Design, Construction, and First Tests of a Demonstration Spark Chamber

Lisa Lin,* Robin Peter,† Sophia Vlahakis,‡ and Tara Vogel§

The University of Chicago, Illinois, 60637

(Dated: September 21, 2018)

Once used extensively in particle physics research, the cosmic ray spark chamber is a tracking particle detector now favored for education and demonstration. We have constructed and are nearing installation of a chamber for educational and demonstration purposes, selecting a shelf-type design to prioritize aesthetic appeal and minimal maintenance. An incoming particle triggers a fast high-voltage pulse to metal plates within the chamber that allows sparking along the particle's ionized path. The shelf-style design and careful scintillator placement limits spurious edge sparking. The coincidence circuit comprised of scintillation counters and NIM units successfully delivers a logic pulse within 60 ns, and its 1.3-V pulse is amplified ten-fold in a total delay time of 500 ns from detection to spark gap pulse, triggering an 8-kV pulse in the chamber. The prototype chamber is tested independently with helium and argon gases. A plan is presented to predict lifetime based on the muon efficiency of the prototype chamber. Initial stages of an interactive soft-shut off are presented for safety and consideration of the chamber demonstration's surrounding academic environment.

I. INTRODUCTION

Spark chambers have been used in the past as tracking particle detectors, and are most widely used today for educational purposes [1]. Our group has built a spark chamber for demonstrative purposes outside of the Maria Goeppert-Mayer lecture hall at the University of Chicago (KPTC 106), which was newly dedicated in 2017.

A spark chamber consists of uniformly separated metal plates held between a pair of muon-detecting scintillators that triggers a high-voltage pulse circuit. The chamber is filled with a noble gas (or mix of noble gases), to allow for easy ionization of gas atoms by any through-going particle with enough energy—namely, cosmic rays and minimum ionizing particles (MIPS). As a charged particle passes through, scintillators located directly above and below the chamber emit photons that propagate to attached photomultiplier tubes (PMTs), where they are converted to electrical signals. These signals are then evaluated for coincidence. Coincident signals generate a logic pulse, which is triggered and subsequently amplified to roughly 3 kV. This amplified pulse triggers arcing across the spark gap, effectively closing a switch in the circuit and prompting a rapid discharge from capacitors onto selected chamber plates. By applying high voltage to the plates before the ionized electrons have a chance to recombine, the spark chamber prompts electrical breakdowns along the path of the original high-energy particle.

Development of the demonstration chamber can be organized into four overarching components: the chamber design, the muon trigger unit, the high voltage circuit, and gas flow through the chamber. The Cambridge spark chamber group has produced several prototype components for a similar purpose; we have used their procedure as an initial guide [1]. This document tracks our major considerations, revisions, and additions to their work, and presents that information for the perusal and enjoyment of the general public.

II. OPERATING MECHANISMS

A charged particle of sufficient energy will ionize the gas atoms in its path through the chamber. The electrons produced via this ionization then collide with both with each other and other gas atoms. Each of these collisions liberates more electrons, which in turn produce further collisions in a cascading effect of ionizing electrons called a Townsend avalanche.

Plates in the chamber are held alternatingly at ground and a high positive voltage. Electrons drift from the ground plate to the high voltage plate, while the ionized gas atoms flow in the opposite direction, from the anode (HV) to the cathode (ground). This creates a dipole within the generated electron avalanche (Fig. 1).
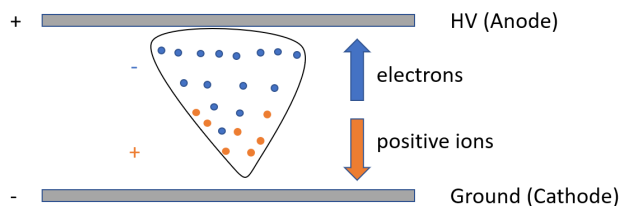

FIG. 1: A dipole avalanche.

Since recombination of an electron with a gaseous ion supposes an increase in electric potential inversely proportional to the radial distance of the electron from the atomic nucleus, the system undergoes a liberation of energy. Recombination therefore results in photon emission

---

* lllin@uchicago.edu
† rpeter@uchicago.edu
‡ sophiavlahakis@uchicago.edu
§ vogelt@uchicago.edu

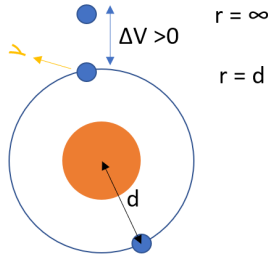that generates yet more ion pairs beyond the avalanche through the photoelectric effect [1].

FIG. 2: Liberation of potential energy when an electron is reabsorbed by a gas atom

The dipole of the initial avalanche encourages the formation of a series of avalanches at the head and tail of the primary, which join together to form a streamer (Fig. 3). Upon reaching the anode and cathode, this streamer forms a highly conductive plasma channel that sparks once the electrical breakdown voltage of the gas in the chamber is attained.
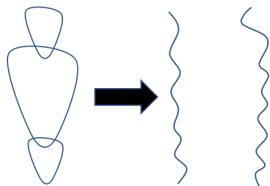

FIG. 3: A streamer, formed by a series of merging avalanches, becomes a conducting plasma channel upon reaching the cathode and anode plates.

By delivering a high-voltage pulse to the chamber upon detection of a muon, the spark chamber creates arcing along the path ionized by the particle (Fig. 4).

III. BOX DESIGN

Design of a functional spark chamber first and foremost must ensure that the chamber (at 100% efficiency) sparks on muons, and only muons.

A. Muons

Muons are chosen for detection because they are about 200 times more massive than their electron cousins. They therefore experience less deceleration when traveling through a medium, so a muon will produce less bremsstrahlung (deceleration radiation at far-field) than other cosmic rays at the same energy. As such, they
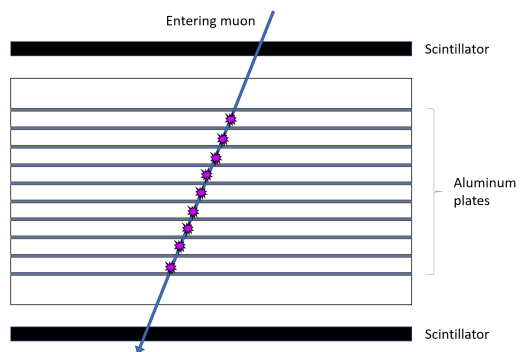

FIG. 4: Basic schematic of a muon traversing a spark chamber. When a particle passes through both scintillators, a PMT coincidence is registered, triggering the high voltage circuit and dumping high voltage across each pair of ground-HV plates. This causes arcing between the two; in order for sparking to occur where the muon has passed through, the total circuitry delay must be shorter than electron recombination time.

are more likely to pass through both scintillators to produce a track through the chamber. For this same reason, other minimum ionizing particles (MIPS) such as protons and pions will also be detected. Nonetheless, muons are especially abundant, with a flux rate of about 1 muon/cm^2/min at sea level. In addition, they have longer lifetimes and deeper penetration into matter than other cosmic particles, making muon tracks particularly clean and straight. Commonplace and easy to identify, muons are the ideal detection candidates for the spark chamber.

To achieve maximum efficiency, our spark chamber takes into account gas flow and purity, edge sparking, and connections to the high voltage circuitry described in Section V. We first experimented with each of these design elements for our spark chamber using a test chamber we built, measuring 6"× 8"×7".

B. Design parameters

1. High voltage connections

Reference [1] outlines two different types of spark chamber designs: a "box" type with plates separated by plastic spacers, surrounded by a larger container (Fig. 6a), and a "sandwich" type, a stack of alternating polished Perspex frames and aluminum plates (Fig. 6b). We opted to combine the best of both designs in a shelf-type design, using an encapsulating container to keep the chamber gas-tight but opting for acrylic instead of metal (as in [1]). This choice optimizes the visibility of the chamber by making all four lateral sides potential viewing areas for any generated sparks.

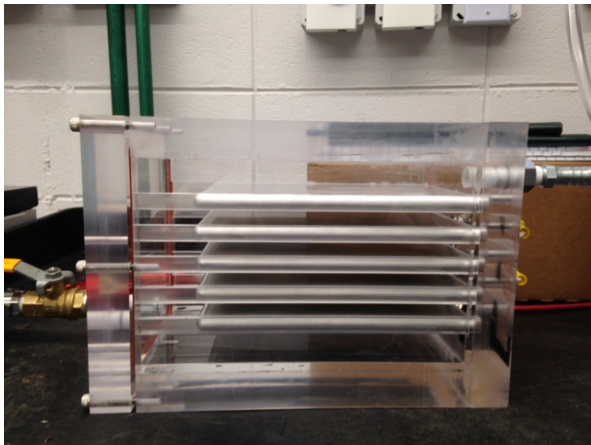To minimize undesired sparking and ionization of the

FIG. 5: Our shelf-style test box in early stages. Aluminum plates slide into grooves in the sides of the insulating acrylic enclosure, allowing for easy insertion and removal. Gas enters the chamber through the inlet in the lower-left of the chamber, and exits through the upper-right inlet.

(a) A box-style spark chamber, with aluminum high-voltage plates held apart by plastic spacers. The plate assembly is enclosed within a relatively gas-tight container, which helps minimize gas leakage. However, this container can also restrict visibility if opaque materials, e.g. metal, are used.

(b) A sandwich-style spark chamber, with aluminum plates wedged between layers of acrylic frame; layers are shown separated. Benefits of this design include little to no edge sparking, since the edges of the metal plates are either in contact with the air or buried between the Perspex. However, the numerous seal locations can lead to degradation of gas purity.
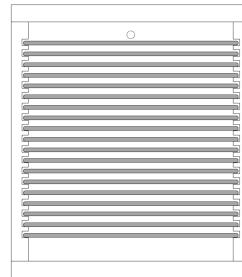
(c) The back view of a shelf-style spark chamber, with aluminum plates inserted into a grooved acrylic container. Note the gaps on all sides of the plates; this clearance allows for the plates to be smoothly slid in and out of the chamber as needed, and their fully rounded edges, which may help to reduce edge sparking. The back face, containing the gas manifold inlet (see Section III C), is not shown for clarity. The hole in the top center of the front face is the gas outlet hole.

