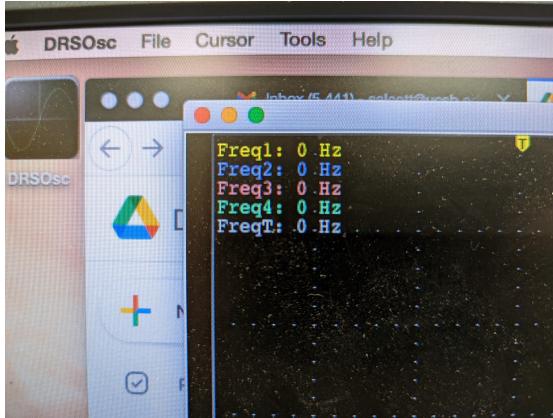


Figure 6.6: Hardware scalars are displayed in the top left corner. All are at 0 Hz when channels are not triggering.

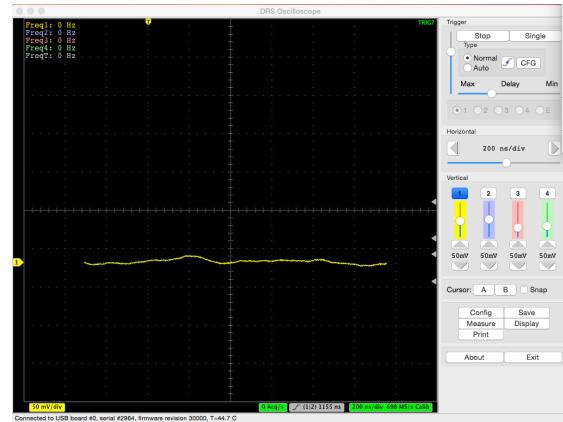


Figure 6.7: A PMT that is attached but not powered on should appear as a wavering, roughly flat line sitting fully within the viewer.

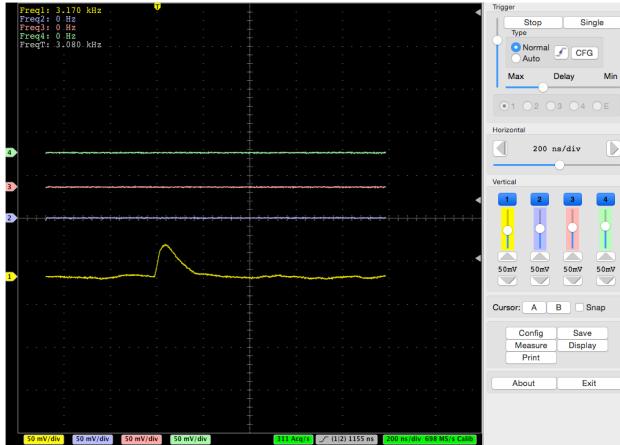


Figure 6.8: Multiple channels can be viewed at once, shown in different colors. When PMTs are powered on, a sample of triggered pulses will be visible on the steam as they occur. Here an SPE pulse appears at trigger time in channel 0.

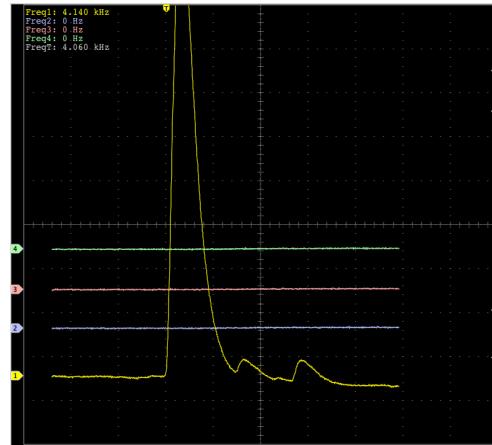


Figure 6.9: A larger pulse is seen in channel 0, likely due to an X-ray from a cosmic shower. After-pulses at SPE level appear in the time window after the initial pulse.

6.2 Gain

The MilliQan PMTs output data in the form of electronic readings, i.e. the signal voltage at every measured moment in the time-window of an event. This becomes an extremely large dataset after even a few seconds of running a single bar at HV, and most of the data is not relevant whatsoever. When photoelectrons produce a signal, they usually do it in the form of short “pulses”, and these pulses can be characterized by their height, width, and area, among a few other features. For this reason all MilliQan data goes through a pulse-finding algorithm, which spits out only these relevant parameters and discards all of the raw voltage readings from the run. All three of the mentioned parameters: height (voltage), area (voltage \times time), and width (duration) are good proxies for the **gain** of the PMT. When data is taken, these parameters are placed into histograms for the many pulses observed, and a spectrum of pulse shapes is created.

6.2.a Single Photoelectron Pulses

To compare the gains of PMTs using, say, the heights of their pulses, you need to make sure you are comparing the same type of pulse. Single Photoelectron pulses are a perfect basis for comparison, because they are generated by the lowest discrete denomination of energy being observed by the PMT: a scintillated photon. Thankfully for us, in the absence of an active radioactive source, SPEs make up an overwhelming majority of PMT pulses due to dark rate and low-energy activity in the bar. This distribution usually takes on the shape of a sharp pseudo-gaussian (actually Poisson) peak, and the value of height or area that that peak is centered around is used to characterize the gain of the PMT.

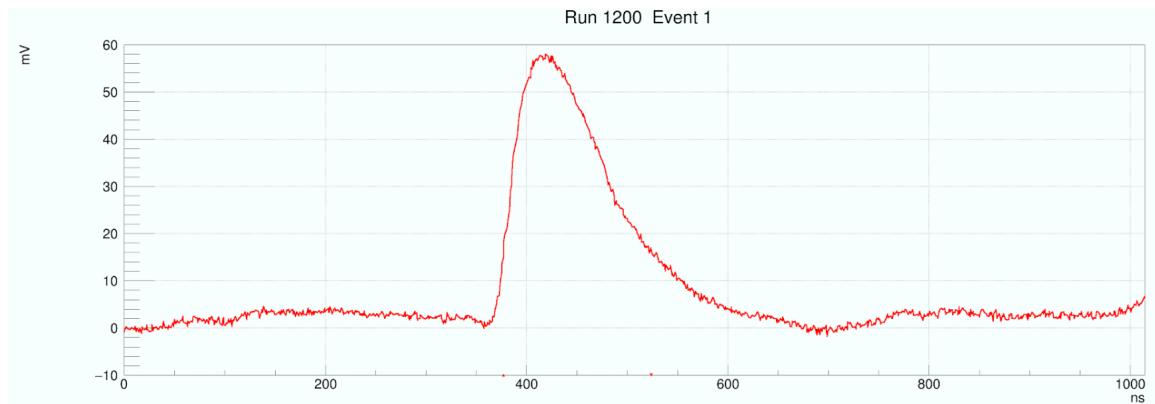


Figure 6.10: An SPE pulse waveform, taken from unprocessed event data.

One is inclined to ask: what about events where multiple SPE are produced simultaneously, or are produced simultaneously as dark rate? This is indeed possible, but it is so statistically suppressed that it can only be identified by the affect it has on the tail of the SPE distribution. If one looks very closely, the SPE tail has barely pronounced “shoulders” at pulse heights that are integer multiples of the peak’s center—remember, energy is quantized! These are photons. Then there is the case of cosmic rays, which are always raining down on us, even deep underground. It is true that these particles, usually muons and the high-energy photons that they produce through scattering, tend to deposit a relatively-significant amount of energy and thus very large pulses, but this does not have a first-order affect on pulse distributions when running R878 PMTs at +1450V. In the case of gammas, these pulses are massively varied in their size, sprinkling the bins in a pulse-height distribution with a few entries across many hundreds of mVs. In the case of direct contact from a muon, the pulses are large enough to saturate the signal, in other words the output voltage of the PMTs reaches its maximum possible value (in the case of our digitizers, 1250mV). This means that for pulse height, muon hits all get lumped together into an overflow bin, and rarely ever affect the rest of the distribution. Muons also produce a significant amount of “afterpulses”—low energy radiation from scattering that rattles around in the PMT, sometimes producing an abundance of SPE-like pulses. In this way the SPE peak grows even larger. In Figure 6.9, you can see what is likely a large gamma-ray (high energy photon) pulse, followed by two coincident SPEs, possibly afterpulses.

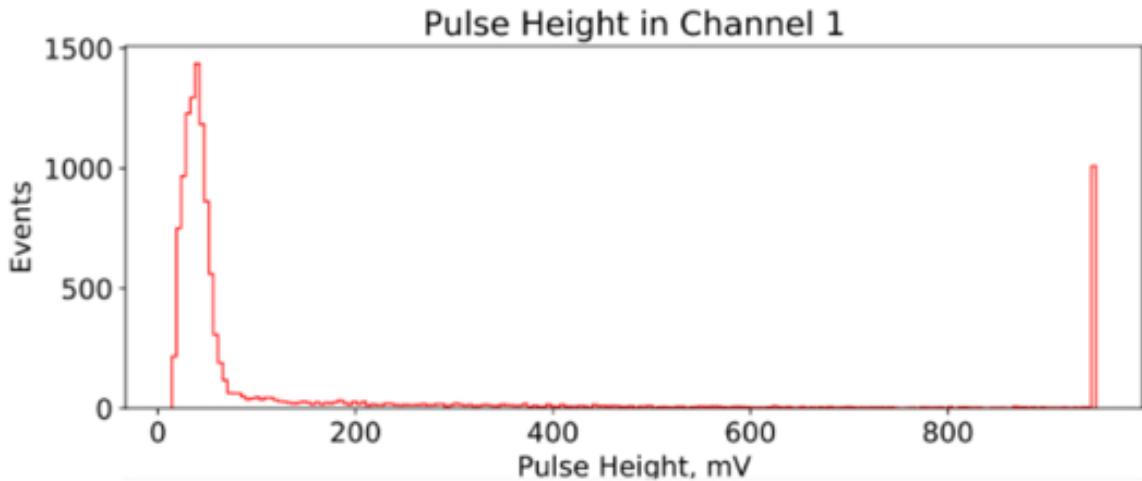


Figure 6.11: A distribution of pulse heights for events in a channel. One large SPE peak is visible at the lower end of the dynamic range, with a spike at the very end in the overflow bin.

6.2.b Nominal Range

In the case of the MilliQan bar detector, PMT SPE pulse height peaks (synonymous with gain) are traditionally in the range of 20-40mV, with a few outliers allowed all the way up to 50+mV, but none allowed under that range, as most PMTs begin to see significant electronic noise below 15mV in our setup. The reason that these values must be constrained and kept above noise is that a standard voltage threshold is used to determine whether a “real” pulse has occurred, and to trigger the DAQ. There is nothing fundamental about these numbers, they are a product of the PMT model and amplification electronics we have used. Standard SPE size for the demonstrator was closer to $\sim 1\text{mV}$.

6.3 PMT Energy Response

The most essential PMT parameter to mCP detection is energy response. In order to maximize sensitivity to low charges, the number of photons *observed* in readout for a given deposit of energy in the scintillator must be maximized. There are a few contributing factors to this effect, the most important of which being:

1. **NP/keV:** the number of scintillated photons produced for a given 1 keV of energy deposit in the scintillator. This is well understood for our chosen scintilaltor.
2. **Light Collection Efficiency:** Not all photons produced in scintillation events make their way to the end with the PMT, with some being absorbed in or penetrating through the wrapping of the bar and escaping. There is also the matter of optical coupling efficiency, since the interface of an R878 PMT is circular while the interface of a bar is square, any air gap between the PMT glass and scintillator may cause additional reflections, and because any obstruction (scratches, fingerprints, etc) will potentially cause additional losses. Great care has been taken toward the quality of design, wrapping, and PMT mounting to maximize this efficiency.
3. **Quantum Efficiency:** Once a photon has entered the PMT, it is still possible that a signal will not be produced. This may occur if the photon does not successfully

produce an electron at the photocathode, if the photoelectron is not correctly guided by the focusing electrode to the first dynode, or if the photoelectron does not hit the first dynode in a sufficient location to produce a successful cascade. Very little can be done to change or directly measure this effect, but it causes the largest variation in energy response between different R878 PMTs.

Studies of energy response in our bars measured directly the calibrated NPE detected by the PMT per keV of known energy deposited, or NPE/keV, a measurement that encompasses all three of these effects in some proportion, distinct from the NP/keV in the scintillator.

6.3.3.a Calibration with Cadmium-109



Figure 6.12: A CERN Cd-109 source.

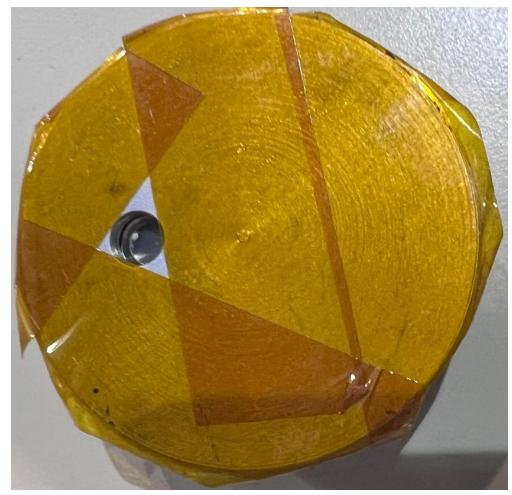


Figure 6.13: Hole for columnation.

If our goal is to measure the NPE/keV of our PMTs, or at the very least to compare the NPE/keV of PMTs to each other, we need a way to expose our bars to radiation of a consistent energy that produces pulses within the limits of our signal voltage limits. There are many choices here, but for the MilliQan bars the collaboration used Cadmium-109, an isotope that primarily emits 23 keV X-rays, with much rarer 80 keV X-rays and 2-3 keV X-rays that are absorbed before reaching a bar or contribute more SPEs.

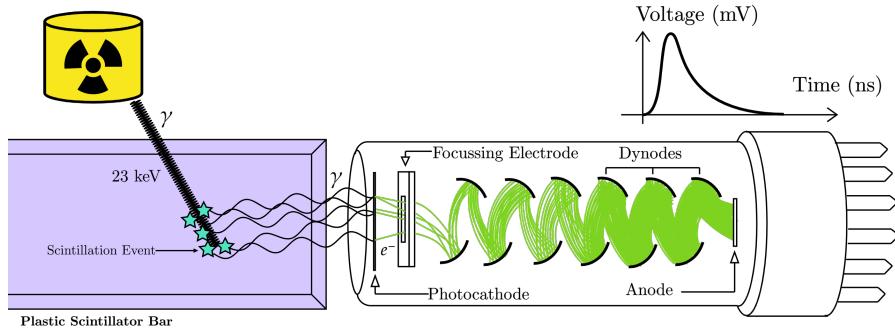


Figure 6.14: The process by which a large number of photoelectrons (NPE) are created by an X-ray from Cd-109.

When a Cd-109 source is placed along the bulk of a MilliQan bar, these X-rays cause scintillation activity that usually produces 6-10 photoelectrons per event, meaning that the pulse height distribution now shows two peaks: a narrow SPE peak and a broad Cadmium peak. In a sense both of these peaks are Cadmium peaks, since Compton-scattered Cadmium X-rays and 2-3 keV X-rays produce a frequency of SPE events that dwarfs darkrate for even a moderately active source. The task of measuring the NPE/keV is now relatively straightforward. All one has to do is fit these two peaks (pulse area is a more consistent measure in this case, due to fewer anomalous pulses from electronic noise and practically no saturation peak), and divide the center of the Cadmium distribution by the center of the SPE distribution. This result can then be divided by ~ 23 keV to approximate NPE/keV. For the sake of calibration the actual calculation is not meaningful, because the NPE is sufficient for comparing PMTs and determining the best available for use.

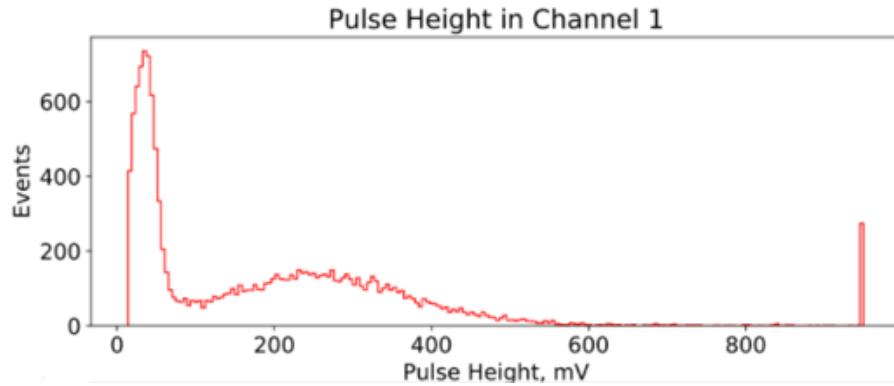


Figure 6.15: A pulse height distribution for a bar exposed to a Cd-109 source. A second peak has appeared, corresponding to large pulses from X-rays. The SPE peak has also grown, due to additional low-energy scattering from Cadmium radiation.

To obtain a better measurement, an additional step is taken to avoid biasing: the SPE peak is *scaled* and *subtracted* from a source run distribution, leaving a 23 keV Cadmium peak not biased by the tail of the SPE distribution in fitting. This works because the SPE distribution is the same shape with a sufficient number of events between a no-source and a source run, meaning they cancel out so long as they are normalized.

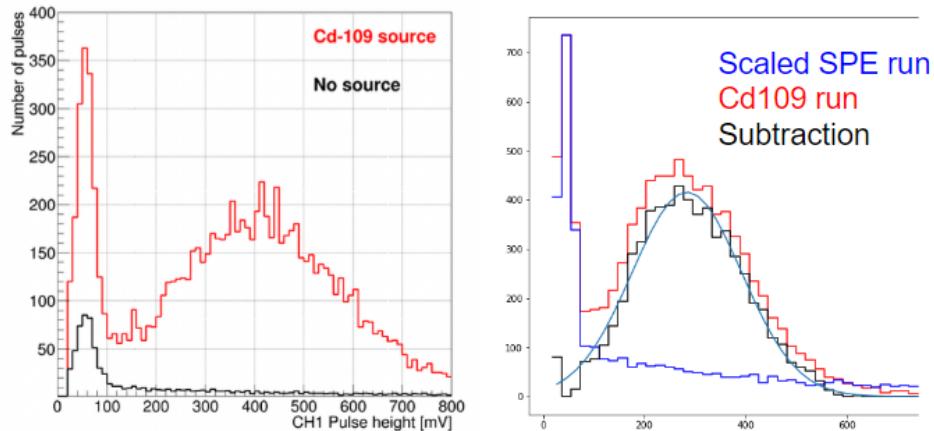


Figure 6.16: (Left) No-source vs Cd-109 pulse height distributions [24]. (Right) A no-source run is scaled up and subtracted off a Cd-109 run.

The source peak that remains is now additionally easy to fit with an approximate Gaussian, because it is the only peak in the data.

6.3.b Height, Area, and Duration

It is worth clarifying the differences between pulse height, area, and duration in the context of calibration. When a PMT's signal rises above its trigger threshold, the DRS or other digitizer records a window of data around the pulse. The offline software finds the highest and lowest points, and records that as the pulse height. It looks at how much time it took for the signal to fall back down to the trigger threshold after passing it on the rising edge of the pulse, and records that as duration (or pulse width). Pulse area is the integrated area under the waveform. Pulse height is very useful for measuring gain, because it is a direct indication of what voltage level SPEs and other signals will rise to, which can be used to set sensible triggers and cuts in data collection and analysis. It comes with the disadvantage, however, that pulse height is maximized at 1250mV, and many high energy particles will far surpass that, saturating the signal and leaving a pulse that plateaus. This is where pulse area may come in handy. After the height of a pulse saturates, its area may continue to grow as the plateau gets wider and wider. While some information about the shape of the top of the pulse is lost, this means that area is preferable to height when studying cosmic rays and other high-energy interactions. Eventually area plateaus as well, when the size of the pulse is large enough to take up the entire window of saved data. Pulse area is also a better general measurement of energy response, because it is in proportion to the number of photoelectrons produced, and is more reliable than pulse height for non-SPE pulses. Consider Figure 6.17: here two SPEs have been produced in short succession by a blinking LED, causing their pulses to overlap in the digitized waveform. Because they overlap above the trigger, the measured pulse height is close to that of a single photoelectron, when the pulse actually contains two. Pulse area will still measure the combined area of both pulses, whereas pulse height underestimates the energy response.

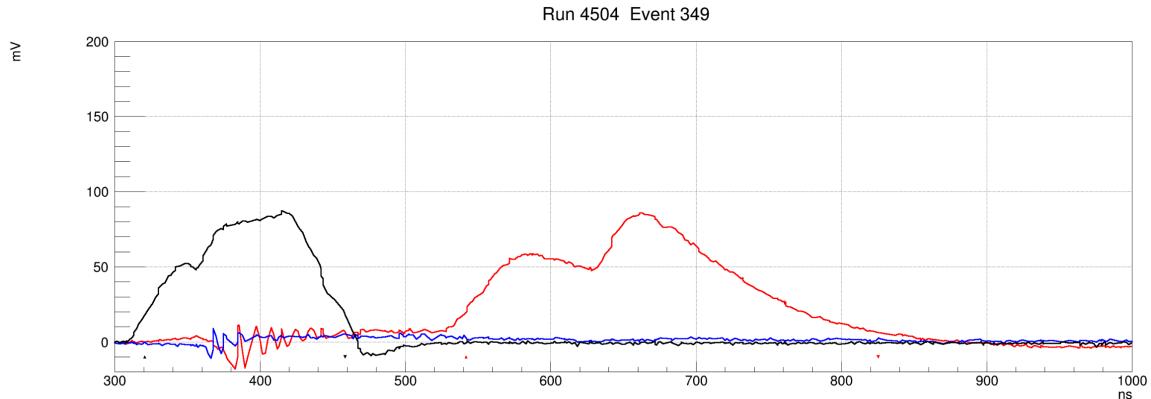


Figure 6.17: An LED pulse (black) causes two scintillation events, which happen to be so close in time that they overlap. This is counted as one pulse with the height of the 2nd peak but the area and width of the combined shape.

This brings us to pulse duration. Figure 6.17 will register as a single pulse with an unusually large pulse duration, and that makes it a useful parameter for making quality cuts on data. SPE pulses have fairly consistent rise and fall times, meaning that pulse duration can be used liberally as a filter to remove large time-scale drifts in GND, sinusoidal interference from WiFi, sharp spikes due to electronic noise, among other unwanted false pulses.