FIG. 6: Schematics of three spark chamber designs; we ended up combining designs (a) and (b) in our final design, (c). See Appendix A for full drawings of our final chamber design.
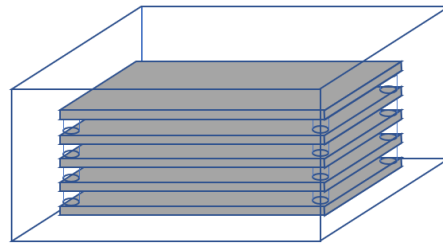
gas as seen in several designs presented by [1], and in particular to prevent corona discharge in the electronics, we decided to keep all high voltage circuitry and connections outside of the acrylic chamber. O-ring screws, fitted into drilled and tapped holes in the plates, provided the electrical connections between the chamber plates and external high voltage circuitry. Attachment of large, fast discharge capacitors to the backs of these screws not only prevents wiring from entering the chamber, but also minimizes inductance of the circuitry, and so speeds up the rise pulse onto the high voltage plates (see Section V C for more details).
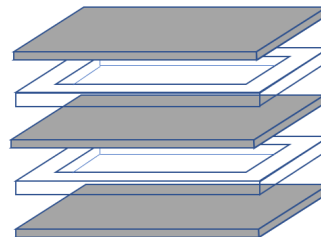
2. Edge sparking

Taking inspiration from the sandwich-style design, we machined lateral grooves in all sides of the acrylic container save the front face to isolate the aluminum chamber plates. Burying the plates in this manner inhibits edge sparking by isolating any edges conducive to divergent electric fields [1][2]. As an additional measure, we fully rounded the edges of the plates with a radius of 0.125" to prevent sparking across sharp edges.

Another step we took was to resize our scintillators to create a fiducial volume (Fig. 7). Reducing the scintillator dimensions to be smaller than the chamber plates filters out any muons that might have passed through the chamber right at the fringes of the plates. Such muons could leave carbon traces on the acrylic walls, which serve as low-resistance conductive channels for sparks. A chamber with carbon traces will likely spark on the traces instead of on actual muons.

Also notice that although the rightmost end of the aluminum plates slots neatly into grooves in the back face of the chamber, a gap remains between the leftmost edge of the plate and the front face. We allowed the front

face of the chamber to be removable, to allow for easy insertion and removal of the plates into the grooves for cleaning purposes, etc. The orange O-ring seal between the front face and the rest of the box creates an air-tight fit upon tightening socket screws drilled into the front face. Rather than trying to guess how much the O-ring would compress, we provided a small gap between the front face and the front of the aluminum plates. This avoids the creation of more points of contact with the conductive plates, which could cause spurious sparking. Furthermore, even if the chamber plates did extend into grooves in the front face, there would be insufficient space for an O-ring seal.
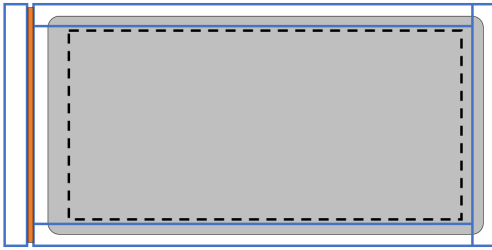

FIG. 7: A bird's eye view of the spark chamber - note the inset of the aluminum plates into three of the four acrylic (blue) sides, and the gap left between the front face and the plates to account for O-ring seal (orange) compression. The black dashed line represents the location of the scintillation counters, creating a fiducial volume for the spark chamber. This further reduces the amount of edge sparking the chamber generates. Screws and gas valves not shown; figure not drawn to scale.

C. Gas-related design choices

We discovered a number of flaws with our original test chamber design with regards to gas flow-through. Gas purity of a noble gas in the chamber is critical as it lowers the breakdown voltage at which a spark will be created between each pair of high voltage-ground plates. As seen in Fig. 8, the gas input valve of the test chamber is located at the bottom of the chamber because the two tested gases, He and Ar, rise. Following the path of least resistance, the gas being flushed through the chamber tends to rush through the large BSPP inlet pipe to the large bottom gap closest to it. Since the air already in that space has no outlet, gas cannot enter and air cannot leave efficiently. The same goes for the smaller gaps between the plates, so that only the top gap has effective flow-through.

We instead sought to equip the final chamber design with a vertical gas inlet hole through the top of the back face, and small gas flow-through holes above and below each plate. Though a single gas inlet hole would have been the most elegant solution, machining of a 15"-deep
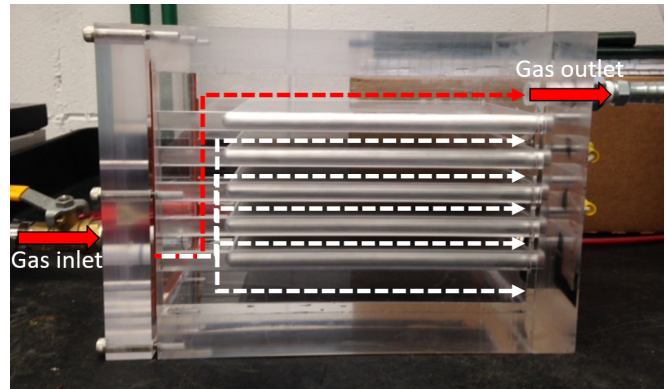

FIG. 8: Paths of gas flow through the test chamber. White denotes "dead ends -" though gas may flow into these gaps, air has no way of exiting the chamber, as the O-ring screws provide a nearly perfect seal.

hole was unreasonably expensive. The modified design consisted of a milled groove in the back face, with a chamfered manifold bar that nestled into a shallower groove on the back face. This arrangement leaves a large gap allowing for easy flow-through to the gaps between the chamber plates (Fig. 9). The gas input location remains near the bottom of the chamber, because although a gas inlet hole in the center of the manifold can still force gas through the smaller inlets at high flow-through pressures, the spark chamber itself will not be running at such high pressures (see Section VI C for more details on flushing the chamber).
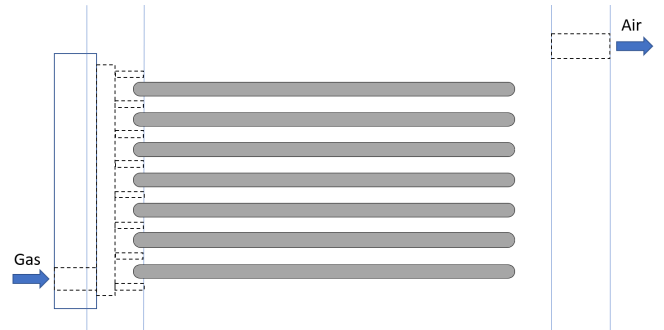

FIG. 9: The gas flow schematic of the full-sized spark chamber. Dashed lines represent drilled through-holes in the chamber. In this improved gas flow design, gas enters through a hole in the gas manifold; a deeper groove in the back face creates a channel connected to smaller inlet holes, enabling gas to flow through all routes in between the plates. Air can then be pushed out through the gas outlet centered near the top of the front face.

D. Construction of the Chamber

The acrylic faces of the chamber and aluminum plates were custom machined by an external shop. We then bonded the faces together with epoxy.

During construction of the test box, epoxy was applied on one acrylic face, then quickly pressed against a neighboring piece within the allowed 90-second drying period of the adhesive. Though we initially used a paintbrush to apply the epoxy, it proved incapable of brushing on the epoxy in large enough quantities, and we observed the formation of many air bubbles upon joining two walls. After noticing this, we switched to an eyedropper to apply a thicker layer of epoxy. This method covered a larger surface area at once than the paintbrush and gave us greater room for error in regards to the time constraint from epoxy evaporation and specific placement of each acrylic piece.

The optimal assembly procedure was to first attach the sides to the bottom, then the top on the sides, then the back face. After washing both the test chamber and the aluminum plates with detergent and water, the plates were slid into the grooves in the sides. The large O-ring seal was placed carefully on the front face, making sure that it did not interfere with any of the holes drilled in the front, and the front face screwed in with socket screws.

When tightening any threaded object - whether that be a BSPP pipe connector or the socket screws in the front face (See Appendix A for detailed drawings), special care must be taken to apply an equal torque to each screw; this evenly pressurizes the seal, minimizing trapped air bubbles. On the other hand, to prevent the acrylic from cracking, the screws should be tightened only to the extent that they can no longer be loosened by hand.

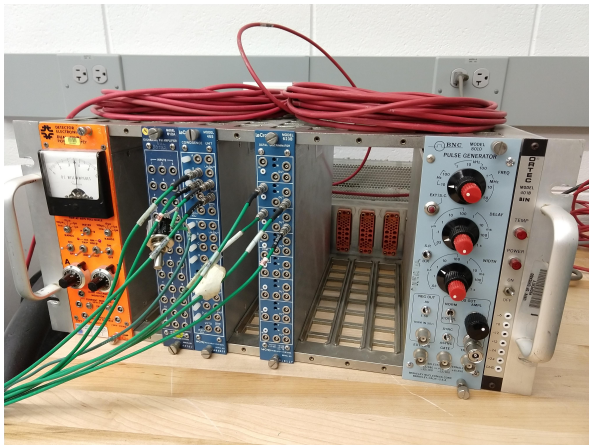## IV. MUON-TRIGGER SYSTEM

FIG. 10: The NIM crate with (from left to right) the PMT high voltage unit, the logic unit, the discriminator, and a pulse generator.

The muon trigger system refers to the scintillator, PMT, NIM crate, NIM modules, and coincidence amplifier. These elements collectively detect incidence of a muon and initiate the high-voltage pulse circuit with minimal delay.

A. Scintillation Counters

FIG. 11: Our final wrapped scintillation counters.

A scintillator is a plastic material that fluoresces under the influence of ionizing radiation; coupled to a photomultiplier tube (PMT), the assembly is called a *scintillation counter*. A charged muon that passes through a scintillation counter temporarily excites the electrons of the scintillation material, causing photon emission. This light is detected by the PMT and converted into electrical signals that in turn initiate the rest of the coincidence circuit. The scintillator is wrapped tightly in light-tight black tape to reduce superfluous noise from photons from ambient lighting that would otherwise leak into the plastic.

We received several pre-assembled scintillation counters for the project, passed down from an older detector built by Jim Cronin's group. We adapted these counters for use in our spark chamber.

Unlike the standard paddle-style scintillation counter, these historic counters featured enormous scintillators coupled with relatively small PMTs. This discrepancy contributed to some inefficiency in signal production per actual muon, since photons reflecting within a scintillator would not necessarily traverse to the small collection window at the PMT. The scintillators were cut to 24"×12" to reduce the disproportion and to better fit the dimensions of the new chamber.
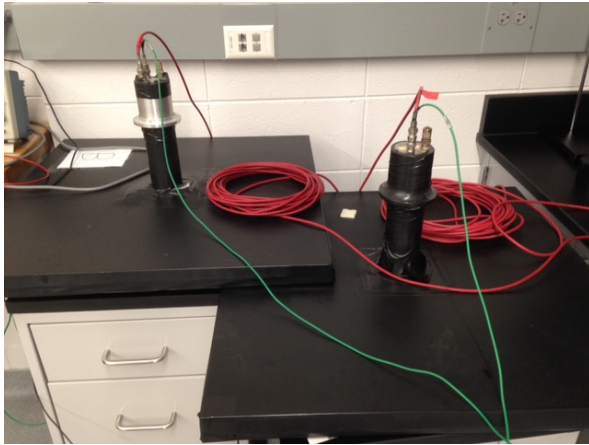
FIG. 12: The original twentieth-century scintillation counters.

B. Scintillation Counter Construction

1. Cutting and polishing the scintillators

Our preparation and assembly procedures were based on the work of the Berkeley Lab group [3]. We lacked the recommended rotating polishing pad and powders, and substituted clog- and water-resistant sandpaper sheets instead. The resultant scintillators are functional, if less visually pristine than those polished by the Berkeley Lab (ref. Section IV C). Refer to Appendix B for a more detailed procedure.

We cut the Cronin scintillators in half with a band-saw to obtain two pieces of scintillation material approximately 24" L × 12" W. To determine the optimal procedure for polishing the scintillators, we milled the cut edge of one of the pieces before polishing, and polished the other directly after cutting.

600-, 800-, and 1200-grit sandpaper was used in three stages to polish the scintillators, alternating the direction of sanding each stage to be orthogonal to the last, removing the striations left by the previous grit. The final product is polished such that tilting the scintillator allows visibility of the other sides.

After polishing our first scintillator, we noticed striations left on the unpolished sides of the scintillator by the sandpaper (Fig. 14). To minimize scratches, we would recommend strips of painter's tape on the sides, a pencil line's width below the edges of the side being polished. Regular scotch tape will work in a pinch, but be aware that it will leave sticky marks that will need to be removed with isopropanol (Fig. 15).

Milling the cut edge of the scintillator gave a flatter, smoother starting surface and dramatically reduced the time required to polish. There was no visible difference in the final state between the milled and unmilled scintillators, however. Since the milling step itself is slow and tedious, future scintillators should be polished without
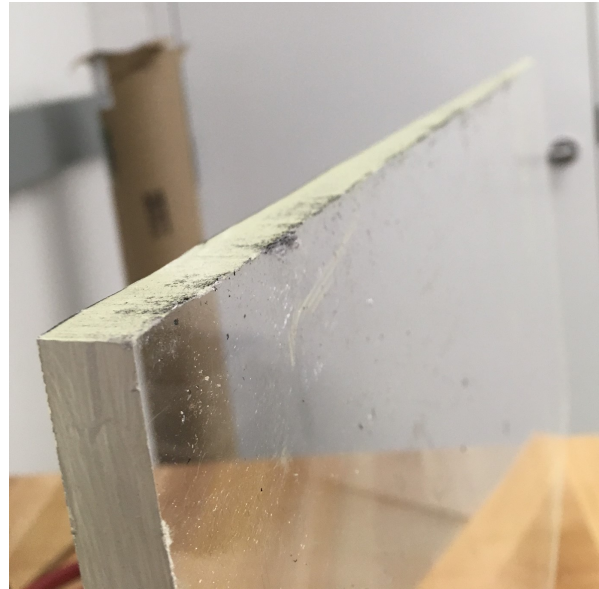

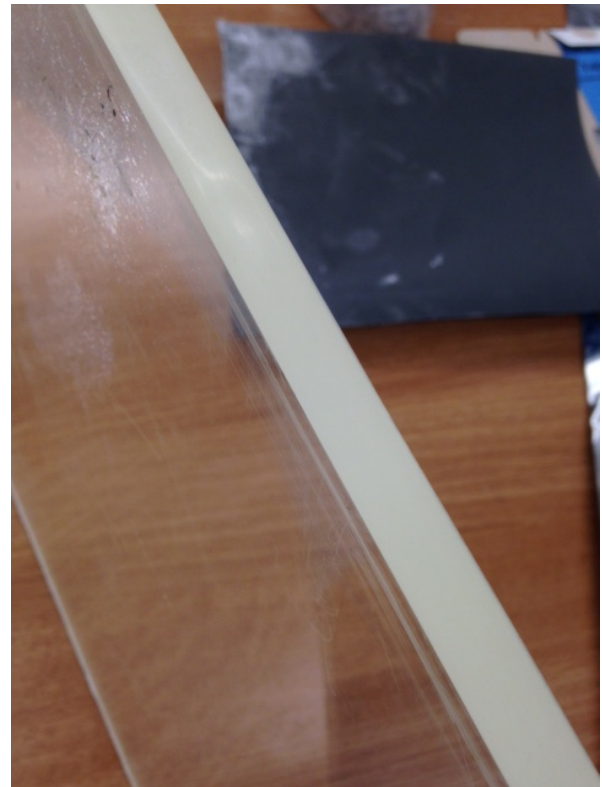FIG. 13: Scintillator just after being cut with a band saw.

FIG. 14: Scintillator after being polished with 600, 800, and 1200 grit sandpaper. Careless sanding left scratches on the sides.

milling but with an additional coarse sandpaper stage (around 400 grit) before beginning the 600-grit step.

The scintillators were tightly wrapped in 2-mil (50 μm)

FIG. 15: Adhesive left after taping the sides of the scintillator to prevent scratching by the sandpaper used to polish them.

(a) A PMT decoupled from its scintillator, with a visible layer of rubbery adhesive.

(b) A fully cleaned PMT, after some intensive scraping with a razor blade.

FIG. 16: Process of cleaning off the Cronin PMTs and preparing them for use in our own scintillation counters.

aluminum foil with a hole to give ample clearance for the end of the PMT to be attached to the scintillator, which operates at high voltage. Overlapping edges of the foil were sealed with scotch tape.
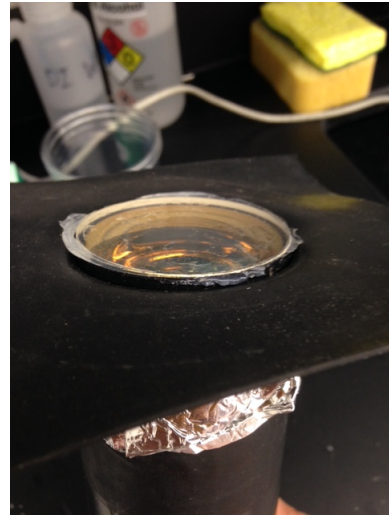
2. Wrapping and cleaning the PMTs

We decoupled the Cronin scintillation counters and carefully used a razor blade to peel away the layer of optical cement still attached to the PMTs. Special care was taken to not scratch the surface (Fig. 16b).

3. Coupling the PMT and scintillator

The scintillation counters are comprised of layers of (from bottom to top): cardboard, black paper, black plastic cover, scintillator, and black plastic cover, with the sides wrapped using thick black electrical tape. The PMTs themselves have a ring of black rubbery plastic around the lenses (Fig. 16a).

As a substitute for the black cardstock used by the Berkeley group, we cut and taped up cardboard boxes to create two rectangular panels slightly smaller than the scintillator, around 11 7/8" × 23 7/8". We taped a piece of black plastic to one, and glued a piece of black paper on either side of the other. A layer of black electrical tape around the edges ensured a strong connection between the plastic (or paper) and the cardboard. The PMT and its plastic shell were fit through a 2.5" diameter hole through the center of the paper-covered scintillator (Fig. 17). The extra support of the cardboard pressing down on the PMT's plastic insulation secures the lower PMT, which will be upside-down when set up.

Lightly sanding and cleaning the exposed face of the scintillator with isopropanol prepared it for coupling to the PMT with a layer of optical cement. The surrounding area was secured with electrical tape to allow for a strong seal during the drying period.