6.4 Dark rate

The dark rate of a PMT can be measured in real-time by connecting a bar to the DRS, setting a sensible SPE trigger threshold (\sim 15-20mV), and observing the scaler trigger rate in the upper-left corner of the DRSSosc program. This quantity is not *really* dark rate: it is merely the rate at which the DRS triggers and saves data, which it does every time the signal it receives rises above the set threshold. This can be caused by radiation, cosmic rays, light leaks, electronic noise, low-energy scintillator activity, and indeed true thermionic dark rate. This distinction is not very important for calibration, because factors other than light leaks and dark rate will affect the bars equally, and we wish to check for and minimize both of these affects. It does mean that dark rate should be measured only from no-source runs though, of course, as placing a radioactive source next to the bar like Cd-109 increases the event rate significantly: beyond a factor of 100 for the source we used at CERN. There is one other nuance to measuring dark rate, one that is especially relevant for a “test-bar” setup. Test bars were used in early construction, when few bars had been wrapped. In these tests many PMTs would be swapped out from the same bar, held in by a loosely-attached PMT mount that could be opened like a clam-shell to allow PMT removal and insertion. These test bars were not light-tight, of course, so they were placed in “dark boxes” covered with felt and often calibration was done with the lights off. Light leaks during tests were actually not an issue once the bar was fully covered by a dark box, but the light that entered the scintillator while PMTs were being switched *was an issue*. This exposure leaves activity with longer decay timescales, and can even double the observed dark rate if scalers are taken immediately after swapping in a new PMT. For this reason, it was important to wait for 5-15 minutes after swapping in a new PMT to measure scalers.

6.5 Bar Procedure

6.5.a Materials

A few materials will be needed to calibrate a MilliQan bar:

- (x1-4) Bar detector component.
- Positive HV power/readout base board (x1-4).
- HV Power supply (positive polarity, up to 1450V, 0.2mA required per bar).
- 2-chan LV power supply (or two 1-chan supplies) and appropriate cabling.
- Ribbon cable LV power splitter board
- 8-pin ribbon cables
- DRS4 Evaluation Board
- Computer+Monitor+DRS4 software
- SHV-SHV HV power cable (x1-4)
- HV splitter box (as needed per bar)
- SMA-SMA signal cable (x1-4)

- Thick cloth large enough to cover bar fully
- Puck-style Cd109 radioactive source (preferably $> 1\mu\text{Ci}$)

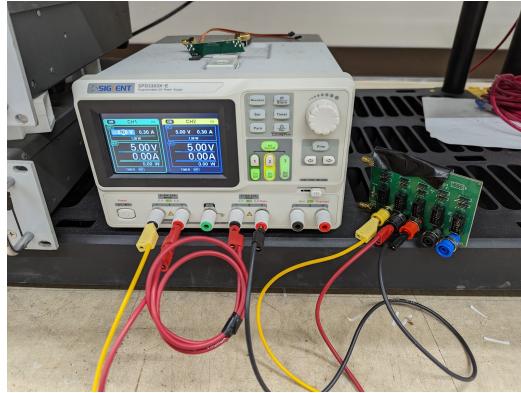


Figure 6.18: A much smaller power distribution board with 5 channels is used for PMT testing and calibration.

6.5.b Preparing the Bars

Before powering on the bars, they need to be cabled and set up.

- First unpack the desired number of bars to be tested from their crate or storage location (up to 4 for one DRS). They are fragile and the wiring for the calibration LEDs on the PMT-end can be easily torn without caution.
- Inspect the bars visually to find any obvious damage, such as torn wrapping or damaged/missing PMT dynode prongs. Do not test any obviously damaged bars.
- Without any cables plugged in, attach a PMT base board to the dynode prongs of the PMT on the end of each bar.
- Attach a ribbon cable to the black connector on the base board, an SMA cable to the SMA port closest to the inserted ribbon cable, and an SHV cable to the SHV port.



Figure 6.19: A fully wrapped bar with no base.



Figure 6.20: A base attached to the end of a bar.



Figure 6.21: Proper cables inserted into a bar's base.

- (e) **Without power turned on**, attach the other end of these cables to their corresponding supply/readout. Ribbon cable goes to LV power board, SHV goes to HV supply/splitter, and SMA goes to DRS4 readout. DRS4 takes at max 4 SMA cables, and LV power board takes at most 5 ribbon cables.
- (f) Configure the LV power supplies so that they share a common ground and have ports held at relative +5V/-5V. Without power, run +/gnd/- to appropriate ports on the LV power board.
- (g) Make sure HV power supply is set with positive polarity (negative polarity would be disastrous).
- (h) Launch the DRS4 software on its attached computer, and make sure it is set up correctly.

6.5.c LV Light-Leak Safety Tests and Inspection

It is possible that the wrapping of the bars could have been damaged in shipping, introducing large light leaks. For this reason it is important to first test the bars for light leaks in a very low-light, low-voltage environment to avoid oversaturating and destroying them. Four bars at a time can be tested for safety, following the below steps:

1. If a visual inspection has not been completed, check each bar thoroughly for any signs of torn tape/paper in the wrapping. If any of these are visible and can be patched easily, do so.
2. Once a bar has had a base attached with all of the necessary wiring, but before any power is turned on, place the bar detector fully under a large felt/blanket or into a dark box. If possible, turn the lights off and minimize sunlight.
3. Power on the low voltage, ensuring that flat, static waveforms appear on the DRS display and cover the entire event window. If a large amount of noise, or a periodic waveform appears and is accompanied by high hardware scalars but the HV has not been powered on, power-cycle the LV supply and check all of the connections.
4. Set the HV power supply to **POSITIVE** +300V, and turn it on. This is a very low-gain configuration, and SPEs should not be visible. Cosmic muons, however, should be visible. These will not appear as perfect pulses, but they should be easily discernible. The trigger threshold may need to be increased to see these events properly.
5. Gradually increase the voltage in small \sim 100-200V steps up to +1450V, monitoring the hardware scalars and shape of event pulses after each step. **Never exceed +1450V**. At no point should the trigger rate from scalars rise above 15 kHz at a 20mV trigger threshold, and it will usually be much lower. If increasing the voltage results in a rate over 15kHz, a light-leak is present and the bar should be set aside for patching.
6. Once at +1450V and triggering at 15-20mV, there should be small but visible SPE pulses in each event, interlaced with the occasional tidal wave of a cosmic muon. The trigger rate from scalars should be between 3-15 kHz. If these criteria are met, remove the light-blocking felt and monitor the scalars. If they increase, there is a light-leak and the bar should be set aside for patching.

- Finally, turn on the lights. If the trigger rate from scalars increases, there is a light-leak and the bar should be set aside for patching. Otherwise, the bar is light-tight and safe for testing.

6.5.d No Source/Dark rate Testing

The first data to collect for a bar is called a “no source”, or dark rate test. The purpose is to obtain a clean measurement of the PMTs dark rate and the position of its SPE peak as a measure of gain. This peak can then be used to measure NPE/keV for the PMT when compared to source tests later. The following procedure is used:

- Once bars are deemed safe, set them all up with low-voltage on, high-voltage at +1450V, and a trigger threshold at 15-20mV.
- If testing multiple bars, make sure all of their waveforms are being displayed in the DRS **or their data will not be saved**.
- When saving no-source data for a set of PMTs, it is important to take $\geq 10k$ events in order to have proper statistical resolution. If all of the PMTs share a similar dark rate (indicated by the hardware scalars), this can be achieved by triggering on the OR of every channel for 40k events. However, it is often simpler and easier to only trigger on a single PMT at a time, and save 10k events for each. I recommend doing the latter, and simply switching the bars via the trigger menu. This keeps the data clean but does not require any rearrangement of the SMA cables.
- Save 10k events for each PMT as a properly formatted file to the proper directory, depending on your analysis setup.
- Make sure to turn off the high-voltage before changing any HV settings, and before attaching/unattaching a base to a PMT.
- Because there is no source present, and because cosmic events are easy to cut from the data, these tests serve as a representation of a PMT’s dark rate/dark-current, a portion of the experiment’s background that can not be shielded.

6.5.e Cd109 Source Testing



Figure 6.22: A bar sits with a small Cd-109 source puck sitting along its bulk at UCSB.

- Follow the same procedure as the no-source testing, but this time placing a small Cd109 on the bar, approximately 15cm from the non-PMT end. The exact placement will not make a great difference as long as it is along the bulk of the scintillator and that there is nothing physically blocking the two (such as the plastic PMT mount on the PMT end).

2. For this testing it is required that each run of PMT events is saved individually per bar, as the Cd109 source must be directly resting on a bar for the recorded data to be useful (it can not “catch” multiple bars at the same time).
3. When a Cd109 source is placed on a bar, its hardware scalars will shoot up appreciably, depending on the activity rate of the source. This does not pose any risk to the PMT or bar.
4. Because Cd109 has consistent deposits of energy, the NPE/keV of this data when compared against the no-source runs will serve as a good proxy for quantum efficiency. This can be approximated by measuring the ratio between the pulse height distribution peak values between the SPE peak (usually $\sim 50\text{mV}$) in a no-source run and the Cd109 peak (usually $\sim 400\text{mV}$) in a source run.

6.6 Panels

The panels of the detector are given less attention than the bars during calibration, since they essentially are not a part of the fiducial volume of the detector. Their gain should be within appropriate ranges and they should be given the most efficient PMTs possible, but their performance takes lowest priority overall. The only factor demanding attention on the panels is their dark rate. Panels are veto triggers for events, meaning every event containing dark rate activity in a panel will be thrown out, even if there was no cosmic shower or other activity involved. Panels are also very prone to light leaks, so it is important to check and minimize their dark rate to the best of ones abilities before installing them.

6.7 Units

Once bars are assembled into units following the instructions in the previous chapter, they need to be calibrated an additional time with and without Cd-109. There are two reasons for this: for one, it is possible for the wrapping to be torn or the PMT mount to be jostled or damaged during the unit assembly process, affecting dark rate or light collection efficiency. Perhaps more importantly though, a second type of Cadmium testing needs to be performed for each of the bars, where a Cadmium source is placed not on the bar directly, but between the gap between two bars on the far side from the bar in question. The reason for this is that some bars will not be directly accessible once they are in supermodules and installed in-situ, meaning a method like this must be used to calibrate the final detector. By recording data with the same relative positioning in this earlier test, we obtain a direct comparison between the far-bar measurements and the on-bar measurements using the same equipment, which can be used to extrapolate all of our desired quantities once we are calibrating the full detector. This far-bar positioning is demonstrated in Figure 6.23, where data will be taken both with the bars in contact with the source and those only visible to the source through gaps. In Figure 6.24, we see that there is a noticeable affect on pulse height distributions between the on-bar and far-bar tests. This is why it is important to measure all three during this stage: no source, Cd on-bar, and Cd far-bar.

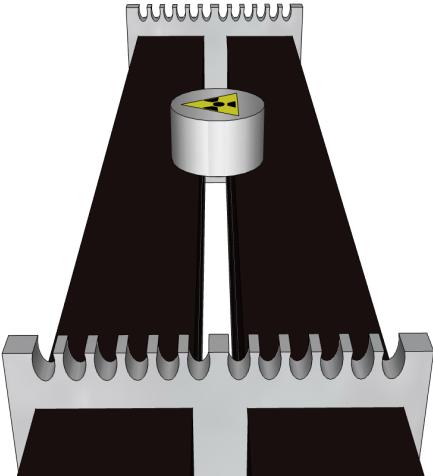


Figure 6.23: A 3D model shows a Cd-109 source positioned between two bars in a unit, so that its collimated X-rays can reach the two lower bars.

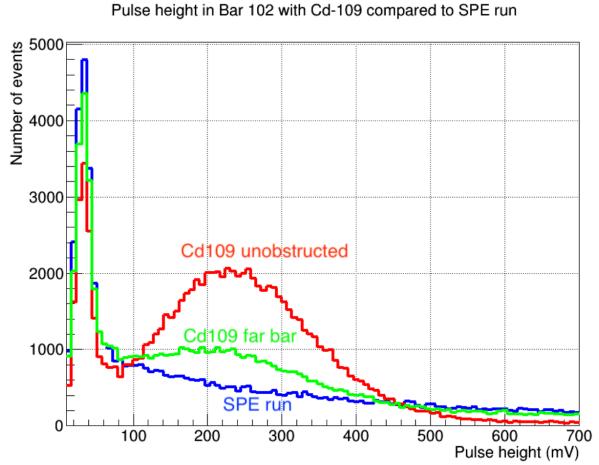


Figure 6.24: A pulse height distribution is compared between bars on the left, with one touching and one further away. The “far bar” peak is diminished, but is still measurable and can be used for calibration [24].

The actual procedure for testing the units is nearly identical to the procedure for the bars. Be sure that you are testing and triggering on one bar at a time, first taking a run with no source, then following up with a run taken with the source on that bar, then moving the source to the opposite side of the unit and taking an identical run in that position.

6.8 Supermodule

Once units are safely fitted and cabled into supermodules, and those supermodules inserted into the detector, it is time for the final round of Cd-109 calibration tests. This final test is performed to ensure that the performance of each channel is well understood even in their final configuration, that no damage occurred during installation, and to measure the effect of the CMS magnetic field on each of the PMTs’ performances. Because the detector is compact and shielded, this is the most complicated calibration to perform.

6.8.a Source Testing Procedure

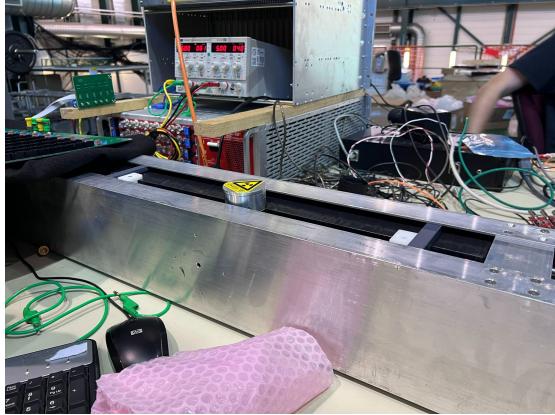


Figure 6.25: A Cd-109 source is placed between two bars in a unit, in a supermodule.

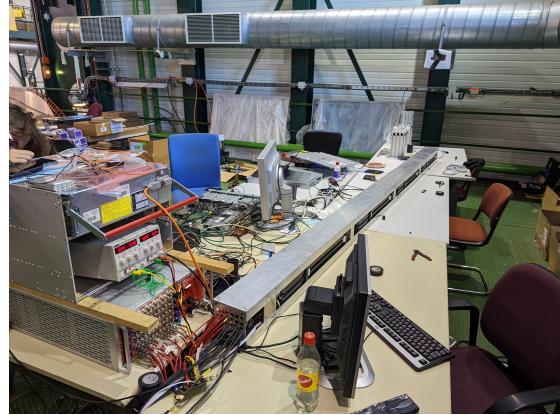


Figure 6.26: A fully cabled bar lies at the upstairs workstation in CMS-S.

1. Level the detector to horizontal position. Turn off HV for the two top panels, and carefully remove and set aside their bases. Unscrew their support bars and remove the top panels from the detector. Set them safely on a table with foam sheets to support them. Side panels do *not* need to be removed. Return the detector to 43.1° towards IP. Lifting and transferring is a 2+ person job.

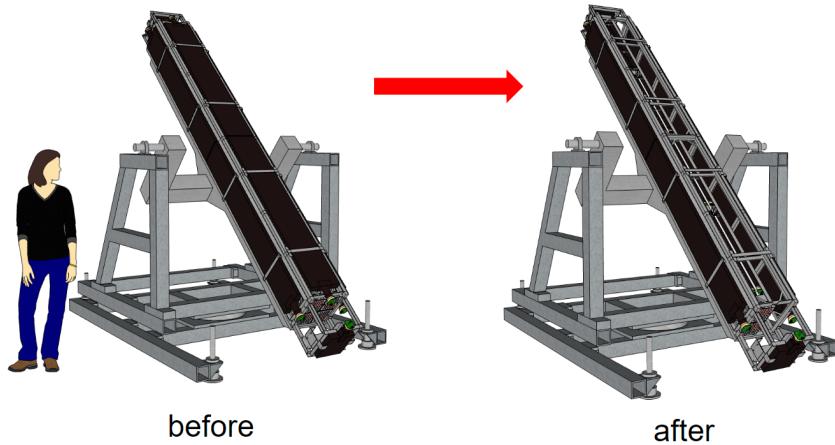


Figure 6.27: The top panels of the detector are removed in this 3D model to allow for proper source calibration of all its channels.

2. In order for X-rays from the Cd109 source to penetrate the bars and give accurate calibration, the hole on the bottom of the source container must be aligned with the glass opening. Be sure that the source was aligned, taping it securely so it does not need to be re-aligned. This can be seen in Figures 6.12 and 6.13.
3. Next prepare the DRS4 for testing, and prepare a log. Configure settings for the DRS4. To calibrate a supermodule, each test will be 10k events, triggering at 10mV. Specify all of the variables for each run. Every bar in the supermodule will have one dedicated no-source run and one dedicated source run, totalling 32 runs (saved as .dat files).
4. When preparing to test a bar, its BNC signal cable must be moved from its spot on the digitizer (leaving the BNC-MCX adapter inserted) to a BNC-SMA adapter inserted into one of four channels on the DRS. It is best to do this one at a time for each bar, to avoid mix-up. The ordering of calibration is best done starting with layer A (closest to floor) and ending with layer D. In each layer, the order of calibration runs will be left-right, top-down. On the signal cable labels this will look like Su1SA2-Su1SA4-Su1SA1-Su1SA3 (supermodule 1, layer A example) for supermodule 1 and 3 because they were rotated anti-clockwise to go into the detector, and like Su2SA3-Su2SA1-Su2SA4-Su2SA2 for supermodules 2 and 4, since they are rotated clockwise for insertion. As discussed, this ordering is due to the fact that the supermodule cables were routed with the cables facing up, prior to insertion, when in the detector they are on the left and right sides. A review of this convention can be found in Figure 5.61.
5. Take a no-source run for a given bar. When that is done, a source run must be taken. The positioning of the source should be consistent, and will vary for each bar and each

supermodule. For any source test, it is ideal for the hole on the bottom of the source to touch each individual bar being tested directly, at about half way down the bulk of the bar. For supermodules 3 and 4 this is possible, because gaps in the supermodule plating allow access from both above and below.

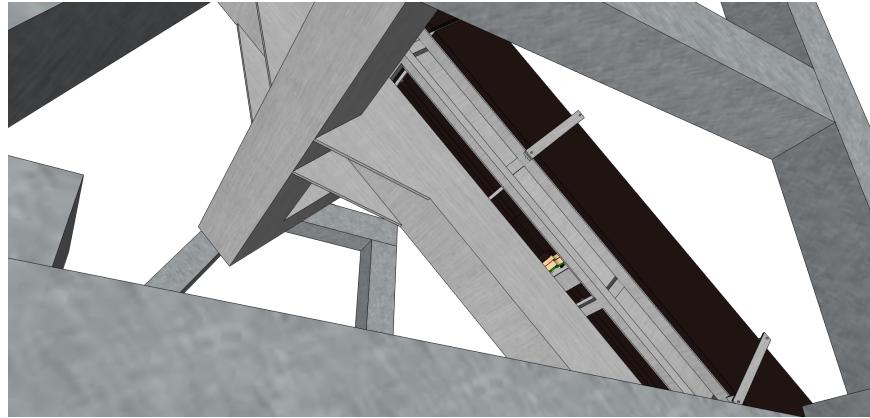


Figure 6.28: Exposed bars can be seen from the underside of the detector. A source can be placed up against these bars, necessary for the lowest row.

For supermodules 1 and 2 this will not be possible, because their undersides are blocked by supermodules 3 and 4. That is why through-gap calibration has been performed on units previously, allowing access from the top alone to be sufficient (though not ideal). For these supermodules the source will be positioned normally for the accessible bars on top (so that the hole directly touches their center), and will be placed in a central location for the other two, so that the source hole is aligned with the gap between the bars, giving “line-of-sight” to x-rays for the bottom bars. Be wary that the source hole is NOT aligned with the center of the source cylinder (offset in one axis)

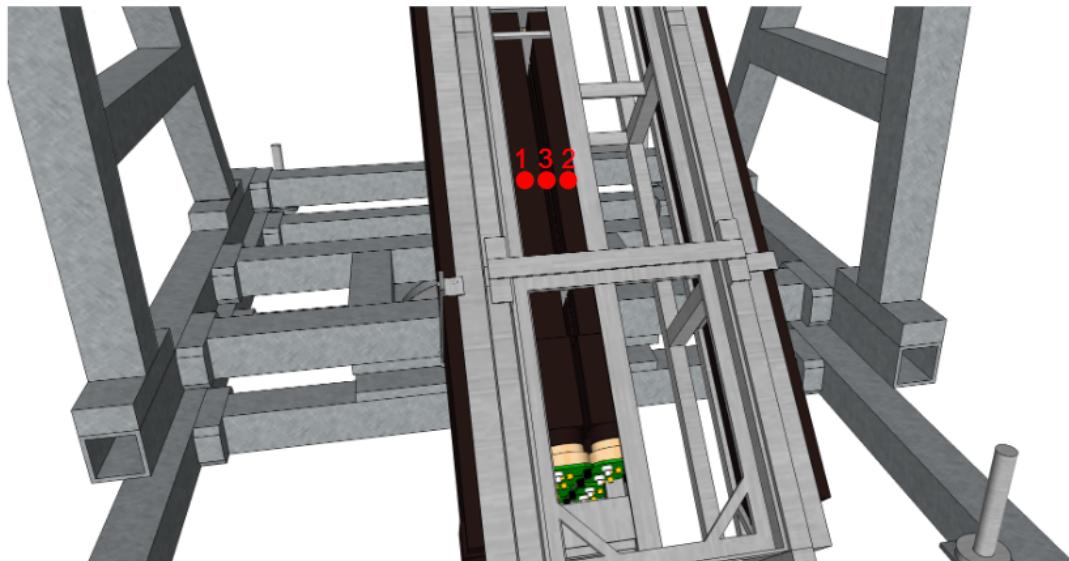


Figure 6.29: A 3D model shows the three Cd-109 positions at which data should be taken on each unit in the detector.