Once the PMTs were cured and safe to handle, the plastic-covered cardboard was attached to the foil-wrapped scintillator with scotch tape. The entire stack was sealed with black electrical tape around the sides (Fig. 18).
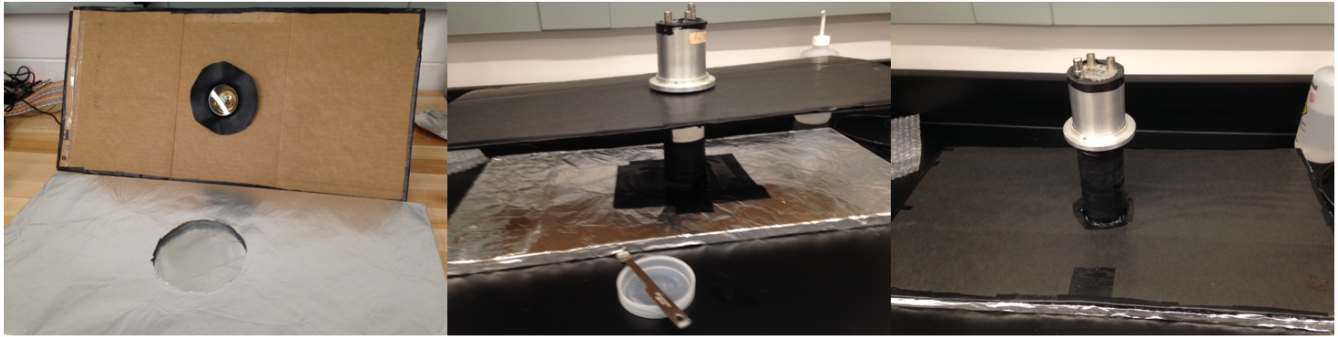
FIG. 17: Process of coupling the PMT with scintillator. The rubber ring around the PMT lens insulates it from contact with the foil, while cardboard provides extra support for keeping the PMT in place (left). After applying the mixture of optical resin and hardener, the PMT is taped down, and any visible spots of scintillator covered (middle). Keeping the edges of the black cardboard just shy of the scintillator edges ensures that the scintillator makes good contact with the black tape that will secure the entire assembly (right).
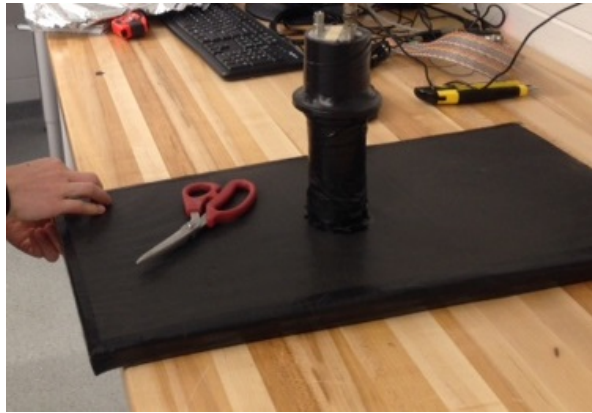

FIG. 18: Taping around the edges of the scintillation counter assembly.

C. Pulse Characterization

1. Rates and Light-Leak Testing

Each scintillation counter was tested for light-leaks through analysis of the trigger frequency. The rate of triggering is highly dependent on the PMT voltage. Oscilloscope rate measurements were observed at 1.9 kV so that the frequency tended to be above the scope's 10-Hz distinguishing threshold, while triggers were manually counted at 1.7 kV to provide another confirmation. We turned overhead lights on and off and directed a flashlight towards the scintillation counters; no significant difference was observed. The scintillation counters are suitably light-leak-tight.

TABLE I: Brief light-leak check of scintillation counters.

| Voltage (kV)	Lights On	Lights Off
1.7	44 triggers/min	50 triggers/min
1.9	20-110 Hz	20-80 Hz

2. PMTs

The new scintillation counters were tested in a voltage range of 0-2 kV. Both individual counters began reliably triggering at around 1.7 kV. Each PMT provided a pulse of magnitude ranging from 0.5-4 V. A 50-ohm BNC terminator attached to the PMT anode reduced the pulse width from a time on the order of microseconds to about 100 ns.

One of the PMTs was roughly handled before this project received it and displayed a significant amount of noise (Fig. 19b). This oscillation was consistent amongst fluctuating pulses, and the magnitude of this PMT pulse showed the same fluctuations from 0.5-4 V. Both pulses were fed into a discriminator module to determine whether the damaged PMT would affect the coincidence circuit.

3. Discriminator

The discriminator module reduces and separates noise from muon signals by only accepting inputs above a certain threshold. It outputs a pulse when the input crosses the voltage threshold, ending the pulse when the input falls back below. The discriminator outputs corresponding to both scintillation counters were identical, despite the noise of the damaged PMT; we continue to use the PMT for this project. The oscilloscope has an impedance of 1 MOhm. A 50-Ohm terminator acts as a temporary
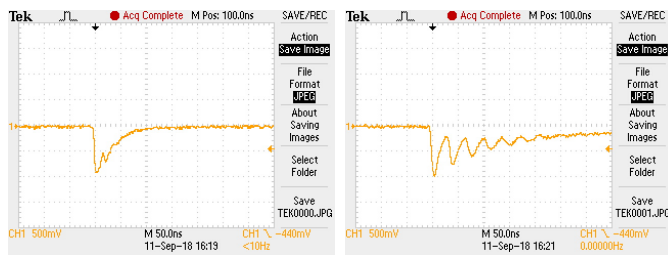
(a) A PMT pulse of roughly 0.5 V and 100 ns.

(b) The noisy pulse from the damaged scintillation counter.

FIG. 19: Two pulses taken from the two scintillation counters at 1.8 kV.

substitute for the load of the amplifier circuit and allows for better characterization of the PMT pulses.

4. Coincidence

The two discriminator outputs are read by a logic unit, which searches for a two-fold coincidence between the two PMTs. This coincidence ensures that the pulse circuit only triggers on muons that leave an ionization track through the entire chamber for a visually complete spark path, as opposed to a muon that passes only through one scintillator and leaves a track through only part of the chamber.
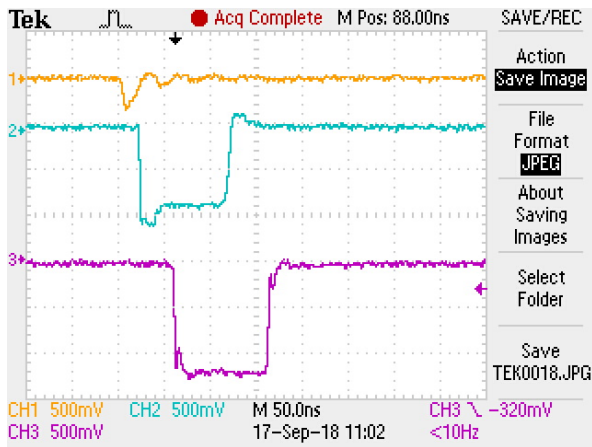

FIG. 20: From top to bottom: PMT pulse, discriminator output, and coincidence pulse, terminated with 50Ω. 60 ns are elapsed between muon incidence and the leading edge of a coincidence pulse, possibly attributable to 15 ft (cumulative) of LEMO cables from a PMT to a coincidence pulse input on the oscilloscope.

5. Coincidence amplifier

The high-voltage pulse circuit requires a threshold input pulse in order to trigger the BJT (see Section V A.

We probed the amplifier circuit's behavior with an input pulse from a NIM pulse generator on the positive complement setting to determine this threshold. Fig. 21a superimposes several oscilloscope pictures to demonstrate the effect of input pulse width on BJT saturation. There is a threshold around 70 ns, when the BJT collector sustains its highest peak voltage (11V) before falling again. Note that this effect could also be attributed to the larger negative peak of the coincidence pulse that occurred when the width was adjusted, as we were unable to fully isolate the two parameters. The output of the logic unit is a 1.3V pulse. The amplifier circuit we built requires at least a 1.7V pulse, though we were unable to further increase the pulse magnitude with the existing mechanisms (Fig. 21b); more in Section V.

V. HIGH-VOLTAGE PULSE CIRCUIT

The high-voltage circuit delivers a high-voltage (6-10kV) pulse to the chamber before gas recombination occurs–within 500 ns for 100% sparking efficiency or within 600 ns for 85% efficiency in helium gas [1]. Our circuit, based on a design first used in a spark chamber built at the Dutch Institute for High Energy Physics and then in a chamber built by University of Cambridge High Energy Physics, is composed of three major components: an amplifier circuit, a transformer, and a spark gap.

A. Amplifier circuit

The amplifier circuit uses two fast-switching transistors–a Bipolar Gate Transistor (BJT) and an Insulated Gate Bipolar Transistor (IGBT)–to amplify the -1.5V pulse output from the logic unit to a -100V pulse within 100 ns. This amplified pulse serves as the input to the transformer.

During our construction of this circuit, we utilized LTspice circuit simulation software [4] to serve as a reference for understanding how the operation of our circuit compared with the ideal case.

1. Transistors

To gain even a cursory understanding of how the amplifier circuit works, one must understand how each of the transistors we used functions.

The BJT (S1)

BJTs come in two flavors: PNP and NPN; however, the amplifier circuit used a PNP-variety BJT, so only that type is relevant to the discussion here. A BJT acts as a current amplifier; one uses a small current at its base to control a larger current from emitter to collector.
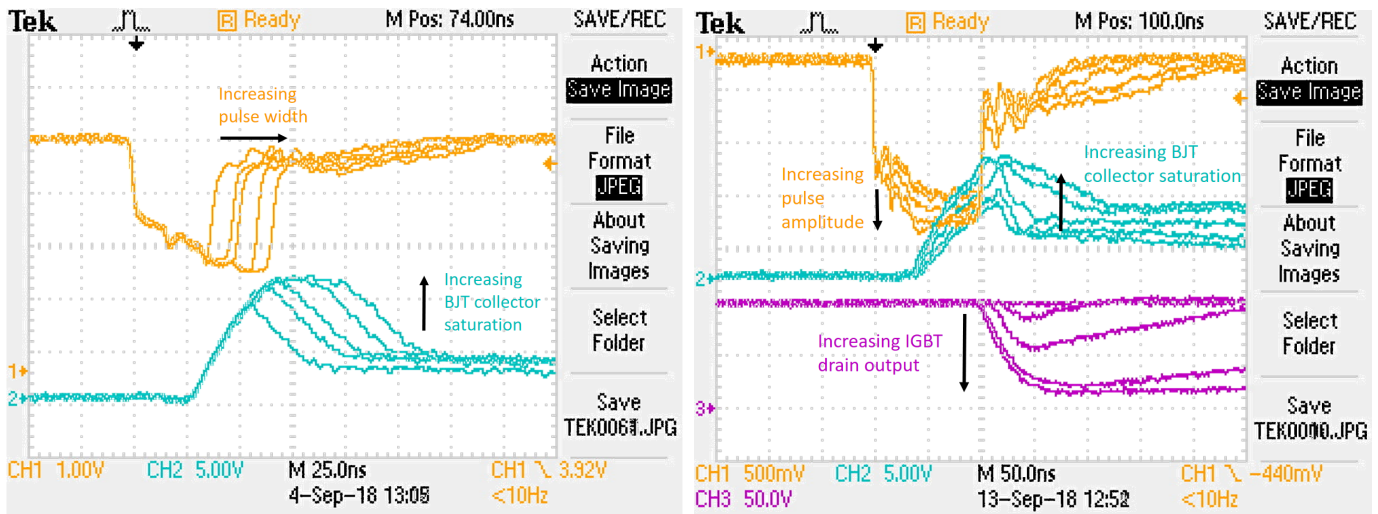
(a) A wider pulse increases the BJT pulse magnitude.

(b) A higher-amplitude pulse increases the BJT pulse magnitude and thus the IGBT pulse magnitude.

FIG. 21: Superimposed images displaying effects of adjusted input pulses on the amplifier circuit. The generated coincidence pulse is shown on top in yellow, the BJT collector is below in cyan, and the IGBT drain is shown at bottom in purple.
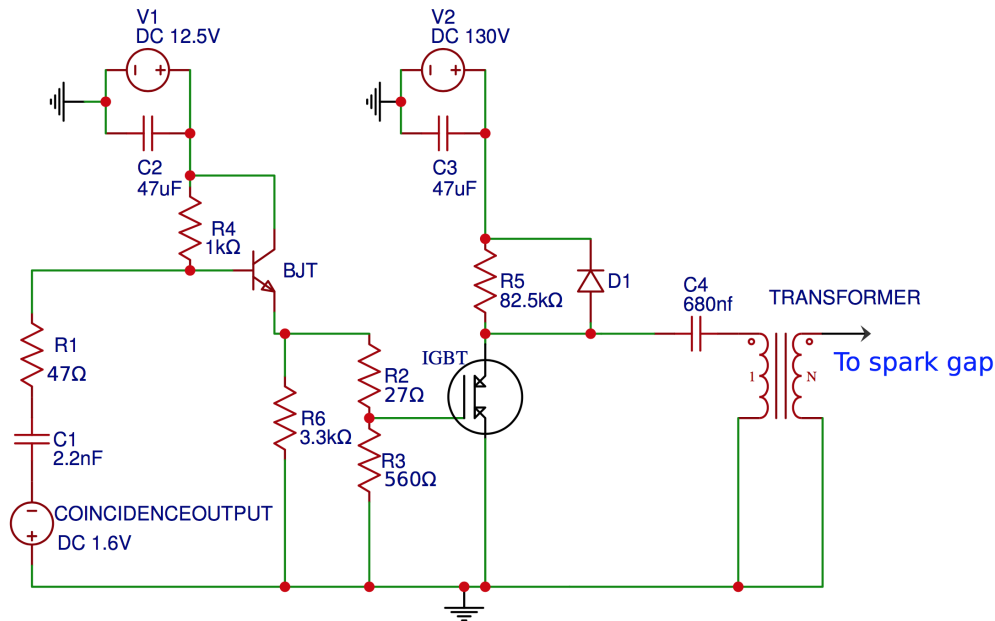

FIG. 22: The circuit diagram of the amplifier circuit. The pulse from the coincidence unit causes the voltage difference between the base and emitter of the BJT to reach threshold value. Thus, the switch closes and the BJT drops V1 from emitter to collector. This in turn causes the gate of the IGBT to reach its threshold voltage and close the switch between drain and source. Then, there is a path to ground for +V2, which causes C4 to discharge and deliver the newly amplified pulse to the next stage: the transformer.

When the base and emitter are at the same voltage (V_{BE} = 0), the switch is open; no current flows between emitter and base, so no current flows between emitter and collector. When V_{BE} reaches a threshold value, which we found through experimentation to be roughly 1.5V, there is a current between emitter and base, leading to a larger current between emitter and collector; the switch is closed. Ideally, the BJT will drop V_E across the collector load [5] [6]. In our amplifier circuit, the pulse from the coincidence unit causes V_{BE} to reach the threshold
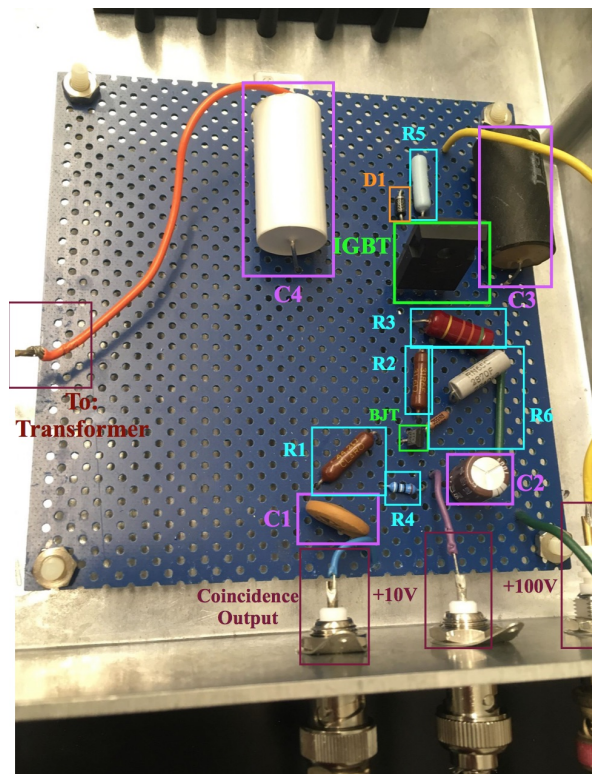
FIG. 23: A labeled picture of our completed amplifier circuit. The green wire in the bottom right corner indicates ground connections, and it is attached to the shield of the SHV jack (mostly cut off in this picture) delivering the +100V input. Note that care was taken while building the circuit to leave ample space to isolate the portion of the circuit at high voltage.

value, which closes the switch and drops the 12.5V from V1 across to the emitter, ultimately putting a positive voltage on the gate of the IGBT.

The IGBT (S2)

IGBTs behave like a combination between a BJT and a MOSFET, and thus are useful in high-voltage, high-current situations. They contain a gate that is electrically isolated from the drain and the source. When the gate is below its threshold voltage, the drain and source are insulated from one another, and no current can flow. When the gate reaches its threshold, the connection between drain and source becomes conducting [5] [6] . In our amplifier circuit, this creates a path to ground, discharges C4, and causes a voltage change across the transformer primary.

For an explanation of each circuit component's purpose, see A.

2. The Input Pulse

To cause the BJT to perform optimally, we had to tailor the input pulse coming from the logic unit. To determine optimal pulse conditions, we used a pulse generator to trigger the amplifier circuit instead of using the muon trigger setup and logic unit. We tested circuit behavior while varying the pulse in 3 ways: width, amplitude, and frequency. We found that the minimum width for the input pulse was 70 ns (Fig. 21a). Shorter pulses did not allow the BJT switch to remain open for long enough to apply voltage to the gate of the IGBT for sufficient time for the output of the circuit to reach its optimal minimum. From LTspice simulations, we discovered that longer pulses would cause the circuit to become oversaturated and draw too high a current for the IGBT.

In addition, we found that the BJT required an input pulse with an amplitude of at least 2V for optimal performance, though the circuit behaved passably with a 1.7V pulse (Fig. 21b). Thus, 2V is the threshold value for V_{BE} in order for the switch to close. However, the output of the logic unit is only 1.3V. In order to address this, we needed an intermediate amplifier to increase the coincidence pulse amplitude to that threshold. We are currently using a potentiometer to increase the cable impedance and thus to increase the voltage amplitude. However, this solution mildly distorts the pulse shape. As the project progresses, we would like to replace the potentiometer with a small circuit created using resistors and a differential amplifier chip.

Further, the frequency at which the amplifier circuit is pulsed has a pronounced effect on its behavior. When pulsed at a higher frequency, the IGBT output was observed to reach a higher magnitude than when pulsed at a lower frequency. This is possibly because with a more frequent pulse, key capacitors do not have time to discharge fully, causing the IGBT output to be more saturated. Similarly, we found that, when testing the circuitry, it took a few pulses to reach maximum magnitudes after all inputs were turned on. This is likely because capacitors take a few pulses to charge up. However, a few limiting factors to how much one should increase pulsing frequency exist. Too-frequent pulsing could damage the circuitry.

3. Timing and Output

Because the entire circuit must deliver a 6-10kV pulse to the chamber before gas recombination occurs (ideally under 500 ns), timing is key; every nanosecond counts. According to the data sheets for the parts we purchased, the activation time of the BJT can be as short as 45 ns, and that of the IGBT as little as 45 ns. Thus, under ideal circumstances, the circuit should produce an output 90 ns after the coincidence pulse, or a bit more slowly if taking into account cable lengths. In an earlier iteration of the amplifier circuit, the timing was mysteriously slow;
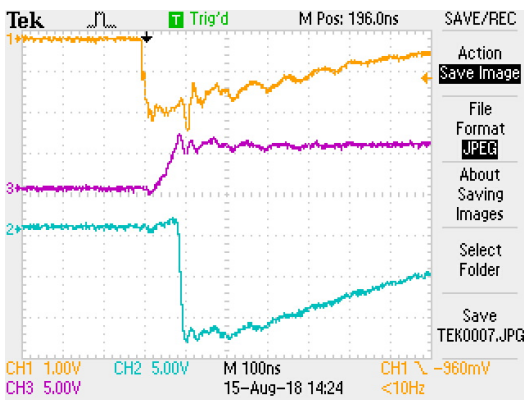
FIG. 24: Oscilloscope traces of an amplifier circuit test. The yellow trace is the coincidence pulse, the purple is the IGBT gate, and the blue is the circuit output. This test was conducted with V2=20V rather than its usual 130V so as to not damage the oscilloscope.
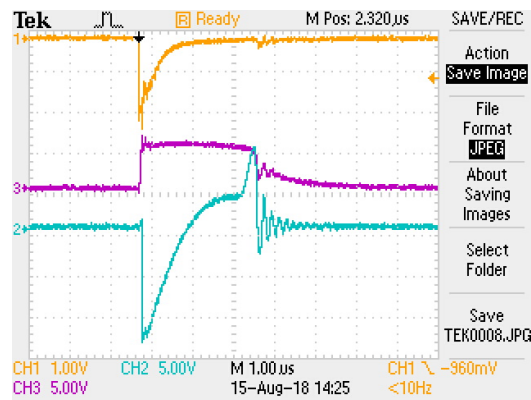

FIG. 25: The same amplifier test as in 24 but shown here on a much larger timescale in order to show longer-term circuit behavior. The color traces represent the same probing locations as in 24. Note the irregular shape of the coincidence pulse in both images.

the amplified output pulse occurred 250 ns after the coincidence pulse. By probing at the outputs of the BJT and IGBT, we were able to find that the vast majority of the delay happened at the IGBT stage. Attempts to tweak the circuit to cause faster behavior failed. Thus, we built a replica of the circuit on a breadboard in order to more easily switch out components to find optimal values for the resistors and capacitors. The breadboard circuit was fairly finicky; sometimes it was slow, sometimes the same combination of resistors and capacitors produced faster results, and sometimes we inadvertently disconnected components, causing it to stop working altogether. It did, however, convince us that we were satisfied with the existing combination of components for the circuit. We then rebuilt the amplifier circuit, this time on a plastic PCB with no soldering connections. We took care to space components adequately and to create a clear division between the side of the circuit with 12V input and the side with 50V. Testing this new amplifier circuit revealed much improved timing; the output now reaches its peak in under 150 ns. In addition, the pulse becomes steeper as the input voltage from the 50V supply increases.