6. One by one, move up the detector and take a no-source and source run for every individual bar, triggering on that bar. One person will need to start the DRS runs, while another person will need to hold the source at each position while runs are going. It is important the source is in firm, unmoving contact for the full duration of a 10k event run. For supermodules 1 and 2, there will be 3 source positions per layer. One on the top-left bar, another on the top-right, and one in the center to reach the bottom two. Expected trigger rate at 10mV trigger for no source is 1-2 kHz, 100-180 kHz for direct source contact, and 20 kHz for through-gap source exposure. I have outlined the source positions in Figures 6.30 and 6.31 with red dots for these supermodules.

For supermodule 4, which will be inserted directly after supermodule 1, there will instead be four positions to place the source in each layer. Positions 1 and 2 as shown above, with positions 3 and 4 directly mirroring them on the underside of the supermodule. The source will be held upside-down against the underside, being careful as always that the source hole is properly aligned and that the rates are as to be expected. The source will need to be slid carefully into the gap between the supermodule and the support structure that holds up the cage.

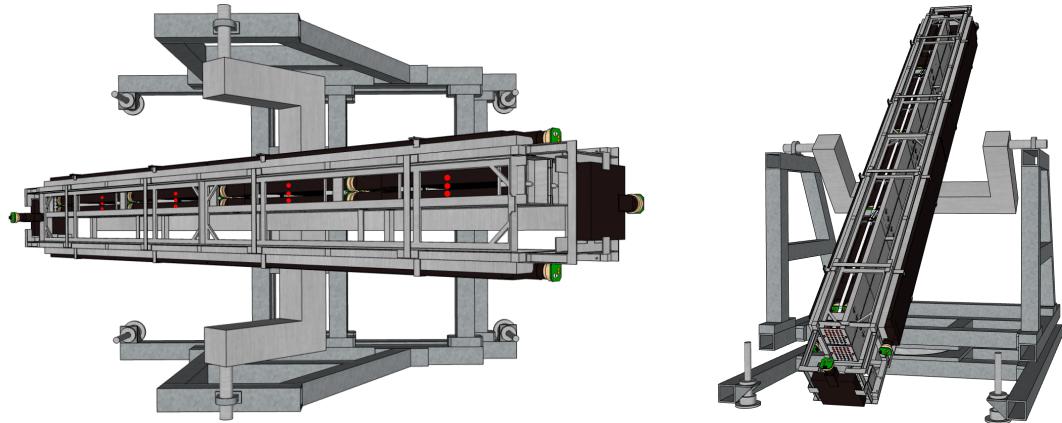


Figure 6.30: An above view of the detector 3D model, with 12 source positions along the top of the detector shown.

Figure 6.31: The detector 3D model, as seen from someone standing in the Gallery.

7. Lastly, upload the clearly documented runs to a safe server and share with the collaboration.

6.9 Monitoring and Future Calibration

At each stage of the bar detector's development calibration on darkrate, gain, and NPE have been performed and monitored for changes. This included the preliminary testing and selection of PMTs, testing on fully-wrapped bars at CMS-S, repeated calibration after unit assembly, and in-situ data taken from the gallery. This data was indispensable for identifying behavioral changes in a few bars after their shipment or assembly, and allowed us to intervene with those bars before they were installed. It also helped us identify systematic changes and understand our equipment better. Thankfully, the results indicated that the performance of our PMTs overall was fairly consistent, and that the variations between stages were well understood and accounted for.

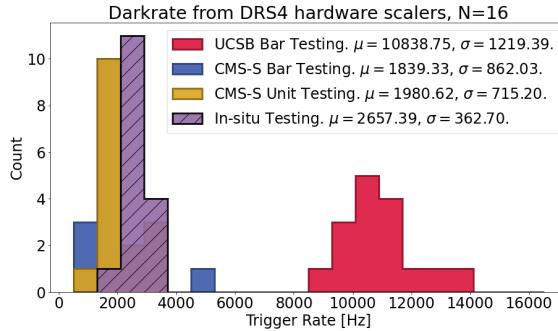


Figure 6.32: A distribution of darkrates for 16 PMTs in the first inserted supermodule of the detector, shown at every stage of calibration. Darkrate reduced substantially following the tests taken at UCSB, because the bars were fully wrapped and left in the dark for months.

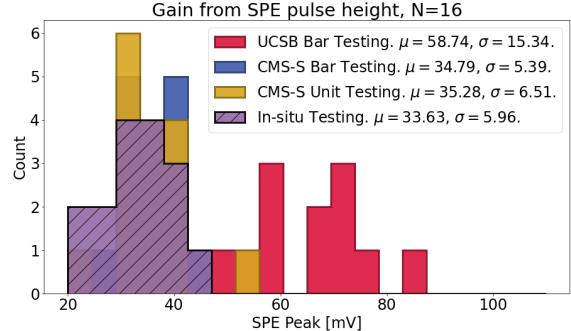


Figure 6.33: A distribution of calibrated gain in each of the same 16 channels at each of the same stages of testing. Gain reduced between UCSB and CERN due to a switch from v2 to v3 bases, but remained largely unchanged after installation due to added magnetic shielding.

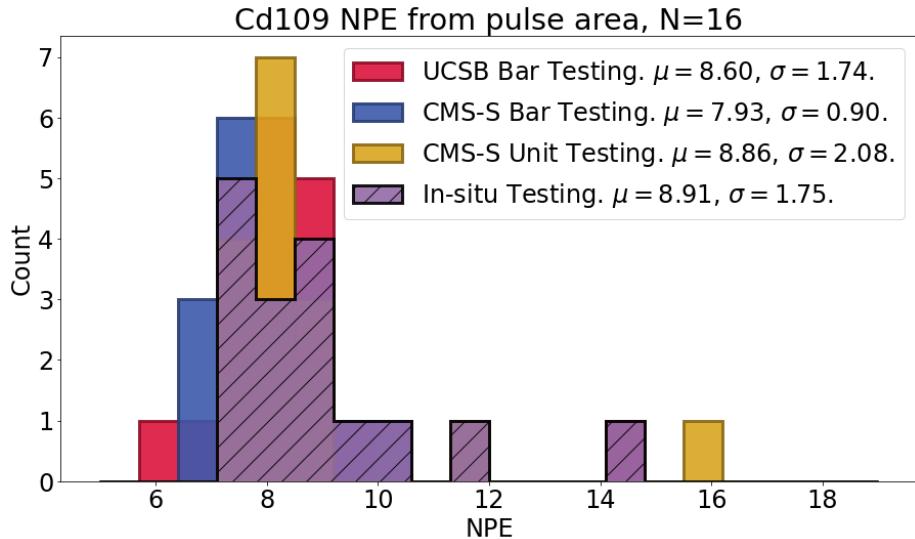


Figure 6.34: The most important parameter, NPE, for the same 16 channels. Thankfully, darkrate went largely unchanged between each of these stages.

After completing hands-on calibration the detector was ready to run, but it will still need to be monitored for changes to its performance in the future. It is not practical for personnel to enter the access tunnel and perform lengthy calibration on every channel using a radioactive source on a regular basis, but thankfully there are less invasive methods of calibration that compliment the thorough data arising from the Cadmium tests. One of these methods is the LEDs, which were installed into each PMT with the explicit purpose of remote calibration. These LEDs are controlled by the power distribution board, which in turn is commanded by an RPi to pulse the LED at a certain magnitude and width. These pulses set a customizable and randomizable trigger for the detector, allowing for unbiased measurement of darkrate and pulse height/gain. However, the LEDs do not tell us very much about energy response.

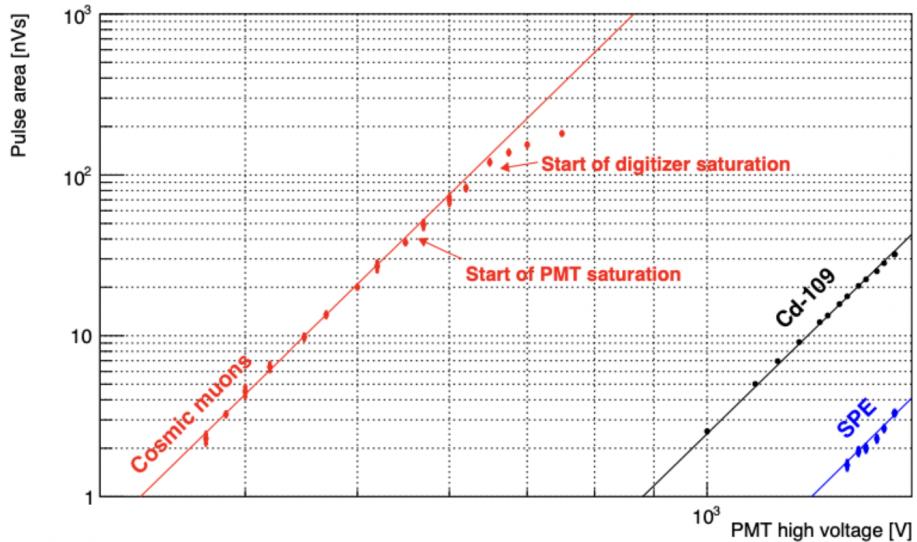


Figure 6.35: Pulse area response in each PMT for various stimuli at different high voltage levels. Both SPE and Cd-109 peaks are visible simultaneously at operational voltages (+1450V), but it is also possible to calibrate using cosmic muons by reducing high voltage significantly, down to a few hundred Volts.

For this we are actually aided by our enemy: cosmic rays. By remotely reducing the power of the detector to the low-hundreds of Volts instead of 1450V, we reach a decreased level of gain that resolves a spectrum of cosmic ray pulse sizes. By performing a small scan of HV levels (such as 200, 300, 400, 500 Volts), one obtains data points that can be used to calculate the NPE/keV of the PMT. This is time consuming and cumbersome, but it can be used once in a while to determine if an obvious shift in the performance of a bar has occurred.

Chapter 7

DAQ

7.1 Digitizers and Triggerboard

Once a detector is fully constructed and cabled, there must be a system that determines when and why each channel of the detector will trigger, selectively collects data from every channel at those times, and converts it to a format that can be analyzed in order to make physics measurements (realistic) and discoveries (optimistic). This process is called Data Aquisition (DAQ), and is implemented in particle physics experiments as large as CMS and as small as MilliQan. For MilliQan, this process starts in a central computer that contains all of the trigger information, called the MilliDAQ. This computer sends information to a custom trigger board, made with an *Altera Cyclone IV* Field-Programmable Gate Array (FPGA). This triggerboard is connected through Low Voltage Differential Signaling (LVDS) and SMA cables to five CAEN Mod. V1743 digitizers, which have a 0.4 GHz sampling frequency, a $2.5 \mu\text{s}$ readout window [33], and 16 channels. Four of these digitizers are dedicated to the four layers of the detector, while the fifth is dedicated to the panels. When the digitizer sees a pulse, it sends that information through LVDS to the triggerboard. The triggerboard makes a decision about whether the pulses it receives achieve certain criteria, including passing a trigger threshold but also programmable geometric requirements for multiple channels. This decision is sent back to the digitizer through an SMA cable, and if the signal is affirmative the digitizer will save that event, and output it through fiber-optic cable back to the MilliDAQ computer to be saved.

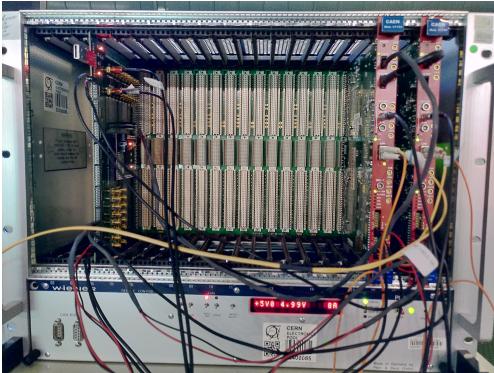


Figure 7.1: The DAQ crate being tested at CMS-S, containing two digitizers on the right and a triggerboard on the left.

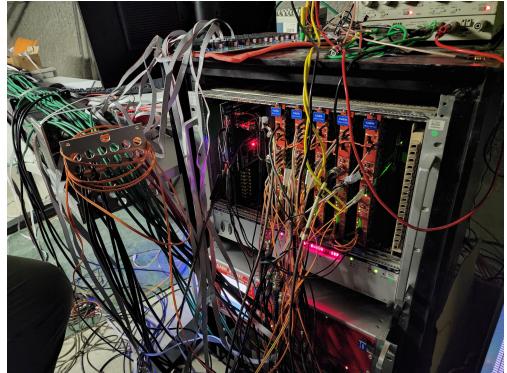


Figure 7.2: The DAQ crate being tested in the Gallery with a cabled supermodule, all five digitizers inserted.

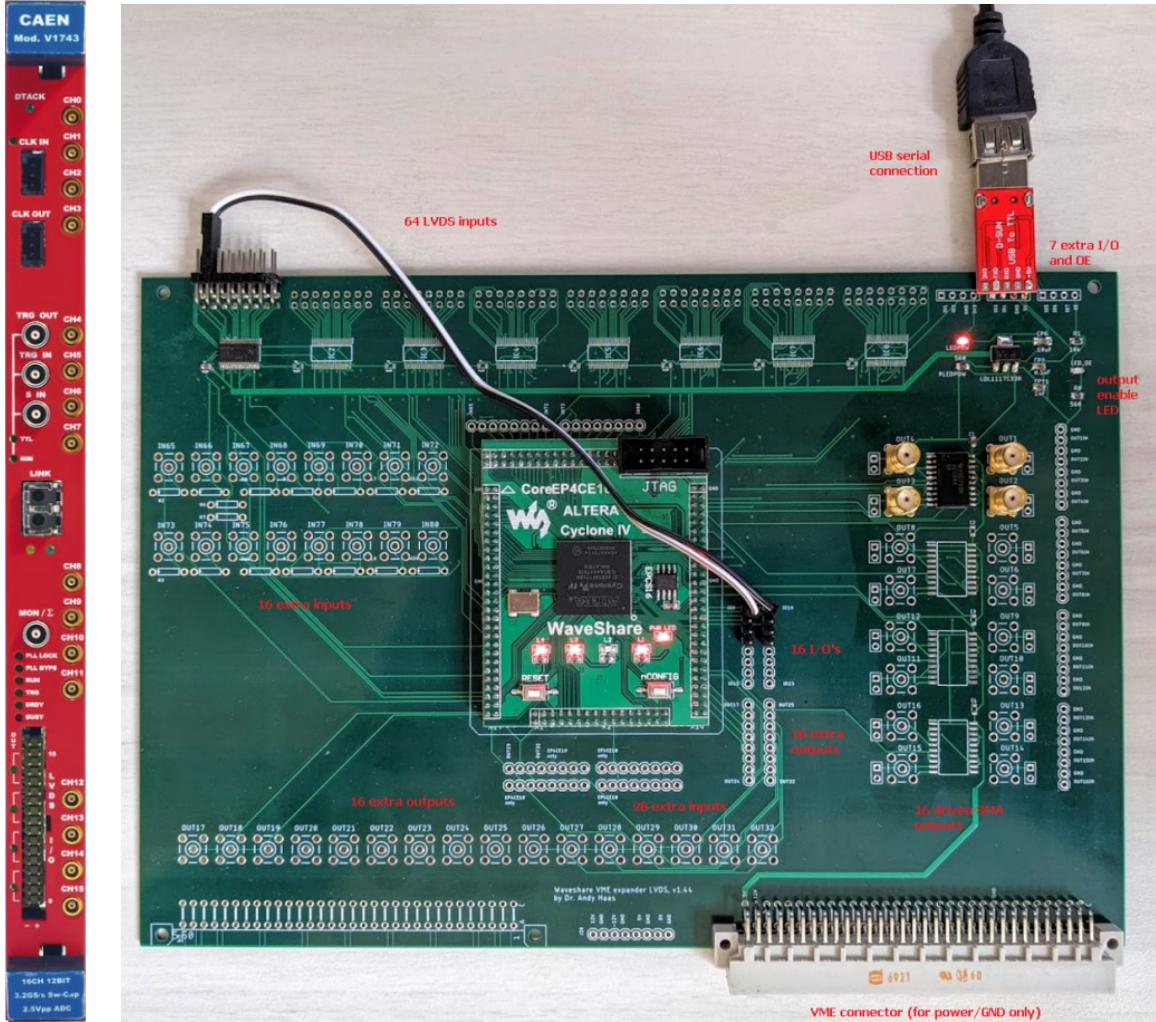


Figure 7.3: (Left) CAEN V1743 Digitizer face, with 16 MCX channels for signal, LEMO for trigger, and LVDS to communicate with the trigger board. (Right) Triggerboard, with trigger information coming from USB connected to MilliDAQ, LVDS and SMA connectors to communicate with digitizers.

7.2 MilliQan Offline

Initially, data is saved as pure voltage output, i.e. the waveforms of pulses themselves. This is far too much information to handle when analyzing data, so instead this signal is run through a pulse-finding algorithm, which records only the most essential features of a pulse and records them in a histogram-based data structure called a ROOT tree. In total 48 variables, of a variety of data types, are recorded for each pulse and put into histograms. Working with these variables is far faster than working with waveform data, and manipulating histograms (for which the only data points are the bins and the bin counts) is significantly faster still. It is from these variables that we get the useful quantities used in calibration: “height”, “area”, and “nPE”. It is also where we can get some new and useful geometric quantities: “chan”, “type”, “layer”, “column”, and “row”.

Name	Type	Description
event	int	Event number
runNumber	long int	The number of the run
fileNumber	long int	Each run is split into 1k event files.
boardsMatched	bool	Are all digitizers matched together
DAQEventNumber	int	Event number stored in DAQ
triggerThreshold	vector<float>	Threshold in mV set for each trigger
triggerEnable	vector<bool>	Status of each trigger
triggerLogic	vector<int>	Trigger logic of groups
sidebandMean	vector<float>	Mean of sideband (defined 0-50)
sidebandRMS	vector<float>	RMS of sideband
maxThreeConsecutive	vector<float>	Max 3 bins in a row of a pulse
chan	vector<int>	digitizer channel associated with this index
chanWithinBoard	vector<int>	Digitizer channel within one board (0-15)
column	vector<int>	Column of bar
layer	vector<int>	Layer of bar
type	vector<int>	Type of channel
board	vector<int>	Digitizer board connected to
height	vector<float>	Height of pulse
area	vector<float>	area of pulse
pickupFlag	vector<bool>	Is pulse flagged as pickup
nPE	vector<float>	number of photo electrons pulse is equal to
riseSamples	vector<int>	Bins between 20% and 80% max before peak
fallSamples	vector<int>	Bins between 20% and 80% max after peak
ipulse	vector<int>	index of pulses in event
npulses	vector<int>	number of pulses in event
time	vector<float>	time of pulse within 2.56ms window (us)
timeFit	vector<float>	Time of fitted pulse
time_module_calibrated	vector<float>	Time of pulse after timing calibration
duration	vector<float>	width of pulse
delay	vector<float>	Time since end of previous pulse
max	vector<float>	Max pulse height
present	vector<Long64_t>	Is there data present in digitizers
event_trigger_time_tag	vector<Long64_t>	Trigger Time Tag from CAEN digitizers (ns)
event_time	vector<Long64_t>	Time standard unix epoch (s)
tClockCycles	float	Number of trigger board clock cycles
tTime	float	Trigger board time (clock cycles * 8ns)
tStartTime	float	Start time of trigger board
tTriggerNumber	float	Trigger bit string, see list of triggers in twiki
tTimeDiff	float	Time diff. between trig and DAQ matching
tMatchingTimeCut	float	Allowed time diff. for trig/DAQ matching.
tEvtNum	int	Event number in MatchedEvents files
tRunNum	int	Run number in MatchedEvents files
tTBEEvent	int	Event number in TriggerBoard files

Table 7.1: A list and description of MilliQan Offline Tree variables, made by Michael Carrigan and Matthew Joyce.

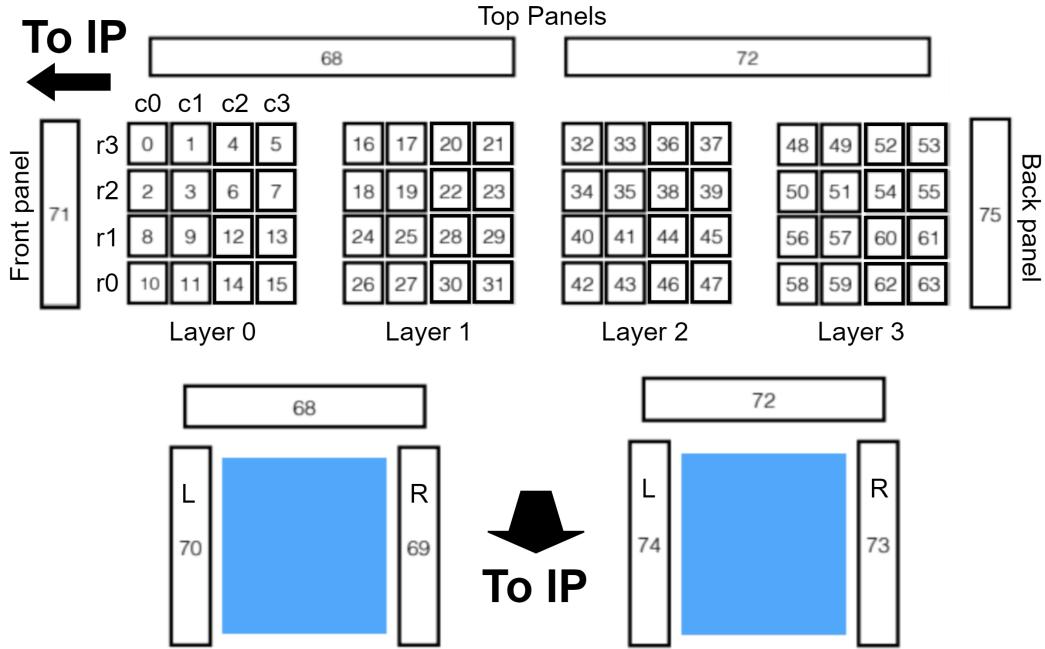


Figure 7.4: A diagrammatic depiction of the detector, with layers and rows mapped out. The top section displays each of the 64 bars, the end panels, and the top panels, while the bottom section displays a rotated view so that the side panels are now visible.