B. Transformer

The transformer has proven to be the most influential and challenging component of the HV circuit with regards to minimizing the time delay. While the coincidence circuit and amplifier circuit contributed 60 ns and 150 ns, respectively, several transformers and transformer arrangements initially held the circuit to a delay on the order of microseconds. The current transformer, under the right conditions, can trigger the spark gap within 410 ns of the coincidence pulse and within 500 ns of muon incidence.

We have used a high-voltage DC voltmeter probe for cursory oscilloscope measurements of the transformer secondary, but it has been unreliable for both qualitative and quantitative measurements, often showing extensive noise without any input or nearby sources. Instead, tracing the noise of a spark on existing scope inputs has allowed for quantifiable evidence of the circuit timing.

Though we originally followed the example of the Cambridge Lab and wound our own copper-wire transformers, they were plagued with a range of problems, including internal sparking, over-delicate and highly breakable wires, and a practically nonexistant load resistance. The ones that did function only reached a voltage transformation of 30:1 rather than the intended 45:1, and provided several microseconds of delay.

We instead tested several commercial transformers in place of our hand-wound ones. One previous commercial transformer has reliably produced an initial spark from spark plug to shield within 900-1100 ns and remains a workable, if slow, option (seen with spark gap in Fig. 27. The current transformer has proven somewhat finicky but quite fast. It was moved several times—soldered in a circuit board, free-floating, clamped to the spark plug, and finally resting on top of the spark plug—and each different iteration came with its own issues. The circuitboard and free-floating operations were concerning and unsustainable on the delicate leads as they were not constructed to relieve strain. Our first solution, clamping the transformer atop the ceramic (insulated) portion of the spark plug, adequately secured the transformer. The metal loop created by the clamp, however, negatively affected the transformer workings, resulting in a $1.5\,\mu s$ delay until the initial spark. The transformer is stable without the clamp (though additional measures to secure it may be taken), and the clamp's removal has allowed us to observe initial sparks below 500 ns and as low as 410 ns with a 130-V input (Fig. 26).
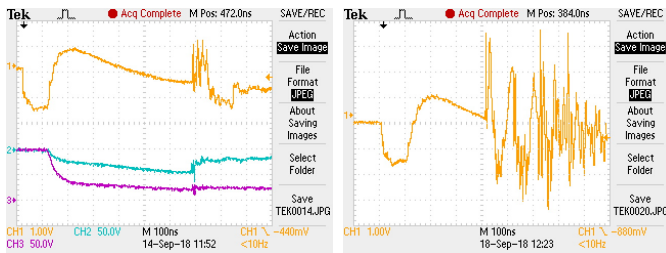
(a) Coincidence pulse and transformer primary, showing a spark as feedback at 680 ns.

(b) One of the faster observed sparks at 410 ns.

FIG. 26: Two oscilloscope pictures depicting spark delays. (Left) shows only the primary spark from iridium tip to shield; (right) shows a secondary spark across the large spark gap, with 2.2 kV applied to the anode pin. The potential difference at this voltage sits so near to breakdown that the secondary spark is practically instantaneous after the first, and visually indistinguishable.

C. Spark gap circuit

The spark gap acts as a switch that closes upon reception of a sufficiently large pulse from the transformer and discharges the capacitors into the spark chamber plates.

Commercial thyratrons and spark gaps are both viable options for spark chamber circuit switches with minimal delay times. Thyratrons are well-documented and can be purchased with the desired short delay times, but the inherent low internal impedance of spark gaps makes them a cheap and reliable option for a fast-switch circuit [2][7]. Thyratrons have also historically had shorter lifetimes than spark gaps [2]. We were able to cheaply construct a small and contained spark gap with a common automobile spark plug and nonspecialized materials in-house and thus opted for the spark gap over a more expensive thyratron.

Our spark gap construction and accompanying circuit diagram both resemble the Cambridge group's designs (Fig. 27, Fig. 28). There are two locations to which the iridium tip cathode may spark: the grounded spark plug shield, and the adjustable screw that acts as an anode.

This two-spark system allows one to control the magnitude of the pulse discharged across the chamber plates. The anode screw is held at a large positive potential (without sparking to the cathode). A transformer input provides enough voltage to spark between the cathode spark plug tip and its shield, and the ultraviolet light from that primary spark sufficiently ionizes the air in the secondary gap to cause a large, robust spark between the anode and cathode. Adjustment of the distance between the electrodes alters the threshold breakdown voltage for the secondary gap, with a larger gap corresponding to a higher breakdown voltage. Increasing the gap thus allows the chamber-side of the circuit to be maintained at a higher potential before the spark gap fires and therefore
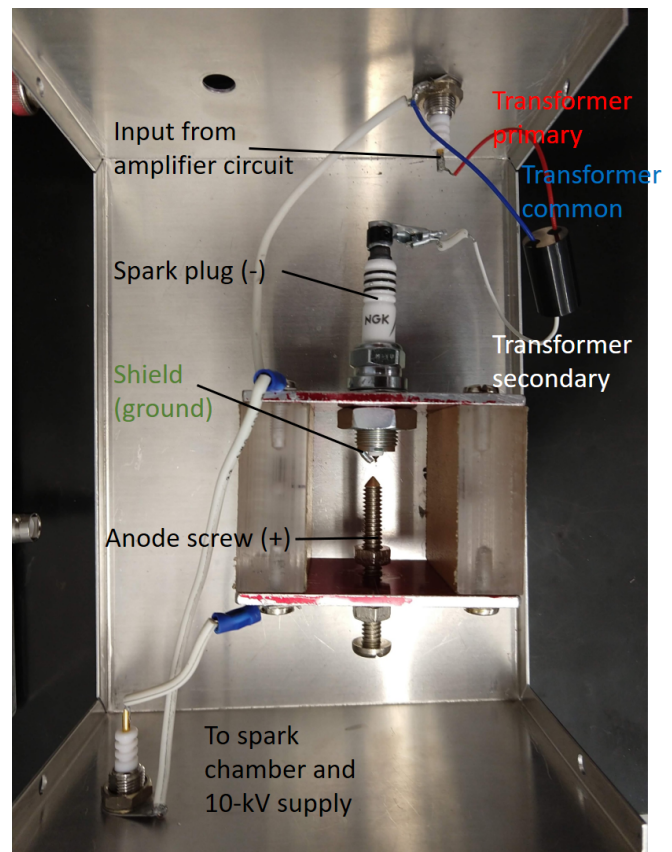

FIG. 27: Homemade spark gap with a previous transformer. The anode screw on the bottom is maintained at a high positive voltage; the screw is adjustable to calibrate the spark gap. The spark plug sits inside a grounded metal plate, electrically connected to the shield, and the iridium tip cathode receives a negative input pulse from the transformer.

increases the discharge on the plates when it does.

The housing for the spark gap was designed with a turnable pin as the anode to accommodate an adjustable anode separation. Figs. 29 and 30 serve to calibrate this distance and operational voltage to optimize timing and match the spark gap to the necessary chamber pulse.

Fig. 29 displays the relation between the spark gap separation distance and the breakdown voltage of the secondary spark at atmospheric pressure, with the primary spark as a trigger. This data is a situational look at the Paschen curve of air with an add ultraviolet-light trigger; see Section VI for further explanation.

Fig. 30 shows faster timing with application of additional voltage above the threshold. The Cambridge Lab observes similar behavior and attributes the relation to complex interactions in the breakdown mechanism beyond the scope of this paper. The gap distance of 0.11in for this trial was chosen during an argon flush of the test chamber to allow a full range of voltages accessible on our 10-kV HV supply. This gap also allowed for sparking of the chamber and observation of the extremes, including
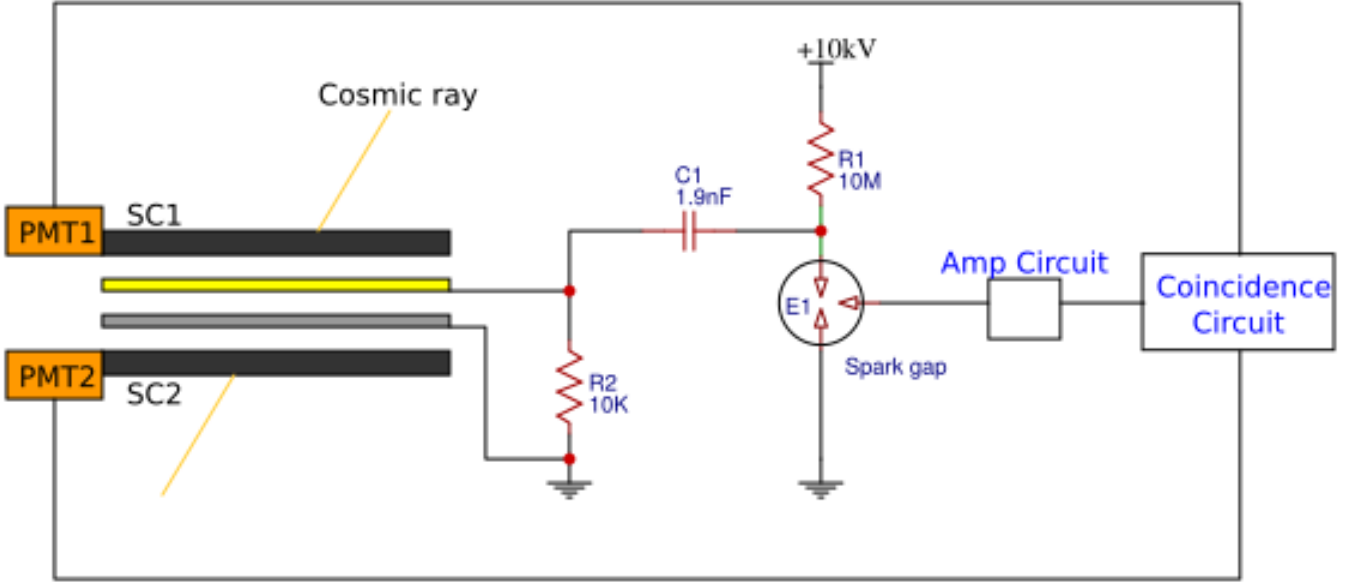
FIG. 28: Simplified circuit diagram for spark-gap system, displaying only one HV-ground plate pair.