Layers, columns, and rows are self-explanatory, corresponding to each of the four layers of the detector and to the grid of 16 bars in each layer. The layers start from the floor at 0, the columns/rows start at 0 from the bottom left corner of each layer, when viewed from the IP. “Type”, meanwhile, is a variable that differentiates bars from each type of panel. Finally, all bars have an associated channel number (“chan”) that starts in the top left corner of the first layer and goes from left-right, top-bottom, but counts all bars in a *unit* before moving to the next unit. It is essential that these channel numbers are mapped to specific PMTs and specific bases, so that NPE can be tuned to match the calibration tests performed on each of the PMTs, and so that issues can be quickly identified and fixed when they arise.

7.3 Event Displays

Though waveform data is not usually useful for analysis of many events, it can be very useful to quickly examine a single event, or a sample. For this reason the waveform data is saved separately, and can be extracted and viewed for the entire detector using a script and website created by Sai Neha Santpur. All one needs to do is indicate the event # within a file, and the file number (both of which can be quickly extracted when filtering events in analysis), and the waveform in every channel for that event will be plotted. Figure 7.5 displays an event that triggered on a small pulse, shown in the middle of the time window, but happened to capture a large cascade of pulses from a cosmic ray muon that hit the detector at the start of the pulse. This event display also shows a color-coded diagram of which channels each pulse appeared in, and the pulse height of the largest pulse in that channel (most saturate at $\sim 1250V$ in this plot).

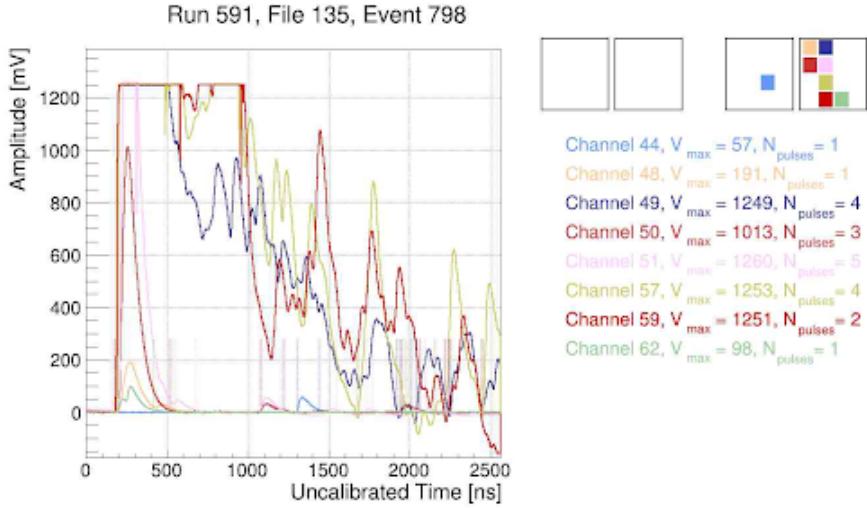


Figure 7.5: An event that triggered on an SPE-sized pulse positioned roughly in the center of the time window, but which happened to capture the typical waveform of a muon passing through the bar earlier in the event.

7.4 Converting ROOT to CSV

ROOT is a language developed by CERN for the explicit purpose of dealing with huge data sets from particle collisions, but as a custom C++-based language not used in any other fields it can be intimidating to Physics students who are novice programmers. As a result, I carried out the majority of my small-scale data analysis studies in Python, since my familiarity with it makes up for its lacking speed and our data-sets have not been massive in the early stages of the project. In order to do this, I had to convert ROOT Tree data into Comma-Separated-Value (csv) data-sets, which can be easily read and iterated over in Python. The first step in doing so is to open a .root run file using root (let's call this hypothetical file `Run.root`), which I do through the Linux distribution *Ubuntu*, as ROOT is not directly compatible with my operating system (Windows 11). Once the file is open, I run the following lines of code to output only the most essential variables to a .log text file:

```
root -l Run.root
[0] t->SetScanField(0);
     .> unprocessed.log
t->Scan("event:chan:row:column:layer:height:area:nPE:duration");
     .>
```

The line `t->SetScanField(0)` makes sure that the data is output starting at the very beginning of the file, the following line changes the ROOT output to a new text file `unprocessed.log`, and then the last line scans the ROOT file and outputs the columns indicated by name, separated by colons. After this is complete, there is usually some junk left over at the beginning and end of the text output, and sometimes these text files are so large that they can not be opened or edited by standard text-editors without exceeding allowed memory. A simple workaround is removing rows using the Bash command `sed`:

```
sed -i '1,4d$d' unprocessed.log
```

This particular line removes rows 1-4 and the last row, which are usually unwanted, but it can be modified to remove any rows. We may also want to remove some columns, which we can do with the Bash command `cut`:

```
cut -d, -f2 --complement unprocessed.log > unprocessed_shorter.log
```

Finally, we will want to check our file to see if the correct lines and columns have been removed, but this can be an issue if our data file is very large. Thankfully, we can use the Bash command `less` to view small pieces of the file without loading it all into memory, and particularly the first and last several lines are easy to view this way:

```
less unprocessed_shorter.log
```

Now we will want to convert this odd `.log` file into a `.txt`, and remove extraneous blank space. I have written a short python script to do this:

```
def scanToTxt(run):
    with open("unprocessed_{}_shorter.log".format(run)) as f:
        lines=f.readlines()

    lines = [line.replace(' ', '') for line in lines]
    lines2 = [line[1:-1]+'\n' for line in lines]

    with open('Run{}shorter.txt'.format(run), 'w') as f:
        f.writelines(lines2)
```

Figure 7.6: A quick Python function for converting a `.log` ROOT Scan output into a text file, removing excess whitespace and adding proper row formatting.

Then we will want to take this file and convert it to CSV, which is easy to manage in python, so I do this with another script:

```
def txtToCsv(run):
    file=open("Run{}shorter.txt".format(run))
    data = csv.reader(file,delimiter='*')
    output=open("Run{}shorter.csv".format(run), 'w', newline='')
    writer = csv.writer(output)
    for row in data:
        writer.writerow(row)

    output.close()
```

Figure 7.7: Another Python function for taking a `.txt` file and converting it to a `csv` file, which are easier to read and iterate through.

This leaves me with data that I am ready to analyze.

Chapter 8

Data Analysis

Throughout the long and time consuming process of constructing and calibrating the bar detector, I have had the opportunity to be directly involved in a number of analysis projects. These included optimization studies, experiments on design, background modeling, and anomaly testing. A few of these studies revealed critical weaknesses of the detector and resulted in design changes.

8.1 PMT Selection

As mentioned previously, our R878 PMTs belonged to another experiment long before they were a part of MilliQan, and as a result they cost us nothing. However, this also meant that their performance was not always consistent, and sometimes they were entirely non-functional. So before doing any final calibrations, all available PMTs needed to be tested briefly for functionality and the best among them selected for the experiment.

Katherine Larina and I began testing PMTs at UCSB on our first day in Summer of 2021, and would test roughly 120 over the following year while also completing scintillator wrapping and performing a few of the studies outlined in this chapter. Data was taken using test-bars, with and without a source, and no fewer than 20,000 events of data were recorded for every PMT tested. Periodically the existing data would be analyzed for gain, dark rate, and NPE/keV, and the best performing PMTs would be permanently mounted to bars as they became available after wrapping. By the end we wrapped and mounted 52 of the best PMTs we could to bars held at UCSB, with other bars having been sent to other institutions for wrapping. The shipment to CERN was packed with as many of the next best PMTs as would fit, for mounting on panels and serving as substitutes should any mounted PMT become damaged.

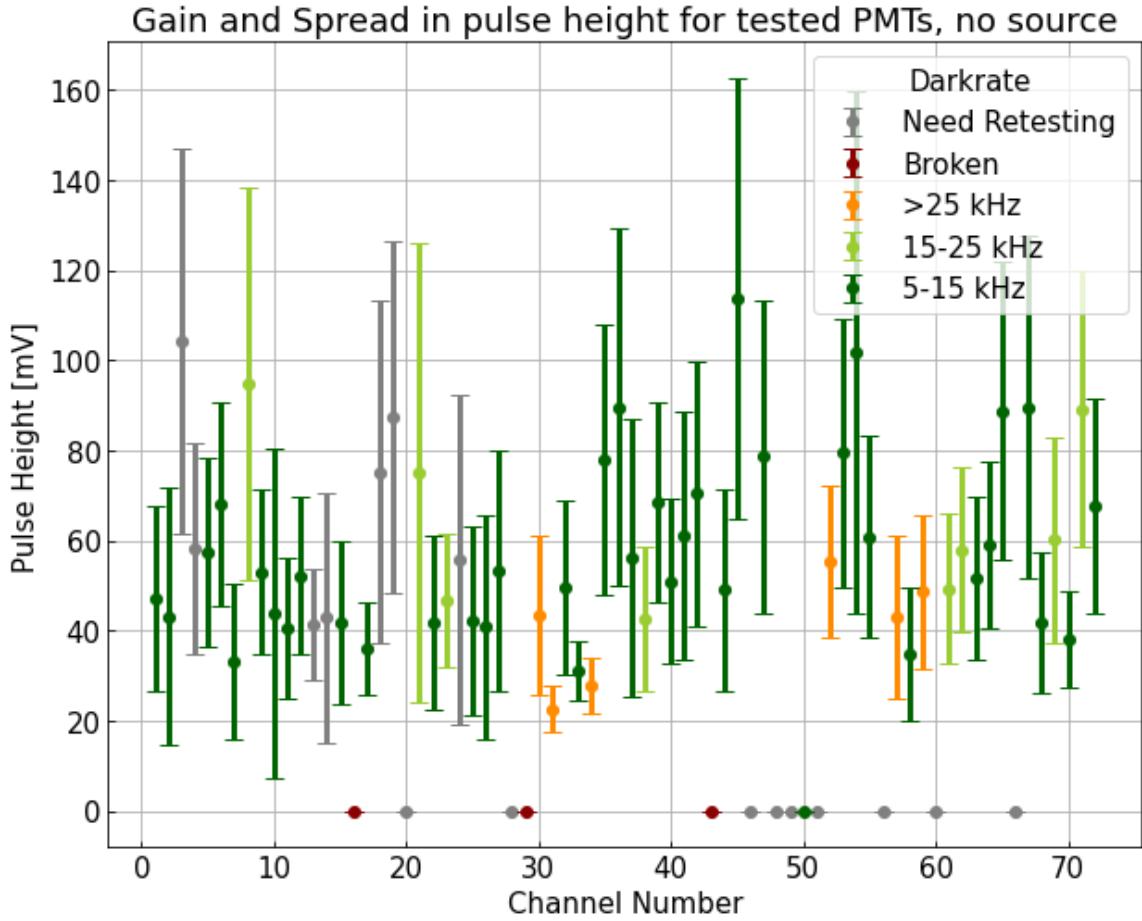


Figure 8.1: The results of an early analysis done in Summer of 2021, graphing the fitted pulse height SPE peaks for tested R878s, with error bars to represent the standard deviation of those fitted peaks. Color indicates the dark rate of the PMT, with some of the more problematic PMTs highlighted lime and orange.

8.2 Dark rate instability study

During the process of testing PMTs at UCSB, it became clear that the dark rate we were measuring from hardware scalars on the DRS was not stable or consistent between tests. Instead, the dark rate of a PMT would initially be very high after inserting it into a test bar and powering on its HV, and over time it would usually decay and stabilize to a significantly lower value. Initially it was speculated by David Stuart that this behavior could be afterpulsing caused by gases that had leaked into the vacuum tube of the PMT and needed to be annealed out by HV over time, and that it was an unavoidable annoyance we would need to adapt our testing procedure to.

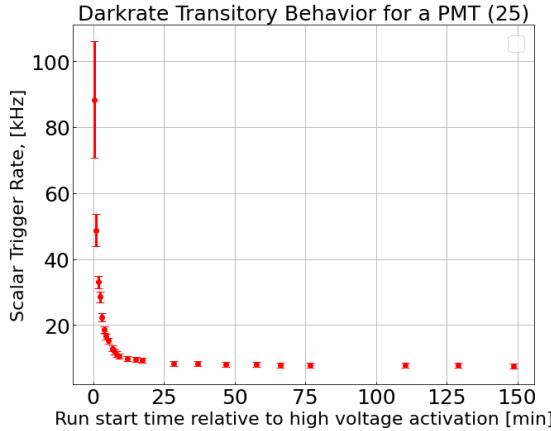


Figure 8.2: Darkrate across many runs taken minutes apart over 2.5 hours, showing a quick exponential decay. This necessitated longer testing procedures.

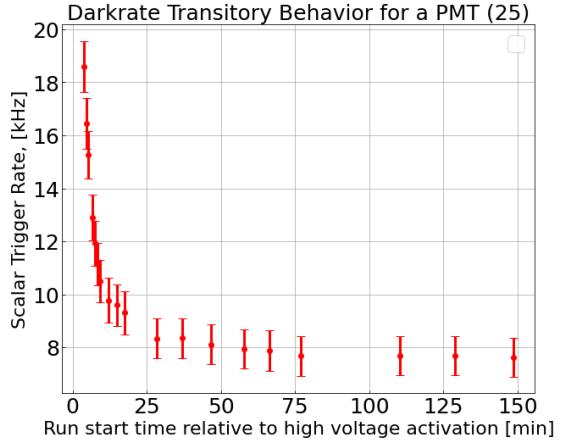


Figure 8.3: Excluding the first few minutes, a better look at the data in Figure 8.2. A slower exponential decay process may be responsible for continued decrease after hours have passed.

The result was that bars had to be powered and left on for up to 15 minutes before their properties were measured, giving time for the dark rate to settle. This significantly increased the time it took to test PMTs, but it was possible to fill that time by wrapping bars. This continued until problems began to get worse, and new behaviors started appearing. The dark rate on some PMTs would rise to a stable level rather than falling, on some it would rise slowly but indefinitely after it had decayed rapidly, and on many the dark rate began to react strongly to the environment around the bar. This was discovered when a section of black felt was thrown over a high-dark rate bar, because it would block potential light leaks and thus bring the dark rate down. Instead, contact with this felt caused the bar's dark rate to skyrocket! Further testing with the felt revealed that this behavior could be replicated even without direct contact, holding the felt suspended inches above the bar. Indeed, even waving ones hand too close to the bar could illicit a reaction.

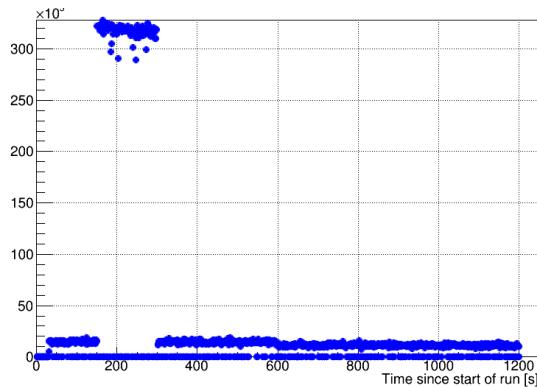


Figure 8.4: Scaler trigger rate in DRS for a PMT taken over twenty minutes, triggering on a pulse generator. Something happened to cause a temporary constant shift up to a significantly higher rate.

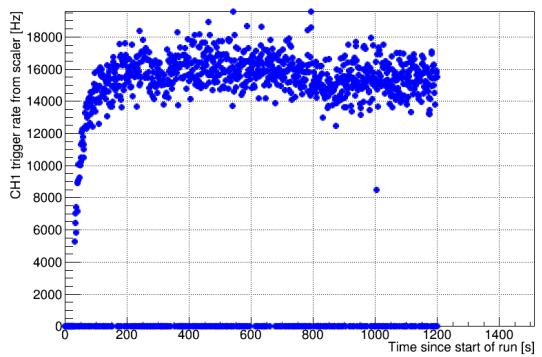


Figure 8.5: Scaler trigger rate for a different PMT over twenty minutes. Rather than decaying, the darkrate of this PMT *grew* to a stable value. This is just one of many anomalies observed in early testing.

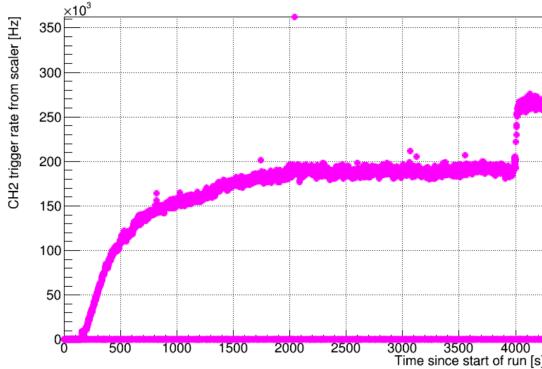


Figure 8.6: The darkrate in *this* PMT grew to a large and stable rate, only to jump to a higher rate when a cloth was gently laid over its bar, demonstrating strange capacitive behavior.

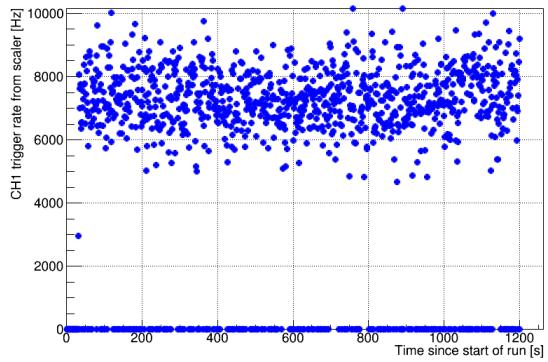


Figure 8.7: Stable PMT darkrate behavior, achieved by reversing HV polarity. Since the -HV issue built up over a long time, it is remarkable how quickly it reversed.

Instability is not good for sensitive equipment in your experiment, especially for a source of background like dark rate. It was critical to identify the cause of this behavior, and the previous theory of residual gases seemed at the very least incomplete. David identified that there was some capacitive behavior occurring in the PMT, where the presence of an object near the bar would cause an electric field that may be interfering with the internals of the PMT, perhaps freeing more photoelectrons than usual. This proved to be partially correct, though I only determined the culprit after investigating the Hamamatsu R878 manual. On page 72, under section 4.3.6: Dark current, the Hamamatsu PMT handbook says:

Some electrons emitted from the photocathode or dynodes may deviate from their normal trajectories and do not contribute to the output signal. If these stray electrons impinge on the glass envelope, scintillations may occur and result in dark pulses. In general, a photomultiplier tube is operated with a negative high voltage applied to the photocathode and is housed in a metal case at ground potential. This arrangement tends to cause stray electrons to impinge on the glass envelope. However, this problem can be minimized by using a technique called “HA treatment”. Refer to section 13.8.2 of Chapter 13 for detailed information on HA treatment. [HA]

Referring to section 13.8.2, we find a section explaining that having a grounded conductive material around the PMT may result in increased dark rate, particularly if there is not a gap of 4mm or wider [HA]. The MilliQan PMTs are surrounded by conductive mu-metal with no gap and are placed in the mounts with no gap, meaning that they may be more sensitive to an external electric potential, like a statically-charged strip of felt, or someone’s un-grounded hand. The section goes on to say:

The above effects of external electric potential can be eliminated by use of the cathode grounding scheme with the anode at a positive high voltage... [HA]

before describing a workaround involving the application of a conductive paint called “HA treatment”. Thankfully, the latter was not necessary for MilliQan. That is because the HV supply given to the PMTs had been negative (cathode), with the anode held at ground, just as the handbook identifies as common but problematic. Thankfully this was not a necessary

configuration for MilliQan, and could be reversed by modifying the bases to accept positive HV, thus inverting the polarity. After discussion with David, this was experimented on with a few bars, including those that had exhibited particularly unusual behaviors in the past. The earlier observed issues vanished almost instantly, and the dark rate of PMTs fell across the board, even for those that had been tested after waiting long periods for dark rate to diminish. I presented on these findings at the collaboration meeting, and it was decided that the HV system for the experiment would be inverted.

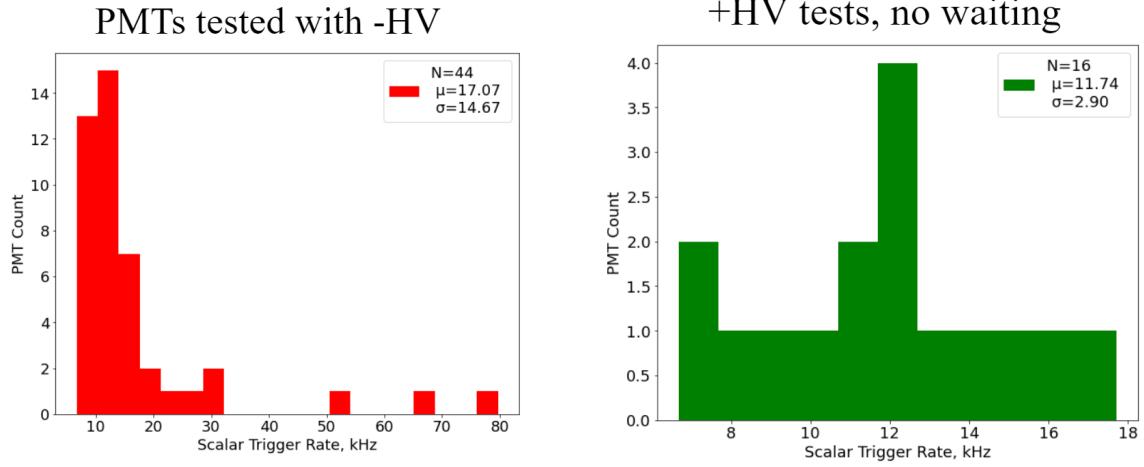


Figure 8.9: Dark rate for 16 PMTs tested with the new positive HV polarity, with no time taken between powering them on and taking data.

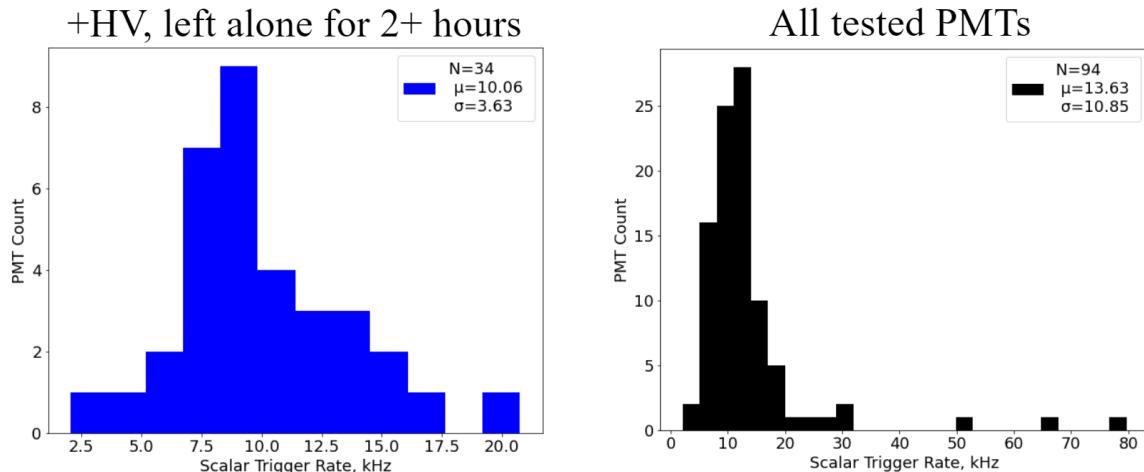


Figure 8.11: The Dark rate of all 94 PMTs tested as of November of 2021, some tested with negative HV and others with positive HV.