pushing the voltage such that the spark gap arced constantly, regardless of the primary trigger; this behavior occurred at 8.2 kV for this distance. Both the distance and voltage calibrations will inform our decisions as we assemble the final chamber and re-implement the circuit.

We also observed whether the spark gap distance affected the timing of the pulse (Fig. 31). There is no clear relation from this data, but one may interpret a mild upward trend. Since we are addressing the secondary spark, the contributing factors are the transformer pulse to the iridium cathode, the voltage applied to the anode, and the breakdown voltage of the gap, dependent on the gap separation. An increased spark gap distance does not affect the speed of the transformer pulse but increases the breakdown voltage of the gap. To reach that higher threshold, the transformer takes longer to produce a higher-amplitude pulse—but we also have maintained the anode at a higher positive voltage. Interpretation of an upward trend would imply that the transformer behavior must cover more ground than the increased anode potential addresses.

Since a sluggish voltage pulse may cause ionized electrons to drift and clear from the chamber gaps, a very fast rise-pulse is desired [1]. To minimize the capacitor discharge time at the plates, the circuit elements were attached directly to the spark chamber's inputs. The capacitors were attached to the O-ring screws by a series of thread adapters. This seemingly convoluted engineering choice was a result of the surprisingly limited selection of thread sizes for screw-on capacitors and the need for sufficiently small tapping screws to fit stably into the O-ring screws (refer to Appendix for detailed construction). On the ground plates without capacitors, similar taps were used to allow attachment of resistors to ground.
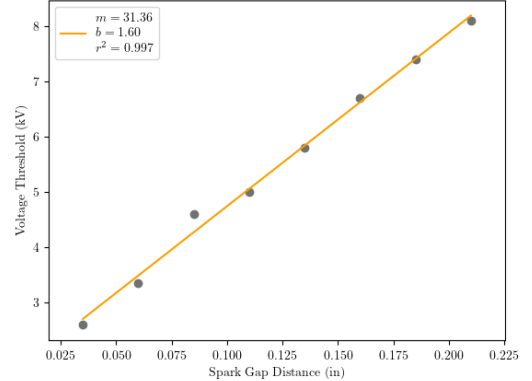

FIG. 29: Calibration of the spark gap separation distance. Increasing the spark gap distance increases the voltage required to break down the air between the anode and cathode. This relation informs the operational distance of the spark gap. Our chamber will operate at least at the breakdown voltage of the chamber gaps, which is dependent on the choice of enclosed gas, and the spark gap will be tuned accordingly.

The initial assembly of this design used copper tape as a high-voltage and ground bus in an elegant but somewhat impermanent arrangement (Fig. [32]). Subsequent iterations opted for the use of split coaxial cables with floating grounds on the capacitor ends in order to minimize noise and external sparking.
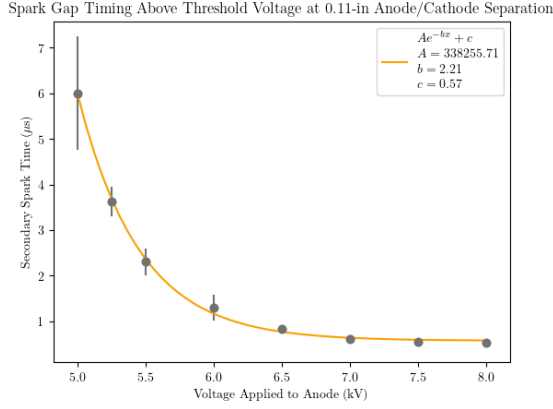
FIG. 30: Calibration of the spark gap applied voltage, taken at a spark gap separation of 0.11 in. Increasing the voltage applied to the chamber and anode will improve the timing drastically with diminishing returns. Uncertainty is given by the standard deviation of the time measurement across ten trials.
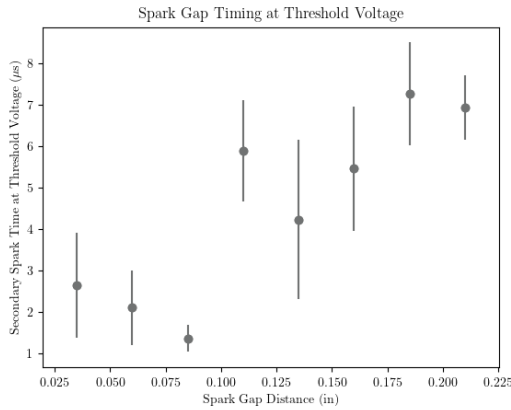

FIG. 31: Plot of secondary spark gap timing at increasing spark gap separation distances. A vague upward trend may be observed but is unclear. Uncertainty is given by the standard deviation of the time measurement across ten trials.

D. Grounding

The grounding of our circuit is depicted in Fig. 33. The three voltage supplies stem from a single power strip in order to minimize the impact of ground loops. In an initial design, the ground plates and spark gap circuit were grounded through the 10-kV supply, while the pulse circuit (amplifier circuit, transformer, and spark gap itself) was grounded through the 100-V supply; the two sides of the circuit were split in order to prevent a large ground loop. To further reduce noise, the entire high-voltage circuit is enclosed in a metal box, with the high-voltage transformer and spark gap also isolated in

FIG. 32: Evolution of the back face of the box with prototyped capacitor-adaptor constructions from an early design with copper tape buses to threaded capacitors (top) to a more secure construction using coaxial cables (bottom, left and right).

another metal box. Testing of this set-up with the test box resulted in sparking across the box, due to the different grounds. At present, all grounds are consistent with the metal enclosure, with no significant noise observed due to the ground loop.

The metal enclosure proved critical to the circuit's functionality as a muon-triggered pulse circuit. In the first attempts to spark the chamber, extreme noise prevented accurate characterization of the pulse circuit timing. Characterization involved viewing a coincidence pulse on an oscilloscope and observing when large feedback occurred to measure the total delay between muon detection and delivery of 10 kV to the chamber plates. The amplifier and spark gap circuits were housed in separate plastic boxes rather than metal, and we observed as little as 10-20 ns of "delay" between coincidence and sparks. Such an unreasonably fast value for the pulse circuit suggested that the chamber was self-triggering: inadequate shielding meant that feedback from the chamber sparks re-triggered the pulse circuit even without a true coincidence. Initial attempts to remedy this issue by covering the circuit's plastic enclosure with grounded aluminum foil had little to no effect on the circuit behavior. Once the pulse circuit was relocated to a sturdy metal enclosure, this self-triggering effect was not observed.
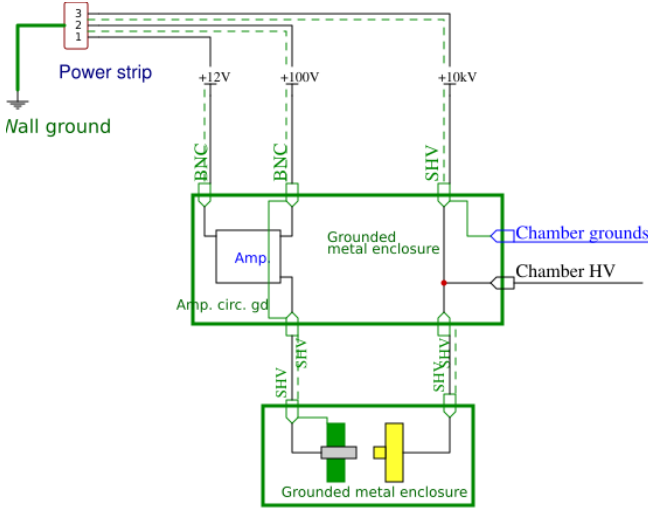
FIG. 33: Circuit grounding schematic. Dashed lines indicate shielded cable. Green indicates ground. Note that the 12V supply is grounded in the other HV supplies rather than directly to the wall. The 100V supply and 10kV supply create a ground loop that has not been problematic.

VI.   GAS PURITY

A.   Theory

The operation of spark chambers relies on the principles of electrical breakdown and Townsend avalanches. Electrical breakdown in a gas refers to polarization of molecules by a strong electric field that effectively transforms an insulating gas into a conductive material. By applying a large potential difference to adjacent metal plates, a spark chamber causes electrical breakdown along the ionized path of an incident muon. The free gaseous electrons accelerate and collide with more gas molecules, resulting in a cascading avalanche of more freed electrons, known as a Townsend avalanche. The path of the muon between the plates is then visible as a bright spark.

Different gases break down at different voltage thresholds. The breakdown voltage of a particular gas is related to its pressure and the distance between the plates by Paschen's law [8]:

$$V_b = \frac{Bpd}{ln(Apd) - ln[ln(1 + \frac{1}{\gamma_{se}})]}$$

V_b represents the voltage threshold where the gas experiences electrical breakdown. A and B represent gas-specific constants determined experimentally, and γ_{se} represents the secondary electron emission coefficient, indicating how many electrons are freed per incident positive ion.