This additionally helps the testing procedure, as it is no longer as necessary to wait between powering on a PMT and recording data. Between Figures 8.9 and 8.10 there is a slight improvement when waiting, but it is within one standard deviation of the mean and not statistically significant.

8.3 Optical Gel Study

Often in experiments involving scintillator and PMTs, optical gel is used at the interface between the two. The purpose of this gel is to reduce loss due to scattering, since any air gap between the glass of the PMT and the scintillator will cause refraction. This was deemed unnecessary for the bars, in part because their narrow shape already confines scintillated light well, but it had not been tested physically whether this gel would improve the angular acceptance in the panels of the bar detector. Inspired by simulations performed by Ryan Schmitz, Katia and I sought to compare the PMTs on a panel with and without optical gel, before we performed any panel mounting. We used two narrow cylindrical pucks of EJ-560 optical interface gel, sized to fit their respective PMTs. Two PMTs were mounted at opposite corners of a large slab of scintillator from a different experiment, held on by an older model of mount. One PMT was mounted using gel and the other was not for the first tests, and then the gel was swapped between them for the next set of tests. When each PMT was being tested, a paper marked with a grid of coordinates was placed on the top of the panel in the area around that PMT, with spacing of approximately 2cm to fit a small radioactive source. For each PMT in each scenario, this radioactive puck would be moved across every grid coordinate, and a run consisting of 20k events would be recorded at each of those points. The idea behind this test is that the number of photons generated by an energy deposit from the source (NP/keV) at each point on the bar would be roughly the same, but the light collection efficiency would vary wildly: at some points the scintillated light would less often make it into the PMT. This has huge implications for sensitivity, since the efficiency of a PMT at producing photoelectrons from scintillated photons does not matter if the scintillated photons do not make it to the PMT in the first place. Particularly in the case of the slab detector, the sensitivity of the detector diminishes considerably if an mCP must pass through a slab directly in front of a PMT for its signal to be seen. Light collection efficiency will never be as high for one of the panels or slabs as it is for the bars, but it is important to maximize.

Setup:

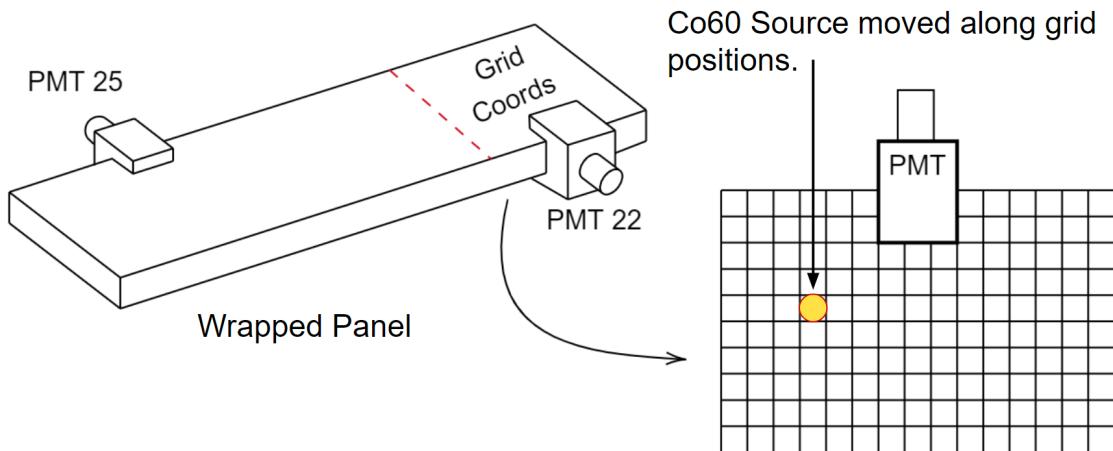


Figure 8.12: A diagram of the setup used for testing optical gel on PMTs for use in the panels and slabs.

Initially the source used for testing was Cd-109, the same used for the bars. This quickly proved unfeasible, however, as the source peak at far positions fell so low that it was overlapping and indistinguishable from the SPE peak. A few other sources were tested, including Cesium-137, Sodium-22, and Barium-133. Finally Cobalt-60 proved most successful, and was chosen for the experiment. The next question was how to quantify the light collection at each point in order to draw a comparison. When using Cd-109 we would measure the NPE produced by its consistent 23 keV gamma rays, i.e. the position of a pulse height source peak, but this method became problematic because the gamma ray spectrum for Co-60 has two close but separate peaks that get merged by low-resolution smearing, and the addition of low-energy Compton scatters means that the resultant pulse height distribution is more of a bump with a long shoulder rather than a Gaussian. Scaling and subtracting a no-source distribution is not applicable here, since the HV/gain of the PMTs needed to be brought down to resolve the Co-60 spectrum, rendering the SPE peak invisibly small. A much cleaner solution was created by Katherine, who ultimately wrote the automatic fitting script. This method was to choose one standard coordinate on the grid (naturally the center) and save a histogram of pulses for the source at that position. Histograms would then be generated with the source positioned at each other position on the grid, and then their x-axes would be scaled up until a best-fit was achieved between them and the distribution at the standard position. This method was successful because every data set contained 20k entries, so the position of the pulse height distribution grew and shrank in proportion to its width, an effect approximated well by simply stretching the graph by a constant factor. The result was a mosaic demonstrating the relative light collection efficiency of each discrete coordinate position across the grid, all normalized to a central point. This gave a clear visual indicator of whether the gel was having a noticeable affect, because one could see visually whether the cone of high acceptance was at a wider angle in the gel tests vs the no-gel tests.

No Gel	7	6	5	4	3	2	1	0	-1	-2	-3	-4	-5	-6	-7
0	0.392	0.465	0.476	0.512	0.555	0.606	0.465	0.555	0.434	0.555	0.512	0.512	0.465	0.425	0.384
1	0.425	0.465	0.526	0.5	0.588	0.588	0.606	0.465	0.476	0.606	0.571	0.512	0.465	0.454	0.416
2	0.416	0.487	0.512	0.555	0.588	0.526	0.769	0.5	0.666	0.645	0.606	0.526	0.512	0.465	0.4
3	0.444	0.465	0.526	0.571	0.645	0.909	1.052	1.052	0.952	0.714	0.588	0.555	0.5	0.465	0.416
4	0.416	0.465	0.555	0.555	0.689	0.909	0.952	1.052	1	0.909	0.740	0.540	0.465	0.425	0.406
5	0.454	0.5	0.540	0.645	0.8	0.869	0.952	1	0.952	0.909	0.8	0.625	0.526	0.476	0.416
6	0.416	0.487	0.555	0.714	0.769	0.909	0.952	1	0.952	0.869	0.8	0.666	0.384	0.454	0.416
7	0.454	0.526	0.571	0.740	0.8	0.869	0.909	0.869	0.909	0.869	0.769	0.714	0.512	0.465	0.408
8	0.406	0.487	0.606	0.714	0.833	0.869	0.869	0.869	0.833	0.833	0.8	0.689	0.571	0.487	0.392
9	0.416	0.5	0.425	0.714	0.740	0.833	0.869	0.869	0.833	0.8	0.769	0.606	0.555	0.425	0.4

Figure 8.13: Scale factors, proportional to light collection efficiency, normalized to a point in the center of the test region for a PMT with no optical gel.

This study ended up taking the longest of any other bench test during my time on MilliQan, for what ended up being fairly uninspiring but nonetheless reassuring results. In Figures 8.13 and 8.14, it is unclear what affect the gel had on the light collection efficiency, if any. In both cases there is a cone of acceptance spreading out from the top of the graph where the PMT is located, excluding the top itself where absorption of radiation in a plastic mount altered the spectrum. To probe for less noticeable differences, each value in the gel grid was

Gel	7	6	5	4	3	2	1	0	-1	-2	-3	-4	-5	-6	-7
0	0.434	0.476	0.512	0.588	0.588	0.645	0.580	0.625	0.588	0.571	0.512	0.540	0.465	0.446	
1	0.434	0.5	0.526	0.606	0.571	0.645	0.645	0.571	0.487	0.606	0.606	0.555	0.5	0.487	0.416
2	0.476	0.5	0.571	0.555	0.625	0.606	0.689	0.645	0.8	0.714	0.625	0.588	0.512	0.476	0.425
3	0.487	0.476	0.588	0.555	0.666	0.666	0.833	0.952	0.714	0.769	0.555	0.571	0.555	0.487	0.444
4	0.476	0.512	0.571	0.606	0.689	0.8	1.052	1.052	1	0.666	0.714	0.571	0.5	0.454	0.434
5	0.444	0.540	0.571	0.5	0.8	0.869	1	1	1	0.909	0.645	0.588	0.512	0.476	0.487
6	0.465	0.555	0.555	0.666	0.8	0.909	0.909	0.952	0.952	0.909	0.869	0.8	0.526	0.625	0.434
7	0.487	0.571	0.588	0.740	0.833	0.833	0.909	0.952	0.909	0.869	0.8	0.689	0.625	0.540	0.444
8	0.487	0.555	0.606	0.714	0.8	0.869	0.869	0.909	0.909	0.833	0.740	0.434	0.555	0.526	0.434
9	0.454	0.5	0.689	0.689	0.8	0.909	0.869	0.869	0.909	0.833	0.740	0.645	0.512	0.465	

Figure 8.14: The same measurements taken with the same PMT, but with optical gel added. There is some visible improvement at the corners of the cone of acceptance.

divided by the corresponding value in the no-gel grid, resulting in a ratio in light collection efficiency at each of those points. This graph (Figure 8.15) reveals that, indeed, the optical gel performed slightly better at very high angles, but roughly the same or even worse for low and medium angles. While useful for future reference, it was determined that the meager improvements at large angles were not worth the added construction challenge of utilizing optical gel, particularly as the efficiency of panels was not important enough to optimize at the level of 10% increases, and because each slab would be covered by four PMTs, with very few if any points at such high angles relative to the PMT.

Gel/NoGel	7	6	5	4	3	2	1	0	-1	-2	-3	-4	-5	-6	-7
0	1.15	1.02	1.07	1.14	1.05	1.06	1.26	1.12	1.43	1.05	1.11	1	1.16	1.09	1.15
1	1.02	1.07	1	1.21	0.97	1.09	1.06	1.22	1.02	1	1.06	1.08	1.07	1.07	1
2	1.14	1.02	1.11	1	1.06	1.15	0.89	1.29	1.2	1.10	1.03	1.11	1	1.02	1.06
3	1.09	1.02	1.11	0.97	1.03	0.73	0.79	0.90	0.75	1.07	0.94	1.02	1.11	1.04	1.06
4	1.14	1.10	1.02	1.09	1	0.88	1.10	1	1	0.73	0.96	1.05	1.07	1.06	1.06
5	0.97	1.08	1.05	0.77	1	1	1.05	1	1.05	1	0.80	0.94	0.97	1	1.17
6	1.11	1.13	1	0.93	1.04	1	0.95	0.95	1	1.04	1.08	1.2	1.36	1.37	1.04
7	1.07	1.08	1.02	1	1.04	0.95	1	1.09	1	1	1.04	0.96	1.21	1.16	1.08
8	1.19	1.13	1	1	0.96	1	1	1.04	1.09	1	0.92	0.63	0.97	1.07	1.10
9	1.09	1	1.62	0.96	1.08	1.09	1	1	1.09	1.04	0.96	1.22	1.16	1.20	1.16

Figure 8.15: The values in Figure 8.14 divided by their corresponding values in Figure 8.13, demonstrating the gain factor achieved by adding gel. The gel helps a little bit in peripheral regions, minimally in the forward region, and not enough anywhere to justify the installation challenge.

8.4 Light Collection in Bars

The natural question that arises after observing the large differences in light collection efficiency across a slab is whether there are major differences along the length of a bar. If so, that would need to be accounted for more carefully in projections, and the exact point at which a source is placed during calibration would need to be kept consistent. I performed a quick test of this behavior by placing a Cd-109 source on a bar and measuring its NPE at 3cm intervals across the length of the scintillator (Figure 8.16). The results of this test are shown in Figures 8.17 and 8.18.

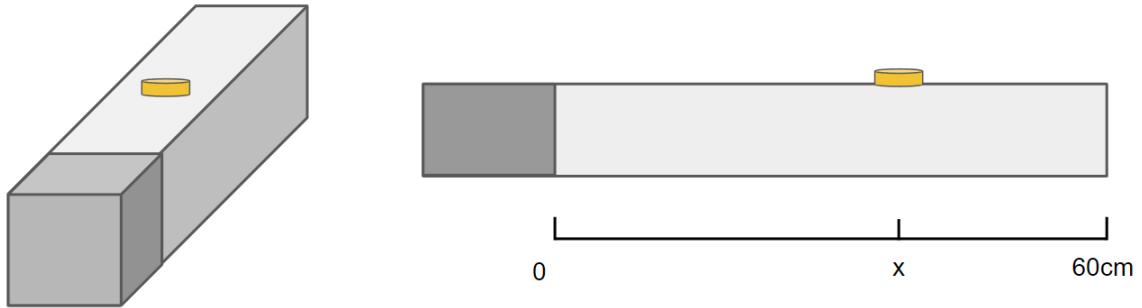


Figure 8.16: The simple setup used to test light collection efficiency along the length of a bar, this time using a Cd-109 source.

We see that the observed activity of the source was maximized along the bulk of the scintillator and lowest at the interface, likely due to absorption and scattering from the thicker wrapping of the mount, and due to the higher angle relative to the bar. The NPE, however, was maximized near the interface, because the (relatively fewer) deposits actually observed by the PMT somewhat expectedly had a higher proportion of scintillated light making its way to the PMT, as there is less space for it to scatter out or be absorbed. This advantage quickly diminishes at distances even a few cm further along, with NPE being fairly consistent across the entire bulk of the bar. Because both the activity and NPE were consistent in this region, with the activity being highest (and thus an NPE peak easier to resolve and fit), it was decided that all bar source calibration would be done with the source placed along the bulk of the bar, and that hairs would not be split over its exact position.

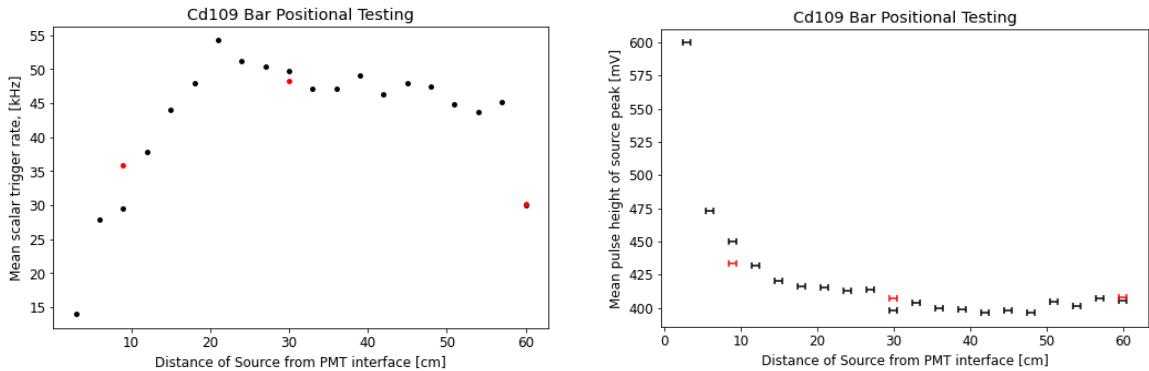


Figure 8.17: The Scaler trigger rate of the PMT when a source is placed at different positions along the bar, activity which is dominated by radiation from the source and not from dark rate. Interestingly, the most activity is found 20cm away from the PMT, indicating more scintillated light is reaching the PMT overall at that angle instead of scattering out.

Figure 8.18: Fitted pulse height peak for the source at each of its positions taken along the bar. Unlike in Figure 8.17, the source peak appears to be maximized at the lowest possible distance from the PMT. While fewer scintillated photons are reaching the PMT overall, they are doing so in bigger clusters. This makes sense, as lower angles are more likely to reflect.

8.5 CMS B-field Study

Calibration of bars becomes more complicated after moving them underground to the gallery, because PMTs are highly sensitive devices that rely on the trajectory of photoelec-

trons passing through a vacuum, and the paths of charged particles tend to curve in the presence of a strong magnetic field (aka B-field) like the one generated by CMS when it is powered on. This is an effect that can be observed in R878 PMTs even on the surface, due to the magnetic field of the earth [24]. This effect is not unknown to Hamamatsu, which is why the glass body of a PMT is almost always wrapped in a shroud of magnetic-field-reducing *mu-metal*, including the R878s. This is sufficient for nominal magnetic field levels, but most standard PMTs are not intended for use in environments as magnetically potent as the tunnels around CMS. For this reason, it was essential to test whether the CMS B-field had an effect on measured properties of the PMTs, how this effect varied with the orientation of PMTs in the gallery, and whether a certain position was more optimal than others at avoiding these potential effects. The PMT is, after all, not symmetrical, and there is some choice as to how it is rotated within each bar and how the bars are rotated (at 90° intervals) within each unit.



Figure 8.19: Me standing beside the support structure, on which a single unit is placed for angular B-field testing.



Figure 8.20: In order to measure the angle of rotation for units in this study, the built-in magnetometer of a smart phone was used.

This question had already been tested by David Stuart at UCSB using Helmholtz coils to produce a static magnetic field at a similar magnitude to the one felt in the gallery, and in fact it was his measurements that were used to select the orientation of each PMT within its bar. But this was a bench test and was made under the assumption that the bars would be oriented with their PMTs mounted on the far side of each bar from the IP, when in the final design that orientation is rotated 180 degrees so that the PMTs are on the near side. To obtain more certain results during commissioning at CERN, bars and units were brought down to the gallery, where they were positioned on the angled detector support structure, given power, and read out to the DRS. Each bar or unit was then rotated through intervals of a certain angle along its axis and two runs were taken at each angle, one with and one without a Cd-109 source. This was repeated over two full rotations of the unit, accounting for overall drift in performance and reducing the effect of human error. The starting angle (0°) was the orientation for which the cable gutters of the units faced upward and the LV ports of each base were facing downward, which was the position 90° between each of the two final planned orientations of units/bars in the detector, corresponding to the left and right sides. Once these tests were complete, each unit or bar was then wrapped tightly with an *additional* layer of mu-metal so that it covered the mount and interface, and the tests were taken again. This way it could be determined whether any negative effects could be reduced by additional mu-metal wrapping. Gain, darkrate, and NPE were compared across

four bars in a unit at angle intervals of 90° (Figure 8.22), and gain alone was measured more precisely for a single bar at angle intervals of 15° (Figure 8.21) without a comparison with additional mu-metal.

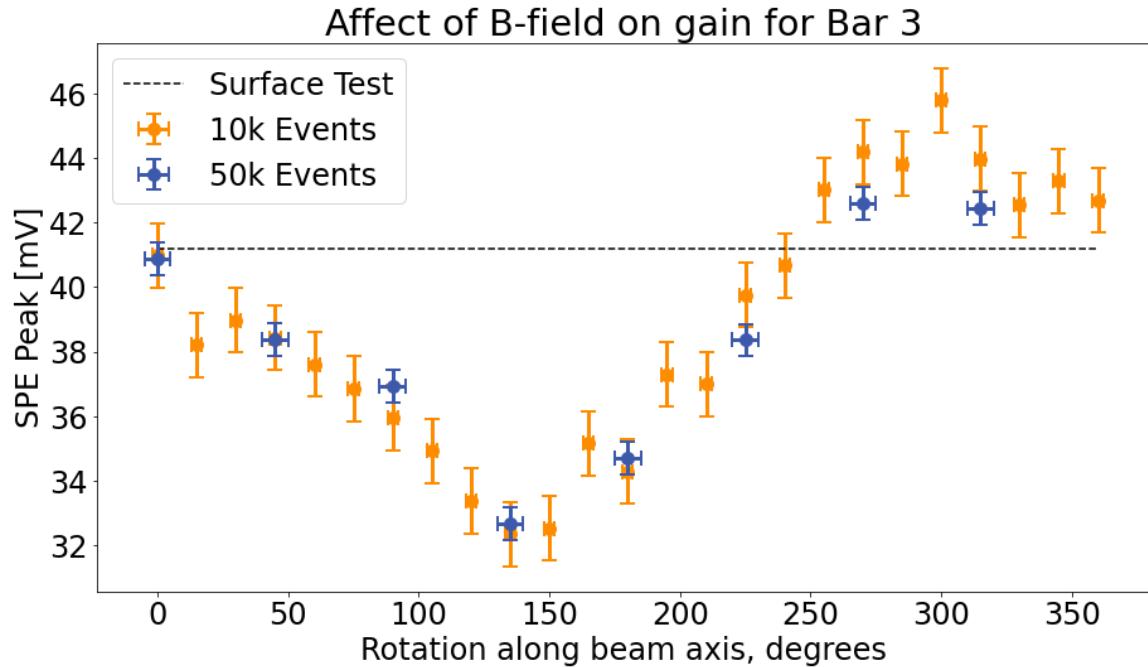


Figure 8.21: Gain measured for a single bar rotated through 15° intervals taken at 10k events each, and then again at 45° intervals taken at 50k events each.

Looking first at the fine-angle gain measurements for bar 3, an angular dependence becomes immediately clear. Gain appears to be maximized at $+270^\circ$, which is the position of bars on the left side of the detector, and if we assume a pseudo-sinusoidal distribution due to symmetry, that would indicate another maximal peak at $+90^\circ$. In the plot this point appears sub-optimal, though this may be due to long-timescale drifts in PMT performance, which have been observed in MilliQan bars. Regardless, it appears that David's selection was still wise, and fairly optimized among the possible rotations for minimizing the negative effects of a B-field. This is just a measure of gain, however, which says little to nothing about NPE/keV, a more critical measurement.

Turning our attention to the performance in bars 1, 47, 120, and 122, we get some mixed results. The same pseudo-sinusoidal angular dependence is repeated consistently across each parameter in each bar, but there is not abundant evidence in favor of any quantity having changed significantly with the addition of extra mu-metal. While the mu-metal data points appear larger in most cases (not good in the case of dark rate), this may also be due to drift, which is somewhat supported by the variation of these measurements from the value measured on the surface. That said, it does appear from these results that adding mu-metal does *stabilize* the performance of each bar somewhat, with a decreased magnitude in angular variation across the board.

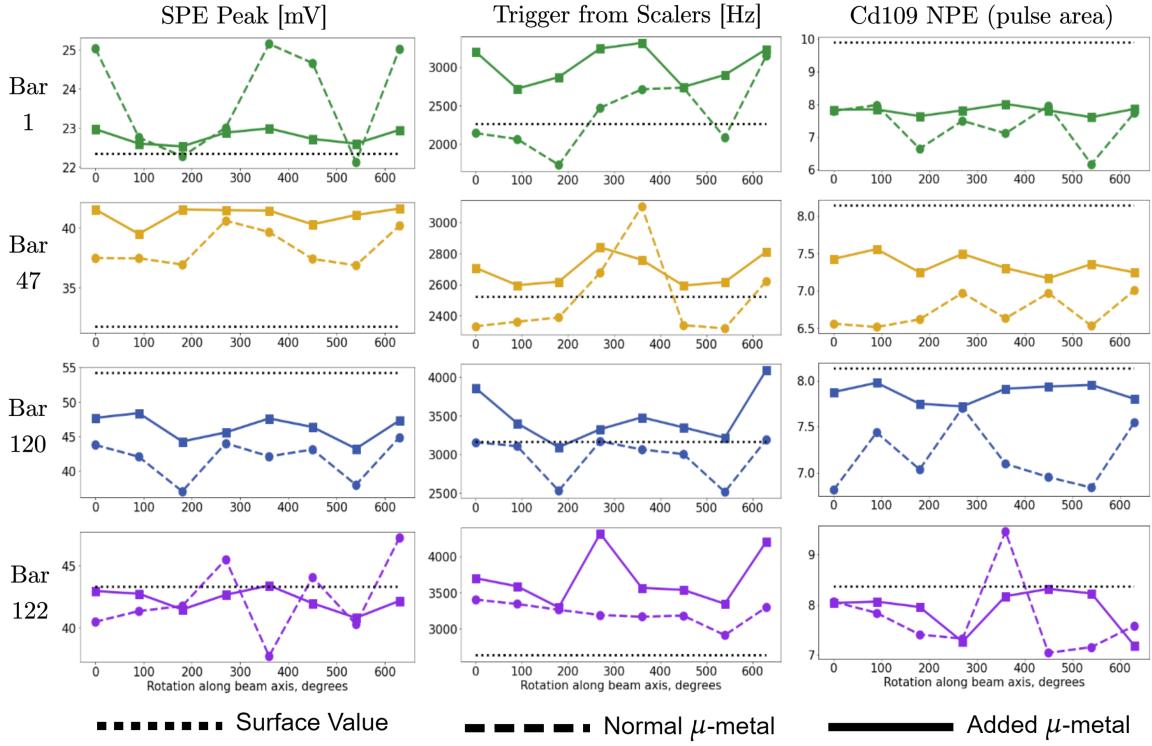


Figure 8.22: Gain, dark rate, and NPE measurements taken for four different bars at 90° intervals in the CMS B-field, pointed toward the IP. Runs are taken with and without additional external mu-metal, and compared to their value at the surface.

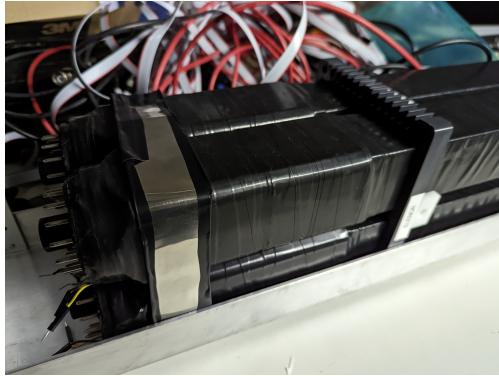


Figure 8.23: A thin strip of mu-metal is used to probe the position in most need of additional wrapping.

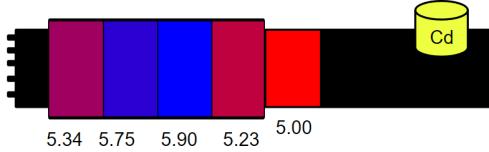


Figure 8.24: The NPE measured of a Cd109 source peak when the thin strip of mu-metal in Figure 8.23 is covering each of the colored regions.

After determining that additional mu-metal may benefit the detector, the question becomes how much is necessary and where it should be applied. To determine this I cut a narrow band of mu-metal (Figure 8.23) and positioned it at five locations on and near a bar's mount, and took data using Cd-109. Figure 8.24 shows the resultant NPE measurements for each mu-metal position, with the maximal location being right over the *interface* between the PMT and the scintillator, with decent results also found at the position covering the bulk of the PMT's dynodes. All results, however, appeared to be reduced by 15-25% from the NPE observed in the bar with full coverage, so it was deemed preferable to cover the mount in its entirety. Thankfully there was sufficient mu-metal to achieve this, so every unit was wrapped in an additional layer before being installed in the detector (as mentioned in Chapter 6).

8.6 Unusual Saturation Issue

In spring of 2023 I began to notice an unusual behavior in certain bars when searching for cosmic showers in data from the almost-complete bar detector. Cosmic showers generate pulses large enough to saturate pulse height, so I was placing cuts on pulse height around the saturation level (1250mV) in order to isolate events in which there was one of these large hits. To my surprise, some of the channels contained no events whatsoever with pulse heights above 1200mV, despite the fact that they were working normally otherwise. Analyzing the pulse height distributions for these particular channels, a consistent deviant behavior was identified: their overflow bins had moved. Instead of a sharp peak around 1250mV, these channels instead had a broader peak around ~ 1080 mV, varying slightly between channels.

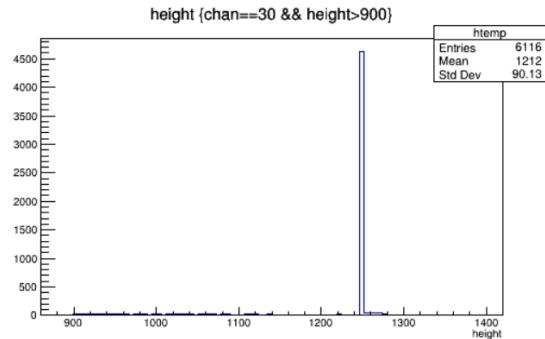


Figure 8.25: What the saturation bins of a MilliQan bar should look like, sharp at 1250.

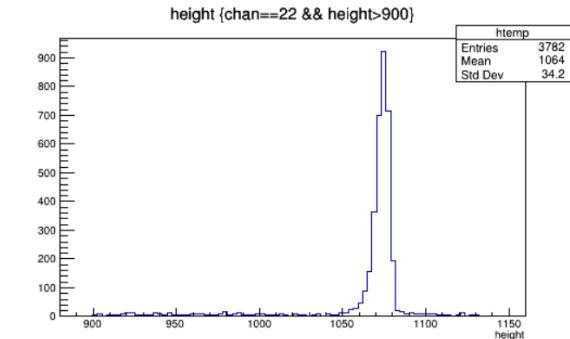


Figure 8.26: The anomalous peaks of some channels, at ~ 1070 mV with lots of variation.

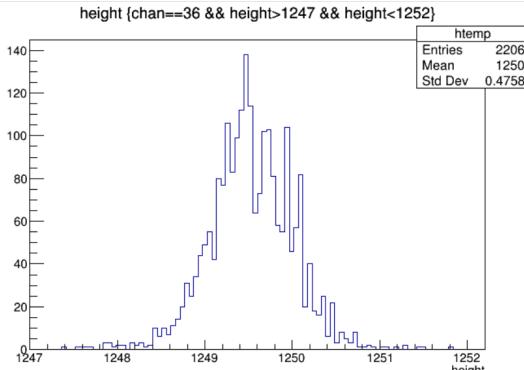


Figure 8.27: A closer view of 8.25.

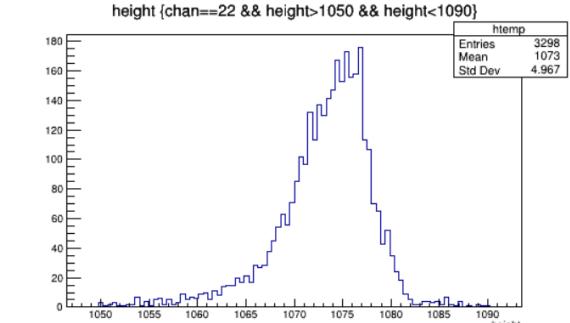


Figure 8.28: A closer view of 8.26.

8.6.a Searching for Cause

The initial instinct of Matthew Citron, David Stuart, and Neha Santpur was that this may be caused by a downward shift in GND for those channels and an increase in noise, two quantities reflected in the Offline variables “sidebandRMS” and “dynamicPedestal”. Checking these variables across a “normal” and “problematic” channel, however, revealed no noticeable deviation to explain the differences. It was not merely an issue with data being saved either, as viewing event displays in those channels revealed saturation of those channels at the reduced height.

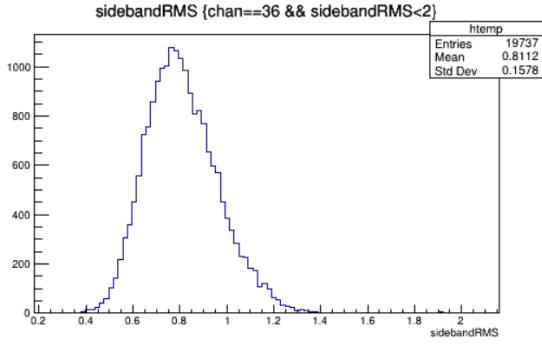


Figure 8.29: Sideband RMS for a PMT with typical saturation behavior.

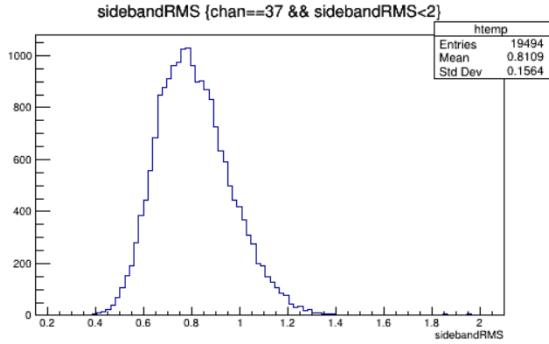


Figure 8.30: Sideband RMS for a problematic PMT, with no clear difference.

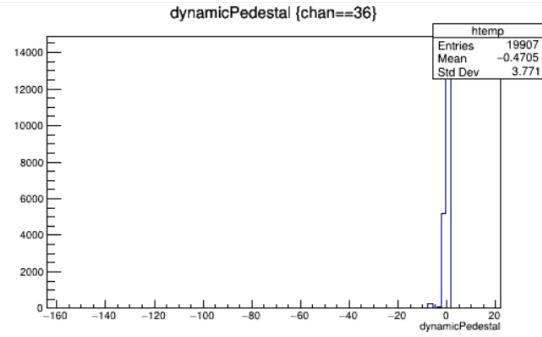


Figure 8.31: Dynamic pedestals for a PMT with typical saturation behavior.

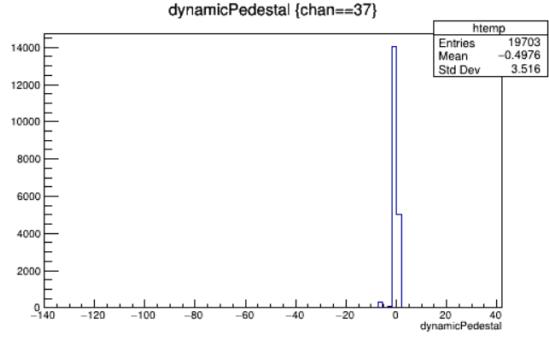


Figure 8.32: Dynamic pedestals for a problematic PMT, with no clear difference.

8.6.b Issue Tracking

While there remains no clear explanation of this issue or how to fix it, I have been monitoring it across recent runs. The saturation appears to occur *almost* exclusively in Supermodule 2, with the puzzling exception of a few channels that behave normally inside that supermodule and a few that behave oddly outside of it. This supermodule utilizes all of the same electronics as those that function nominally, and there is no explanation for why some of its channels are unaffected, but it stands to reason that there is a systemic cause corresponding with the supermodule. One channel, 53, appears to have recovered from this behavior between February and May of 2023, and there are other examples of this variation.

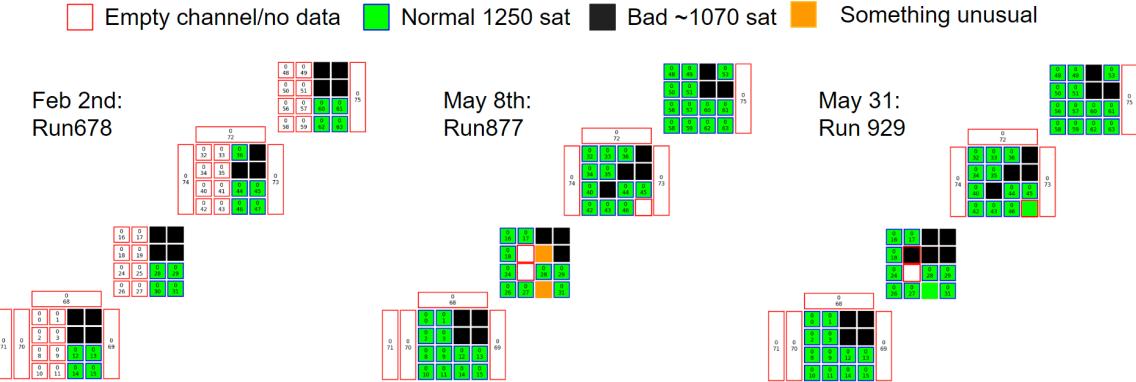


Figure 8.33: A map of normal (green) and anomalous (black) channels in the detector over time, tracking the issue.

Discovering this issue inspired me to take a closer look at much older runs, and while it was not widespread I did discover the exact same behavior in channel 0 during select runs starting in November. Below is an exact timeline of behavioral changes in channel 0 since October 2022:

- **Runs 469-546 (10/4-11/3):** Normal Behavior
- **Run 547 (11/8):** Dead run, no data in any channel.
(note, run desc says “[test] swap signal cables between layer 1 and layer 2 for the 1st installed supermodule more details” starting in this run and going to 552. Assuming the correct conventions are being used, chan 0 is in layer 0 so should not be directly affected)
- **Runs 548-552 (11/8):** Saturation issue appears in chan 0, looks similar (~1080ish), maybe higher, but not much data.
- **Runs 553-583 (11/15-11/24):** Issue gone in 0, saturation back at 1250mV.
(note, log says for Run552 “:[test] normal data taking with all signal cables back to correct position”)
- **Runs 584-591 (11/24-11/25):** Issue back in channel 0, 1080mV saturation
(note, nothing in log indicates change in 584: just says “test”)
- **Runs 592-594 (11/25):** Channel 0 totally empty, dead.
(note: for Runs592-594 channel 0 had LV off/connected to 0 pulse gen)
- **Runs 595 (11/25):** Channel 0 only has peak around ~130mV, but lots of events.
(note: for Runs595 channel 0 replaced with pulse gen at 20ns 2.5MHz pulse)
- **Runs 596+ (1/16-Present):** Channel 0 back to normal, 1250 saturation.
(note: channel 0 plugged back in)

No other channel displayed these issues during these runs, and channel 0 has not encountered the issue ever since. So in total the issue arose twice for channel 0: once when signal cables were swapped between layers (i.e. a major physical intervention) and a second time due to a mystery cause between runs 583 and 584, which were taken 20min apart (so someone was physically down there when the issue occurred). Each occurrence of this issue

appears to have been fixed after some kind of reset, i.e. the signal cables were swapped back to their original positions and the detector was presumably power-cycled; and when channel 0 was disconnected, swapped with a pulse gen, then reconnected.

My hypothesis is that this issue arises when something (likely DAQ or cable-related) is altered physically, likely by accident, and that it can be fixed by “turning it off and turning it back on again”. This theory is only partially spoiled by the fact that there was a major intervention on Supermodule 2 between May 8th and May 31st that involved recabling, so I would expect the effects to have disappeared if that was sufficient for solving the issue.

A cause has not yet been successfully identified, but neither have any serious detriments resulting from this unusual behavior. For this reason the anomaly is considered a low-priority issue for the collaboration, and the mystery behind it may not ever be fully understood.

8.7 Cosmic Shower Radiation Penetration

“Cosmic showers” are one of the largest backgrounds for the detector, so it is important to understand how effective the panels of the detector are at shielding the bars. Furthermore, it is useful for background modeling to understand how well each layer of bars shields the layer underneath from cosmic shower particles, and thus how deep into the detector radiation tends to penetrate. There are so many ways we can do this, but there are also so many different ways that radiation can interact with the detector. A single study filtering data will not be able to account for the many possible trajectories of radiation through the detector, but thankfully CERN’s GEANT4 simulation software can. This means that we can choose a very specific way of building our studies, and then compare the results we find with data to those that are found by running the same study on simulated data. If they agree, this helps to validate the simulation and broaden its scope to account for more detail than we can. That being said, it is still important to craft a search that aligns well with the Physics we are interested in. For one, we only want to look at events in which something very energetic hit the detector (like a muon or X-ray), and we want that particle to be traveling roughly downward. This can be achieved by filtering events to require that a large pulse appears in one of the two top panels of the detector. In order to increase the correlation between examined bar hits and panel hits, we will set the requirement on each of our top panels one at a time, and only look at hits in bars that are physically covered by that panel. We do not then expect for all of that energy to be transferred to the lower rows of bars, instead there will be scattering and energy loss resulting in smaller pulses. For this reason it is sufficient to count all hits in lower bars simply with the requirement $nPE > 1$, as that will cut out most dark rate background. It may interest us to know what this nPE distribution looks like in each row of bars when we require a panel hit, so I have plotted that data in Figure 8.34. The topmost row took the greatest number of hits, as we might expect, but thankfully there is no obvious difference in nPE distributions when we normalize counts across each row of bars.

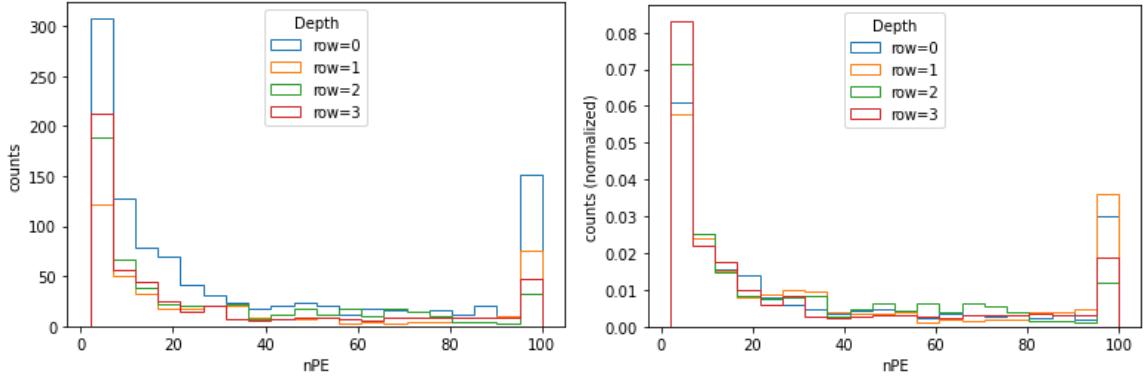


Figure 8.34: A histogram of $nPE > 1$ for pulses in each row of the detector, requiring that the panel above had a pulse with area $> 16.5k$ nVs. Left graph is total counts, right graph is normalized. Once normalized there is no obvious difference.

In summary, we seek to measure radiation penetration in the bar detector by only looking at events with hits above a large threshold in one of the two top panels, then counting the $nPE > 1$ hits in each of the rows of bars hidden directly beneath that panel. The question then becomes, what should be the “large threshold” we set in the top panels? When I conducted this study nPE was not yet calibrated for these panels, so instead I looked at a variety of pulse area cuts, as pulse area is directly proportional to nPE . Modifying this threshold area, I followed the procedure and counted $nPE > 1$ in each row of bars (as well as the total hits in the bar), the results of which are displayed in Figure 8.35. For low area thresholds we see that there is a substantial dropoff between the size of the hit in the panel and the number of $nPE > 1$ hits in the first row of bars, and then each successive row after that, but as we increase the area threshold this story changes. Once we reach area $> 20k$ nVs, we only have 20% fewer $nPE > 1$ counts in the first row of bars than we have hits in the panel.

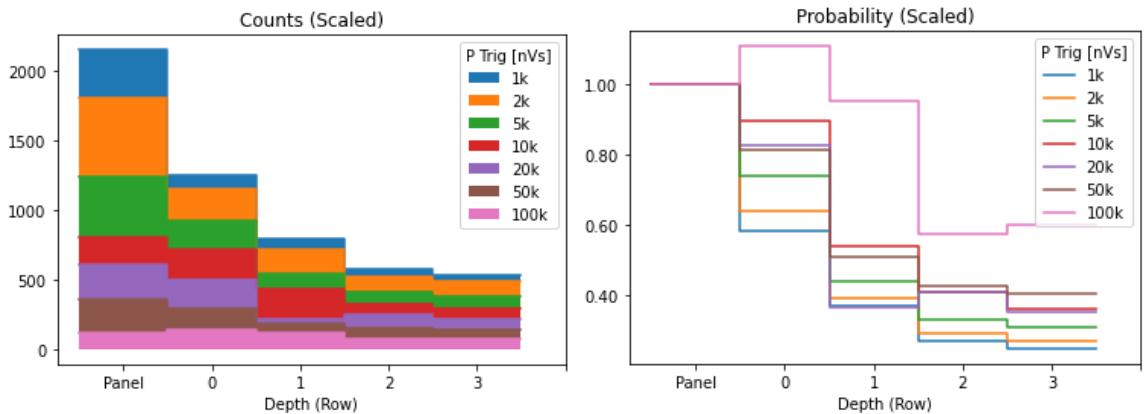


Figure 8.35: The probability of getting a hit in each row depth of the detector, in events with area above each of the thresholds in the legend. Left graph is total counts, right graph is normalized.

Which of these area thresholds we focus on depends on exactly what kind of penetrative radiation we are looking for, at which point there are compelling physical arguments for how much energy each type of particle would deposit, depending on its characteristics and origin. An interesting way of approaching this challenge is in fact to simulate this study, program exactly what type of radiation will rain down on the detector, and then tune the

area cut on data to see if those counts can be matched. Thankfully, resident MilliQan simulation expert Ryan Schmitz conducted exactly this study in GEANT using gamma rays and electrons. Figure 8.36 displays the probabilities that Ryan obtained through his simulations, directly adjacent to the probabilities I obtained from data when setting two fine-tuned area cuts: between 2k and 16.5k nVs area in the top panels, and between 16.5k up to 600k nVs area in the top panels (below 2k nVs is SPE range).

Data		Ryan's Sim gamma/e	
2k-16.5k	16.5k-600k	γ	e
P: 100%	P: 100%	P: 18.3%	P: 100%
0: 52.6%	0: 82.4%	0: 51.2%	0: 82.4%
1: 31.9%	1: 47.7%	1: 44.3%	1: 48.7%
2: 22.5%	2: 38.5%	2: 38.4%	2: 38.5%
3: 22.0%	3: 33.3%	3: 31.9%	3: 32.2

Figure 8.36: My data, taken with two contiguous ranges of area cuts, compared against Ryan's simulation of energy absorption for photons/gamma (γ) and electrons (e).

At first it is natural to question the legitimacy of this method for validating simulation, since it is blatantly fine-tuned to match the simulation! But we should not dismiss these results, because there is nuance to this fine-tuning. When the range of allowed areas in the panel is changed, two numbers are being tuned as inputs: the minimum and the maximum. All we are really controlling is which region of the radiation spectrum we are looking for in our panels. The result, however, is an output of four separate probabilities that all align surprisingly well with simulation! My point is that each of these data points are not being individually tuned, so it is extremely promising that in any chosen area cut we can obtain four relative probabilities that agree so well with the simulation in a preliminary search like this.

8.8 Direct Muon Hit Coincidence

We have examined the likelihood of cosmic shower radiation getting through panel shielding and into the detector, so the next question is how these rogue particles affect our search for mCPs. Recalling the discussions of Chapter 5, Section 3.a, the cosmic ray background can really become an issue when it coincides with random dark rate. The particular nightmare scenario is that a cosmic muon misses the detector, but scatters and radiates low-energy gammas or electrons that hit three colinear bars of the detector, and that a fourth colinear bar has a random dark rate event at roughly the same time. As a result, four bars in a row in different layers would experience SPE pulses near-simultaneously, which is what we expect an mCP to look like. This is not a likely scenario by any means, but it is one of the most plausible ways that a cosmic shower can mimic signal, and the low probability of any scenario is being multiplied by the billions of events that will be examined. The task set out for us is thus to determine just how likely a scenario like this is, and just as in the previous study we can use our results to validate a far more comprehensive

simulation. Just as before there are many ways of doing this, so it is important to frame the question well: If a cosmic muon passes by one location (say, just shy of the detector), how often do we find coincident low-energy radiation in the areas surrounding it (such as in the detector)? A great strategy for doing this was proposed by Ryan, who suggested that I require a muon *directly* hit the detector in one layer, and then see how often a single small hit appears in each of the other layers of the detector. That information can then be extrapolated to determine the scenario where the muon instead just barely *misses* the detector, and of course that is best handled by well validated simulation.

8.8.a Tagging Muons

The first step to a study like this is to decide what will qualify as a satisfactory muon hit when looking through detector events. This may sound trivial at first, since muons themselves pass unimpeded in a straight line through the detector, leaving behind huge pulses. Complications arise, however, due to the many angles at which these muons can approach the detector. These muons come primarily from above, since they are produced in the atmosphere, but a study by Andy Haas and Mukesh Ghmire at New York University confirmed that they still have approximately the same $\cos^2 \theta$ angle distribution as they do at the surface [36] [34]. This means that they are most likely to travel straight down through the detector and significantly less likely to pass horizontally through it, but there are many that will travel at angles between these extremes. This makes it challenging to tag events with muons, because there are many different detector geometries that could indicate one. It is possible to write an algorithm that considers a large number of these scenarios, but this is not practical to run on every event when there are millions being considered. I identified three potential fast filter methods, and compared each of them to determine which was most promising.

The first method was to require down-going full-column hits, like the one shown in Figure 8.37. This produces events that are very easy to understand, but it fails to tag any scenario where a down-going muon travels at more extreme angles from the vertical.

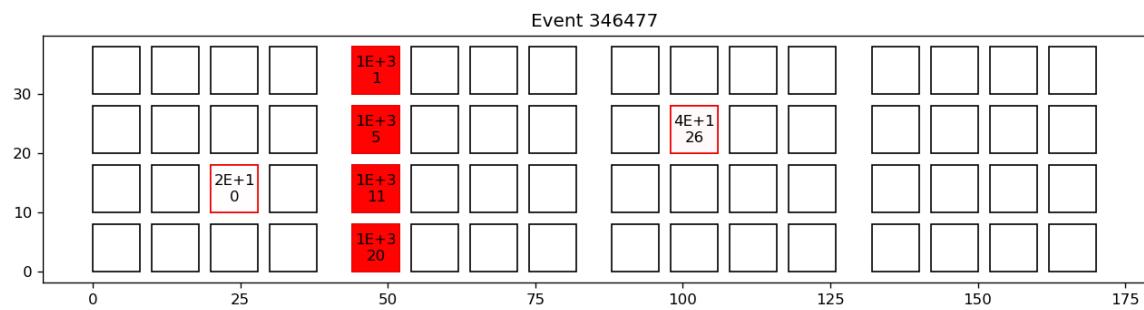


Figure 8.37: A simplified python-generated event plotter for the detector, with the same convention as Figure 8.4. This event has four saturated vertical hits in layer 1.

The second method was to instead require a hit in every row of at least one layer of the detector, regardless of column. This allows for events like the one shown in Figure 8.38, but also allows some nonsensical geometries like 8.39, which is less likely to be a muon.

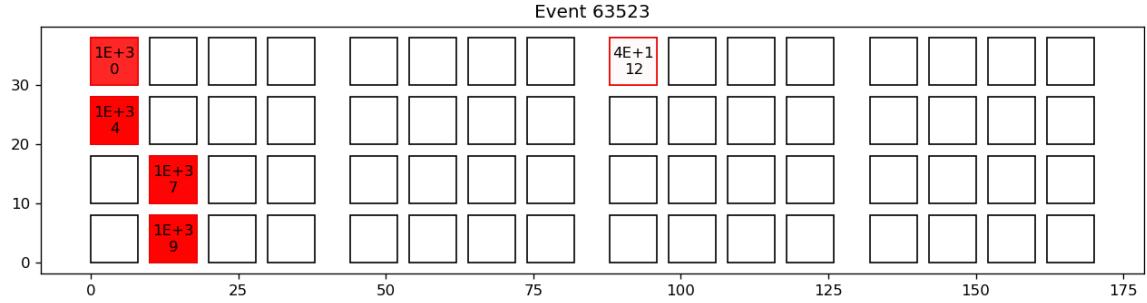


Figure 8.38: This event has a saturated hit in every row of layer 0, suggesting the trajectory of a downward-traveling particle moving at an angle from the vertical.

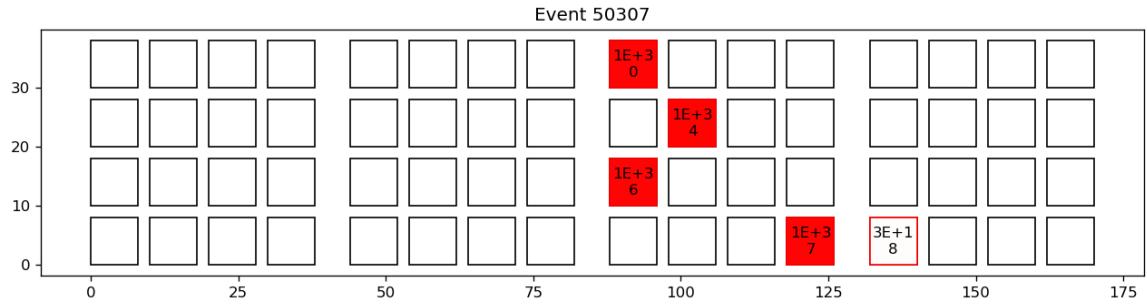


Figure 8.39: This event satisfies the same one-hit-per-row requirement as Figure 8.38, but the position of the hits makes it unlikely this is due to the trajectory of a single particle.

The third method was motivated by the observation that cosmic rays usually leave hits in a great number of bars, often hitting every active bar in a layer. The filter, therefore, requires a hit in every channel in a layer of the detector, as shown in 8.40. This method will not falsely identify any events as being due to muons, but it is obvious that there are muons it would miss, which do not hit every channel in a layer.

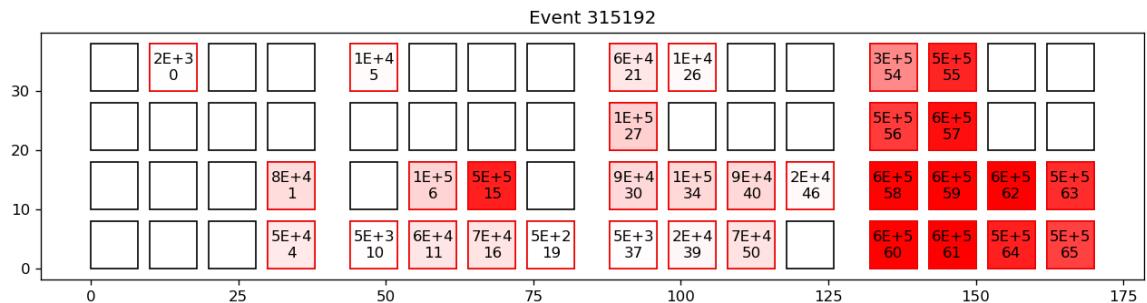


Figure 8.40: An event with every channel saturated in layer 3. The top right corner of each layer is empty because that supermodule had not yet been inserted by this run.

None of these three methods are perfect, but they are all simple and easy to apply to both data and simulation. The task becomes determining which will be the most effective, by analyzing characteristics of the events that they tag. If we look at the muon event shown in Figure 8.5, we see that the area of the cascading pulse created is enormous. This contrasts it with high energy gamma hits, like the one shown in Figure 8.41, which will have the same geometries as muon hits but deposit far less energy. It is therefore a good idea to analyze the area of pulses in events tagged using each method, to see if anything stands out. This

is done in Figures 8.42-8.44, showing area spectrums culminated from every event tagged in a particular run by each method.

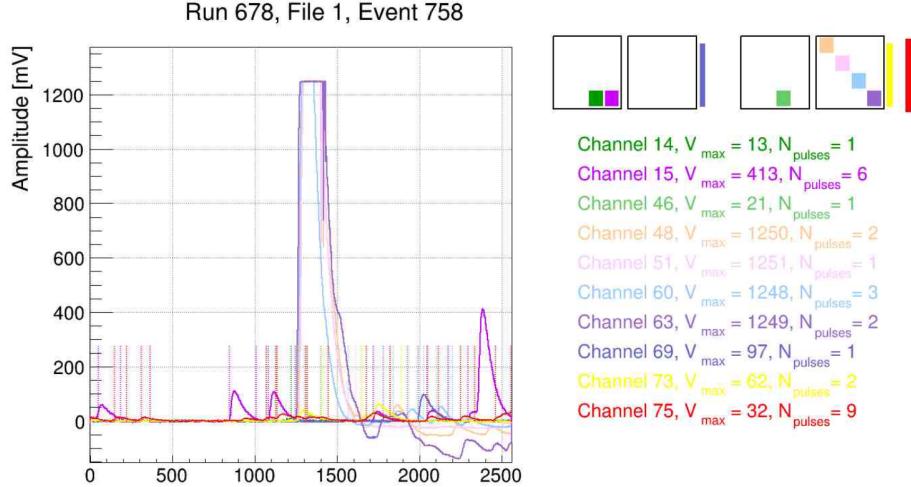


Figure 8.41: A typical “clean gamma” event display, with a very diagonal trajectory through layer 3. Event displays courtesy of [33].

Like pulse height, area distributions have an SPE peak at a fairly low value (typically 1-2 nVs for area). In Figure 8.42 we have not applied any of our geometric methods, but we have cut out these lower area pulses as a control for area distributions outside of the SPE range. In Figure 8.45 we have instead looked at events tagged by muon method 1, with no area cut. The distribution appears to have a small part of the SPE peak in its first and largest bin, likely from darkrate and afterpulses that coincided with tagged events, but then it has a bimodal shape after that with a wide first peak around 100k nVs (still a large area) and a narrow saturation-like peak around 600k nVs. While these peaks are dominant, the bins between them are certainly not empty. If we then apply our second method instead, as we see in Figure 8.44, we get a similar distribution only now with a much more dominant first peak and a less prominent second peak. Looking lastly at Figure 8.45, the balance of these peaks has switched and now the saturation peak is dominant.

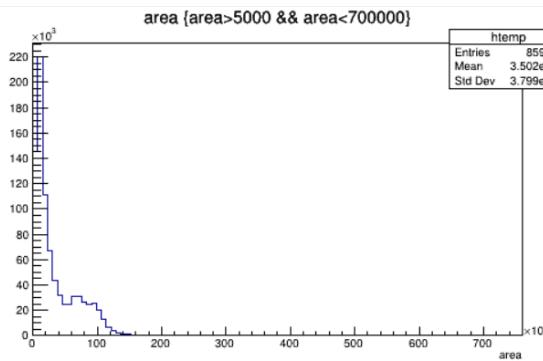


Figure 8.42: Pulse area with only a cut on the SPE peak. No visible saturation bin at this scale.

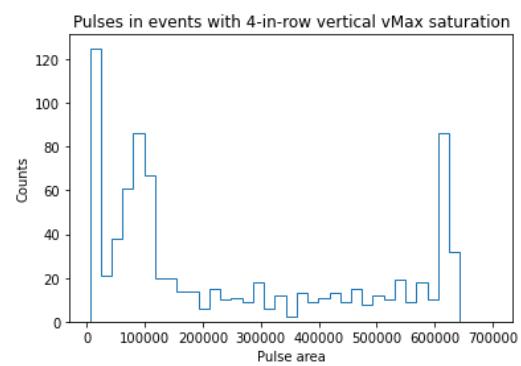


Figure 8.43: Pulse area for four-in-row events. Bimodal with saturation peak.

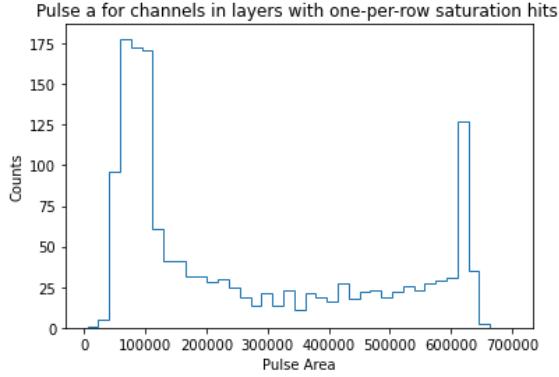


Figure 8.44: Pulse area for one-hit-per-row events. Bimodal, with larger first peak and more events overall.

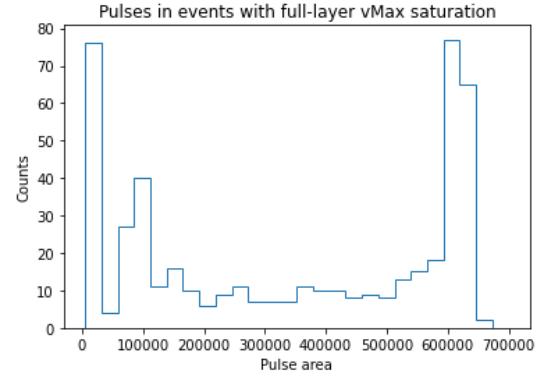


Figure 8.45: Pulse area for one-hit-per-row events. Bimodal, with larger second peak and fewer events overall.

While it is difficult to make rigorous physics conclusions based on these histograms, there is plenty to speculate about. Figure 8.45 makes perhaps the most sense, as any cosmic shower capable of hitting every channel in a layer is likely to deposit a lot of energy, more often saturating at ~ 600 k nVs as we see. Looking at the other two methods we see a lot more pulses with area in that ~ 100 k nVs peak, with more in the one-hit-per-row method than in the four-in-line method. Perhaps this lower peak is due to high-energy gammas, and the tail of that distribution overlaps the tail of a muon distribution which peaks above the saturation voltage, and is thus dumped into the overflow bin. Perhaps the one-hit-per-row filter is more inclusive of these gammas, because its loose geometric requirements are more likely to pick up hits from multiple particles in a shower going in different directions, rather than a solitary muon as we are probing for. That being said, it is still possible that many of the 100k nVs pulses result from direct muon hits, and we can see from the bin counts that the one-hit-per row method is the most inclusive of events in general, including those over 600k nVs. To see how cuts on area effect the data we turn to Figure 8.46, which shows the number of *events* (not pulses) tagged by each method as we increase the size of an area cut that requires these events have *at least one* pulse with an area above the specified value.

	No area cut	>80k area cut	>120k area cut	>400k area cut	>600k
Only area cut	347000	92747	11387	339	151
1-per-row cut	265	256	184	143	80
4-in-line cut	117	115	98	84	60
Full layer cut	27	27	27	26	23

Figure 8.46: Event counts for an intersection of geometric and area cuts. Area cuts only require *at least one* pulse in the event with an area above the specified value.

Looking at each method we see that the one-hit-per-row requirement is the most inclusive across all area cuts, including very large ones that almost certainly indicate direct muon hits. We also see the total number of events that satisfy each area cut alone quickly converge toward the number of events satisfying our requirements as the area cut is bumped up to 400-600k nVs, and that the number of events satisfying each geometric tagging method are effected by less than an order of magnitude at most. A sensible question would be whether

requiring an area cut alone would be sufficient at identifying muon hits, but this chart shows that such a method would be far too inclusive in the $\sim 100 - 400$ k nVs range, and not drastically more inclusive in the > 600 k nVs extreme. It also may include unwanted readout electronic anomalies, absorption of high-energy particles from cosmic showers like electrons (which are not muons and are likely to be absorbed before passing through the entire detector), muons originating from LHC collisions that are a different background entirely, and glancing blows that are harder to characterize the position or trajectory of. It therefore appears that the most practical strategy to tag muons is to *combine* the one-hit-per-row geometric requirement with an area requirement on at least one pulse.

8.8.b Finding Coincident Hits

Now that we have a strategy for tagging muons, we must remember the original goal of this project. We want to see how often a single SPE pulse can be found in each layer of the detector, provided that one layer has been directly hit by a muon in the same event. Figures 8.37, 8.38, and 8.39 are great examples of the type of events we are looking for (although the geometrically-dubious Figure 8.39 gets filtered out when we include an area cut!).

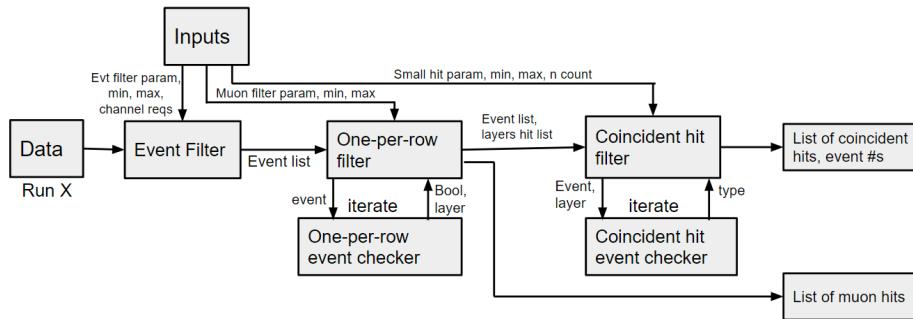


Figure 8.47: Code map of my coincidence-finding program, showing inputs and outputs at each stage.

In these events one layer contains a series of hits that resemble the path of a muon, saturating pulse height in each of those channels. Then in another otherwise empty layer (or more than one, in the case of Figure 8.37) there is a single SPE-level pulse in one of the channels. We want to count how common events like these are, considering every combination of muon-layer and SPE-layer separately to retain as much information as we can about their position relative to one another. Figure 8.47 gives a diagram of the program I wrote to achieve this goal. First data is put through an event filter, which requires that at least one channel in every event has a pulse with area passing a threshold, then a list of suitable events is sent to a second filter. This second filter checks these events to see whether any of the layers have at least one large hit per row (usually at pulse height saturation). This filter outputs a new, refined list of events, as well as a list of corresponding layers that passed the one-hit-per-row test for those events. These go into a third filter, which looks for single SPE pulses in layers aside from the one with the muon hit. As these events are discovered, 28 separate counters keep track of which channel had the muon hit and which combination of other channels had the SPE hit (4 layers can have the muon hit, and for each of those there are: 3 scenarios with an SPE in *one* other layer, another 3 with SPE hits in *two* other layers, and 1 where all three other layers have single SPE hits). The number of muons that

hit the detector are also counted, since not all of them produce events like this, and that makes it possible to model this scenario in simulation.

8.8.c Comparison with Simulation

At the time of this writing neither my analysis code, nor the data taken by the bar detector, is mature enough to produce a compelling simulation validation. One particular issue is that there is an unusually low rate of activity in layer 1 of the detector for recent long runs, which make up the only substantial data with the full detector and all triggers since the detector was commissioned. That said, I have used this data to produce results which may be interesting in light of future studies. This can be seen in Figure 8.48.

		TOTAL EVENTS: 3,235,087		PASSING AREA CUT: 343,302		TOTAL MUON HITS: 7899	
		For layer with 1-per-row hit >1000mV and w/ at least one area>100k hit ,		Small (22-100mV) hits in each layer			
		# events with 0/1 small hit in [layer 0,layer 1, layer 2, layer 4]					
Shower:	[X,1,0,0]	[X,0,1,0]	[X,0,0,1]	[X,1,1,0]	[X,1,0,1]	[X,0,1,1]	[X,1,1,1]
Layer 0	0	6	10	1	0	2	0
Shower:	[1,X,0,0]	[0,X,1,0]	[0,X,0,1]	[1,X,1,0]	[1,X,0,1]	[0,X,1,1]	[1,X,1,1]
Layer 1	0	0	0	0	0	0	1
Shower:	[1,0,X,0]	[0,1,X,0]	[0,0,X,1]	[1,1,X,0]	[1,0,X,1]	[0,1,X,1]	[1,1,X,1]
Layer 2	10	0	17	0	3	1	0
Shower:	[1,0,0,X]	[0,1,0,X]	[0,0,1,X]	[1,1,0,X]	[1,0,1,X]	[0,1,1,X]	[1,1,1,X]
Layer 3	4	0	16	1	3	0	0

Figure 8.48: Results of muon coincidence study on full detector, using a run with 3,235,087 events. Of these events 343,302 survived an area cut requiring that at least one pulse in an event exceeds 100k nVs in pulse area, and of those 7,899 were identified as muons directly passing through the detector. The table represents every geometrical permutation of coincidence between muons and SPEs in the four layers of the detector, with an X indicating a muon, a 1 indicating a single SPE in a layer, and a 0 indicating no activity. These X's, 1's, and 0's are assigned to four values in brackets corresponding to the four layers of the detector, with layer 0 on the left and layer 3 on the right. In other words, [0,X,0,1] represents an event with a muon hit in layer 1 and a single SPE in layer 3. [X,0,1,1] indicates a muon hit in layer 0 with a single SPE hit in both layer 2 and layer 3, etc. The numbers below these spatial labels are counts of how many such events occurred across all of the data. Geometries involving a hit in layer 1 are highlighted red, and appear to be statistically suppressed by the layer's low activity.

It is immediately clear why this study is not prepared for simulation comparison, as the reduced activity in layer 1 has tainted every cell it is involved in and rendered no counts in those scenarios. It is possible that this anomaly is due to an over-active channel in layer 1, which plagued this particular data set. This was partially accounted for by disregarding the problematic channel in the coincident hit filter so that it would not be counted in every single event, but this behavior also biases the trigger and that is not as easy to correct for in analysis. Issues aside, it is promising that so few coincident muon+SPE events were identified in the data even with millions of events taken, as this is, after all, a background study.

Conclusion

The MilliQan bar and slab detectors will collect data over the next few years, and in so doing will push the frontiers of dark matter physics. The MilliQan detectors are simple at their core, taking existing particle detection technology and reshaping it to test new models in ways that have not been done before at the LHC. Even if MilliQan is not successful in detecting mCPs, the beauty of its simple design is that the results it yields are applicable to a large variety of hidden sector and fractional charge theories, and will help guide those searches (among others) in the future. If MilliQan *is* successful, it will serve as a clarion call for high energy physicists worldwide to direct their efforts toward mCPs and their behavior. Indeed, the discovery of an mCP would not necessarily indicate a detection of dark matter—it may be something else entirely! Establishing that connection would need to be the goal of MilliQan’s successor experiments, and many of those there would be. If mCPs are detectable with modern technology, *and* are a feature of a dark sector of dark matter physics, MilliQan has the potential to crack the lid on one of physics’ most enduring and central mysteries.

I am incredibly grateful that I had the opportunity to be involved with this experiment, and that I am able to culminate the work I have done and knowledge I have learned into this thesis. It is rare in high energy physics that student researchers are able to engage with an experiment at every level of its development, and so the experience I have gained working on construction, calibration, and analysis projects will be invaluable for years to come. This is partially a feature of the experiment’s size—far smaller and lower-budget than many high energy collaborations, meaning all-hands-on-deck—but it is also the product of hard work done by MilliQan’s PIs to ensure that graduate and undergraduate students are able to contribute meaningfully on all aspects of the project. I owe so much to my PI, David Stuart, for making sure I got to try everything during my time with the experiment, and for making sure that the other undergraduates and I were treated like equal members of the team. When we were naive or made (costly) mistakes he used it as an educational experience, and he trusted us with big responsibilities that could have easily been given to the more-capable graduate students.

I hope that this thesis may serve as a useful reference on some of the more technical aspects of the detector’s construction and calibration, and that it captures the results of my analysis in ways that can inform future studies on detector data. I also hope that I can continue to be involved with MilliQan in some way as I transition to graduate school, and I eagerly await the publishing of the experiment’s results once there is sufficient data collected from LHC Run 3 to do so.

8.1 Acknowledgements

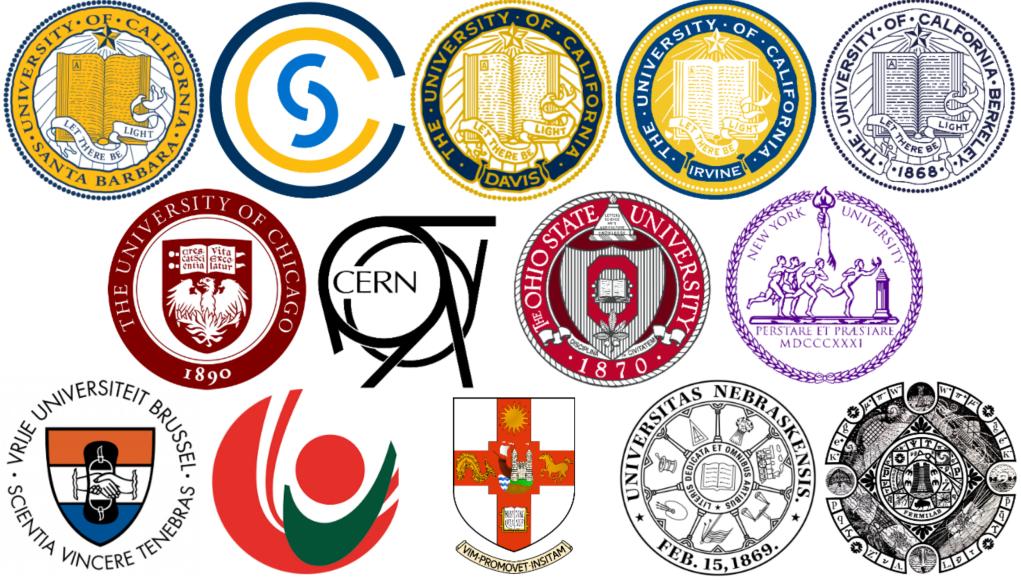


Figure 8.1: Institutional Seals of the MilliQan Collaboration. Top (left-right): UC Santa Barbara, UCSB College of Creative Studies, UC Davis, UC Irvine, UC Berkeley. Middle: University of Chicago, European Organization for Nuclear Research (CERN), Ohio State University, New York University. Bottom: Vrije Universiteit Brussels, Lebanese University, University of Bristol, University of Nebraska Lincoln, Fermilab.

Thank you to David Stuart for recruiting me to the MilliQan Collaboration, for being an amazing advisor, and for proofreading and giving extensive notes on this thesis. Thank you also to Ryan Schmitz for being an excellent Worster mentor, for helping me navigate CERN during our time stationed at CMS, for teaching me cribbage, and for watching bad sci-fi films with me in our vrbo. Thank you to Katia for being an excellent lab partner, constant collaborator, and travel companion. Thank you to Zachary Zerman, who taught me how to solder my PCBs, and thank you to Collin Brahana. Thank you also to Joe Crowley, Keegan Downham, Francesco Setti, Zubair Bhatti, Ian Reed, Mukesh Ghmire, Dariush Imani, Matthew Joyce, and Michael Carrigan for being supportive and brilliant MilliQan PhD students. Thank you to MilliQan postdocs Matthew Citron (now UC Davis faculty) and Neha Santpur who gave invaluable feedback on my writing and help with MilliQan offline all while juggling their many responsibilities with MilliQan, as well as Hualin Mei and David Vannerom who helped me tremendously in installing parts of the detector after the rest of the UCSB contingent departed from CERN. Thank you also to Claudio Campagnari, Christopher Hill, David Miller, Robert Loos, Mouhammad Kamra, Stephen Lowette, and Andy Haas for being insightful and supportive MilliQan leaders and collaborators. Thank you to all of the MilliQan collaboration members that I did not have the pleasure of meeting myself. Thank you as well to my Academic Advisor Prof. Tengiz Bibilashvili, who has been instrumental in my growth as a student and in my navigation of the Physics Research world. Finally, thank you to my friends and to my partner Amalia, who have all supported me through college.

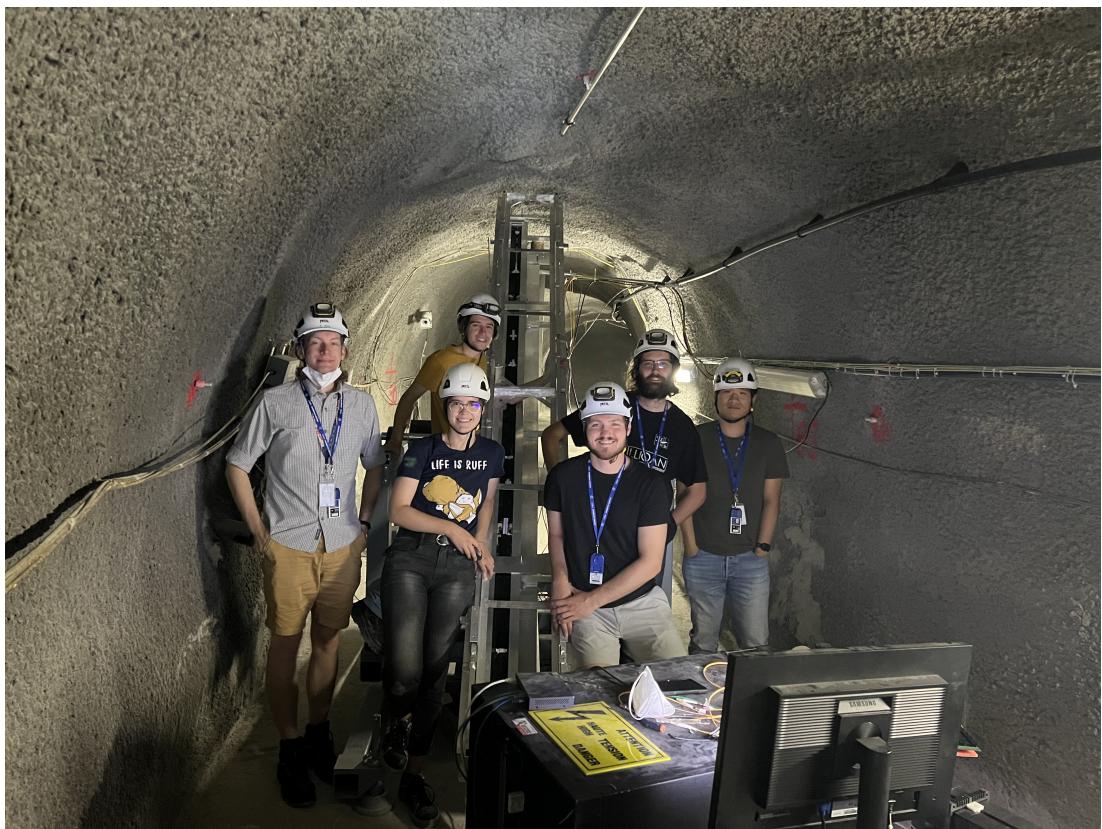


Figure 8.2: The Supermodule 3 Insertion Team. Left to right: Ian Reed, Katherine Larina, Samuel Alcott, Keegan Downham, Ryan Schmitz, Hualin Mei.



Figure 8.3: Attendees of the 5th Annual MilliQan Collaboration Meeting, UC Davis. (Left to right) Jacob Steenis, David Vannerom, Sai Neha Santpur, Stephen Lowette, Matthew Citron, David Stuart, Christopher Hill, Ryan Schmitz, Matthew Joyce, Andrew Haas, Samuel Alcott, Michael Carrigan, Hualin Mei.

Bibliography

- [1] Andrew Haas, Christopher S. Hill, Eder Izaguirre, and Itay Yavin. Looking for milli-charged particles with a new experiment at the LHC. *Physics Letters B*, 746:117–120, jun 2015. URL: <https://doi.org/10.1016/j.physletb.2015.04.062>, doi: 10.1016/j.physletb.2015.04.062.
- [2] Austin Ball, Jim Brooke, Claudio Campagnari, Albert De Roeck, Brian Francis, Martin Gastal, Frank Golf, Joel Goldstein, Andy Haas, Christopher S. Hill, Eder Izaguirre, Benjamin Kaplan, Gabriel Magill, Bennett Marsh, David Miller, Theo Prins, Harry Shakeshaft, David Stuart, Max Swiatlowski, and Itay Yavin. A letter of intent to install a milli-charged particle detector at lhc p5, 2016. [arXiv:1607.04669](https://arxiv.org/abs/1607.04669).
- [3] A. Ball, G. Beauregard, J. Brooke, C. Campagnari, M. Carrigan, M. Citron, J. De La Haye, A. De Roeck, Y. Elskens, R. Escobar Franco, M. Ezeldine, B. Francis, M. Gastal, M. Ghimire, J. Goldstein, F. Golf, J. Guiang, A. Haas, R. Heller, C. S. Hill, L. Lavezzi, R. Loos, S. Lowette, G. Magill, B. Manley, B. Marsh, D. W. Miller, B. Odegard, F. R. Saab, J. Sahili, R. Schmitz, F. Setti, H. Shakeshaft, D. Stuart, M. Swiatlowski, J. Yoo, H. Zaraket, and H. Zheng. Search for millicharged particles in proton-proton collisions at $\sqrt{s} = 13$ tev. *Physical Review D*, 102(3), aug 2020. URL: <https://doi.org/10.1103/physrevd.102.032002>, doi:10.1103/physrevd.102.032002.
- [4] A. Ball, J. Brooke, C. Campagnari, M. Carrigan, M. Citron, A. De Roeck, M. Ezeldine, B. Francis, M. Gastal, M. Ghimire, J. Goldstein, F. Golf, A. Haas, R. Heller, C. S. Hill, L. Lavezzi, R. Loos, S. Lowette, B. Manley, B. Marsh, D. W. Miller, B. Odegard, R. Schmitz, F. Setti, H. Shakeshaft, D. Stuart, M. Swiatlowski, J. Yoo, and H. Zaraket. Sensitivity to millicharged particles in future proton-proton collisions at the LHC with the milliQan detector. *Physical Review D*, 104(3), aug 2021. URL: <https://doi.org/10.1103/physrevd.104.032002>, doi:10.1103/physrevd.104.032002.
- [5] Heinz Andernach and Fritz Zwicky. English and spanish translation of zwicky's (1933) the redshift of extragalactic nebulae, 2017. [arXiv:1711.01693](https://arxiv.org/abs/1711.01693).
- [6] Burkhard Schmidt. The high-luminosity upgrade of the lhc: Physics and technology challenges for the accelerator and the experiments. *Journal of Physics: Conference Series*, 706:022002, 04 2016. doi:10.1088/1742-6596/706/2/022002.
- [7] Gabriel Magill, Ryan Plestid, Maxim Pospelov, and Yu-Dai Tsai. Millicharged particles in neutrino experiments. *Physical Review Letters*, 122(7), feb 2019. URL: <https://doi.org/10.1103/physrevlett.122.071801>, doi:10.1103/physrevlett.122.071801.

[8] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse. Measurement of the electron magnetic moment. *Phys. Rev. Lett.*, 130:071801, Feb 2023. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.071801>, doi:10.1103/PhysRevLett.130.071801.

[9] Sacha Davidson, Steen Hannestad, and Georg Raffelt. Updated bounds on millicharged particles. *Journal of High Energy Physics*, 2000(05):003–003, may 2000. URL: <https://doi.org/10.1088/1126-6708/2000/05/003>, doi:10.1088/1126-6708/2000/05/003.

[10] R. Essig, J. A. Jaros, W. Wester, P. Hansson Adrian, S. Andreas, T. Averett, O. Baker, B. Batell, M. Battaglieri, J. Beacham, T. Beranek, J. D. Bjorken, F. Bossi, J. R. Boyce, G. D. Cates, A. Celentano, A. S. Chou, R. Cowan, F. Curciarello, H. Davoudi-asl, P. deNiverville, R. De Vita, A. Denig, R. Dharmapalan, B. Dongwi, B. Döbrich, B. Echenard, D. Espriu, S. Fegan, P. Fisher, G. B. Franklin, A. Gasparian, Y. Gershtein, M. Graham, P. W. Graham, A. Haas, A. Hatzikoutelis, M. Holtrop, I. Irastorza, E. Izaguirre, J. Jaeckel, Y. Kahn, N. Kalantarians, M. Kohl, G. Krnjaic, V. Kubarovský, H-S. Lee, A. Lindner, A. Lobanov, W. J. Marciano, D. J. E. Marsh, T. Maruyama, D. McKeen, H. Merkel, K. Moffeit, P. Monaghan, G. Mueller, T. K. Nelson, G. R. Neil, M. Oriunno, Z. Pavlovic, S. K. Phillips, M. J. Pivovaroff, R. Poltis, M. Pospelov, S. Rajendran, J. Redondo, A. Ringwald, A. Ritz, J. Ruz, K. Saenboonruang, P. Schuster, M. Shinn, T. R. Slatyer, J. H. Steffen, S. Stepanyan, D. B. Tanner, J. Thaler, M. E. Tobar, N. Toro, A. Upadye, R. Van de Water, B. Vlahovic, J. K. Vogel, D. Walker, A. Weltman, B. Wojtsekhowski, S. Zhang, and K. Zioutas. Dark sectors and new, light, weakly-coupled particles, 2013. [arXiv:1311.0029](https://arxiv.org/abs/1311.0029).

[11] A. A. Prinz, R. Baggs, J. Ballam, S. Ecklund, C. Fertig, J. A. Jaros, K. Kase, A. Kulikov, W. G. J. Langeveld, R. Leonard, T. Marvin, T. Nakashima, W. R. Nelson, A. Odian, M. Pertsova, G. Putallaz, and A. Weinstein. Search for millicharged particles at SLAC. *Physical Review Letters*, 81(6):1175–1178, aug 1998. URL: <https://doi.org/10.1103/PhysRevLett.81.1175>, doi:10.1103/physrevlett.81.1175.

[12] Bob Holdom. Two $u(1)$ ’s and charge shifts. *Physics Letters B*, 166(2):196–198, 1986. URL: <https://www.sciencedirect.com/science/article/pii/0370269386913778>, doi:[https://doi.org/10.1016/0370-2693\(86\)91377-8](https://doi.org/10.1016/0370-2693(86)91377-8).

[13] Julia Woithe, Gerfried J Wiener, and Frederik F Van der Veken. Let’s have a coffee with the standard model of particle physics! *Physics Education*, 52(3):034001, mar 2017. URL: <https://dx.doi.org/10.1088/1361-6552/aa5b25>, doi:10.1088/1361-6552/aa5b25.

[14] URL: <https://home.cern/science/physics/cosmic-rays-particles-outer-space>.

[15] Information@eso.org. The coma galaxy cluster as seen by hubble. URL: <https://esahubble.org/images/heic0813a/>.

[16] URL: <https://home.cern/science/accelerators/large-hadron-collider>.

[17] URL: <https://home.cern/science/engineering/restarting-lhc-why-13-tev>.

[18] URL: <https://timeline.web.cern.ch/timeline-header/89>.

- [19] URL: <https://home.cern/science/experiments/cms>.
- [20] URL: <https://home.cern/science/experiments/atlas>.
- [21] Mortimer Abramowitz and Michael W Davidson. Concepts in digital imaging technology: Photomultiplier tubes. URL: <https://hamamatsu.magnet.fsu.edu/articles/photomultipliers.html>.
- [22] Jun 2004. URL: <https://cds.cern.ch/record/745737?ln=en>.
- [23] Anne Helmenstine. What is a micron? definition and examples, May 2021. URL: <https://sciencenotes.org/what-is-a-micron-definition-and-examples/>.
- [24] David Stuart. Description of milliquan calibration plan, Jul 2022. URL: <http://dstuart.physics.ucsb.edu/Lgbk/pub/E41105.dir/E41105.html>.
- [25] David Stuart. Design changes between v4 and v5 of the milliquan base amp boards, Jul 2022. URL: <http://dstuart.physics.ucsb.edu/Lgbk/pub/E41116.dir/E41116.html>.
- [26] David Stuart. R878 base amp board, Aug 2021. URL: <http://dstuart.physics.ucsb.edu/Lgbk/pub/E41004.dir/E41004.html>.
- [27] Hamamatsu Photonics K.K. Editorial Committee. *Photomultiplier Tubes: Basics and Applications*. Hamamatsu Photonics K.K.
- [28] Photomultiplie. URL: https://www.hamamatsu.com/us/en/product/optical-sensors/pmt/pmt_tube-alone/head-on-type/R878.html.
- [29] Tom Hall. Purpose of optical coupling gel, Aug 2021. URL: <https://www.gammaspectacular.com/phpBB3/viewtopic.php?t=948>.
- [30] R. A. Millikan. On the elementary electrical charge and the avogadro constant. *Phys. Rev.*, 2:109–143, Aug 1913. URL: <https://link.aps.org/doi/10.1103/PhysRev.2.109>, doi:10.1103/PhysRev.2.109.
- [31] Claus Grupen. Physics of particle detection. *AIP Conf. Proc.*, 536(1):3–34, 2000. arXiv:physics/9906063, doi:10.1063/1.1361756.
- [32] Samantha Katherine Dooling. *Differential Cross Section Measurement of Drell-Yan Production and associated Jets with the CMS Experiment at the LHC*. PhD thesis, U. Hamburg, Dept. Phys., 2014. doi:10.3204/DESY-THESIS-2015-002.
- [33] Neha Santpur. Milliquan: First results and plans for run 3, Jun 2023. 6th Forward Physics Facility Meeting.
- [34] Prashant Shukla and Sundaresh Sankrith. Energy and angular distributions of atmospheric muons at the earth, 2018. arXiv:1606.06907.
- [35] Ryan Schmitz. Simulating the milliquan detector via geant4, Oct 2019. MilliQan Technical Information. URL: https://github.com/rtschmitz/milliQanSim/blob/main/Simulating_the_MilliQan_Detector_via_Geant4.pdf.