Plotting Paschen's law as breakdown voltage threshold versus pressure × distance pd forms a roughly V-shaped curve, as seen in Fig. 34. To the left of the curve's minimum, where pd is small, the mean free path of an electron will be long compared to the distance between the plates, meaning that the electron will collide with few other molecules before it reaches the anode. This results in a higher breakdown voltage to be required for a Townsend avalanche. On the right side of the minimum where pd is large, the mean free path of an electron is short relative to the plate separation, so electrons will lose more energy in non-ionizing collisions, therefore requiring a higher breakdown voltage for the avalanche and sparking. In this region, which is where we will be operating our chamber, the breakdown voltage is roughly proportional to pd [9].
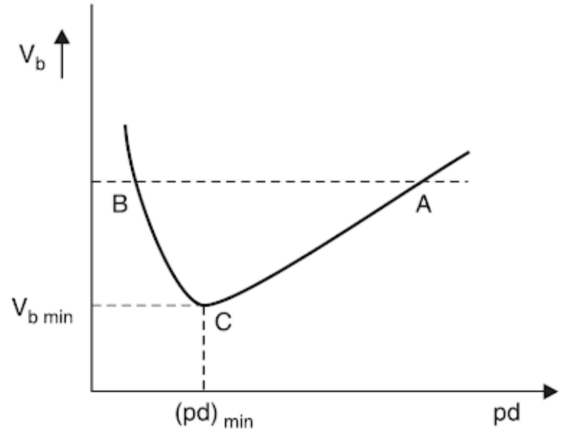

FIG. 34: Paschen Curve relating the breakdown voltage threshold of a gas to the pressure × distance between charged plates.[9]

B.   Gas Choices

We decided to experiment filling our chamber with helium and argon gas. Both gases have the benefit of a high first Townsend coefficient, indicating that many electrons will be freed during a Townsend avalanche. This means that helium and argon will have low voltage thresholds for electrical breakdown, relative to other gas mixtures such as air. Helium and argon are also affordable compared to other similarly performing noble gases.

So far, we have only tested the chamber filled with either helium gas or argon gas separately. As this project progresses, we intend to experiment with using a mixture of gases in the chamber. Helium breaks down at a lower voltage threshold than argon, as seen in Fig. 35. Argon gas has a lower ionization energy than helium, so the addition of argon increases the amount of free electrons during a Townsend avalanche. This should increase the time it takes for electrons to recombine with gas atoms, allowing more time for us to deliver a high voltage pulse before the ionized path of the muon disperses, therefore

increasing the efficiency of our chamber [1].

Spark chambers typically either have gas stagnant in the chamber, or constantly flowing through. We chose to have gas constantly flowing in order to reduce air contamination. The oxygen in air has a high affinity for electrons, which would reduce the amount of free electrons in a Townsend avalanche. This would give us less time to deliver a pulse to the chamber plates before the electrons are recombined with gas ions and there is no more ionized muon path to spark on.

C. Breakdown Voltage Measurements

We measured the voltages where the gas in our test chamber breaks down at different spark ratios, which we defined as the percentage of high voltage pulses sent to the chamber plates that result in chamber sparks. Every pulse corresponds to a muon detected passing through both scintillators.

Fig. 35 shows the spark ratio vs. voltage in the cases where the chamber was filled with helium and argon gas. This plot demonstrates that helium breaks down at lower voltages than argon. Both gases show a higher spark ratio when higher voltages are applied to the chamber plates. This result is consistent with our previous spark gap findings that indicated faster secondary sparks at higher applied voltage on the anode pin of the spark gap (Fig. 30). Therefore, we expect to see more muon tracks at higher voltages, as the pulse arrives before the electrons recombine with the gas ions. The uncertainty expressed in this plot is calculated from the number of pulses counted to determine the spark ratio for each data point. The average number of pulses counted per data point was 29.

FIG. 36: Sparks in our test chamber. The top photo is of the chamber sparking in helium gas, whereas the bottom is sparking in argon gas.
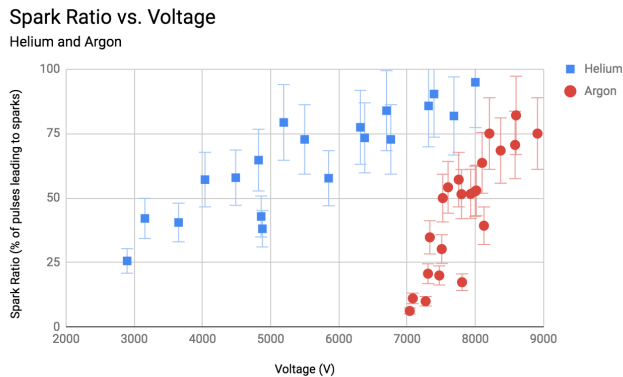

FIG. 35: Measuring spark ratio at different voltages in argon and helium. Helium is shown here to break down at lower voltages than argon. In both gases, spark ratio increases with voltage.

Sparks seen while the chamber was filled with argon gas were brighter than when the chamber was filled with helium, and sparks seen in argon were more blue than the pink sparks in helium, as seen in Fig. 36.
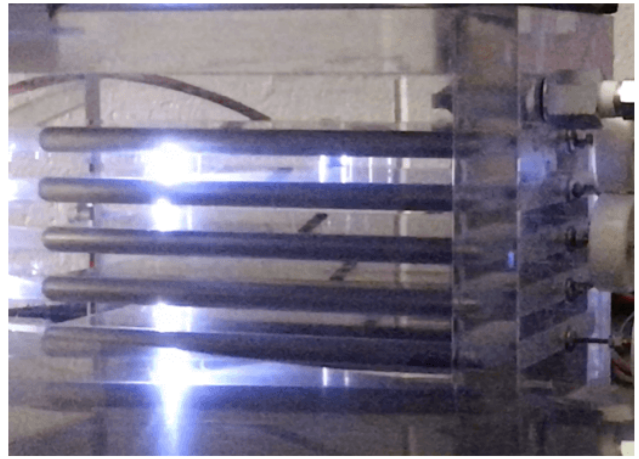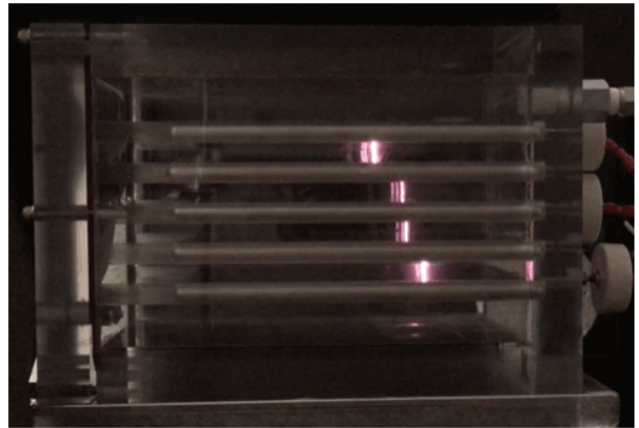
Fig. 37 shows a plot of the spark ratio vs. voltage for the cases where the argon-filled chamber is inside and then outside of the area where the scintillator paddles are overlapping. In the case where the test chamber is sitting outside of the scintillator area, high voltage pulses will not correspond with muons passing through the chamber. These trials aimed to illustrate the differences between spurious sparking and sparking on muon tracks. We see a higher spark ratio when the test chamber was inside the scintillator area. This means the chamber sparks more often when the pulses do correspond with muons in the chamber. From this, we can infer that when the chamber inside the scintillator area, not all of our sparks are spurious and that we are in fact seeing muon tracks. The uncertainty expressed in this plot is calculated from the number of pulses counted to determine the spark ratio for each data point. The average number of pulses counted per data point was 32.

When we attempted to conduct the same outside scintillator area tests in helium gas, the sparks were consistently on the same edges and we stopped the tests in order to prevent carbon traces from forming in those spots
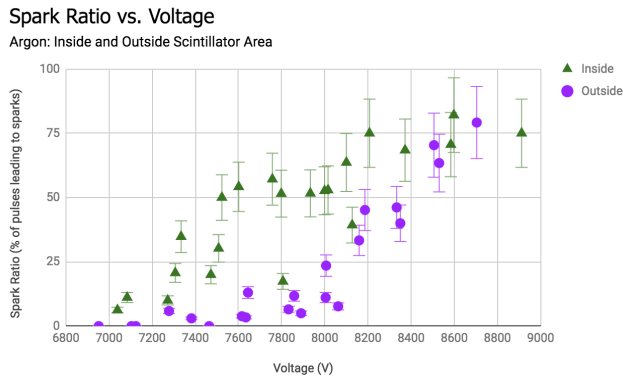
FIG. 37: Measuring spark ratio at different voltages in argon with the chamber inside and outside the scintillator area. The spark ratio is higher when the chamber is inside the scintillator area, implying that not all sparks in our chamber are spurious.

and damaging the chamber.

In the future, we plan to repeat these tests in the full scale chamber with more trials over larger voltage ranges. The larger chamber will have better gas flow and cleaning practices, leading to a higher gas purity.

D. Apparatus for Gas Flow

Our spark chamber is held at a pressure slightly above standard atmospheric pressure with gas constantly flowing through it. In order to gauge the flow rate and prevent air contamination, gas exiting the chamber is passed through a mineral oil bubbler, Fig. 38.

We chose our flow rate to be fast enough to flush out impurities, but slow enough as to conserve gas [2]. We currently operate our chamber with gas flowing at a rate where the bubbler produces one bubble per second.

VII. INTERACTIVITY AND DISPLAY

A. Interactivity

To prevent the (somewhat loud) chamber from detracting from any academic settings near its location, we are currently implementing several soft shut-offs using the Raspberry Pi (RPi). The RPi is connected to the high-voltage pulse circuit in two critical locations; it controls a DC volt- and current-regulator at the input of an in-progress 5V coincidence amplifier supply as well as at the 12V supply. A script allows the RPi to toggle each of these power supplies on and off and to control the precise input voltage. Since the RPi may be remotely accessed, the project is currently designing an interactive terminal or button to place beside the spark chamber display.
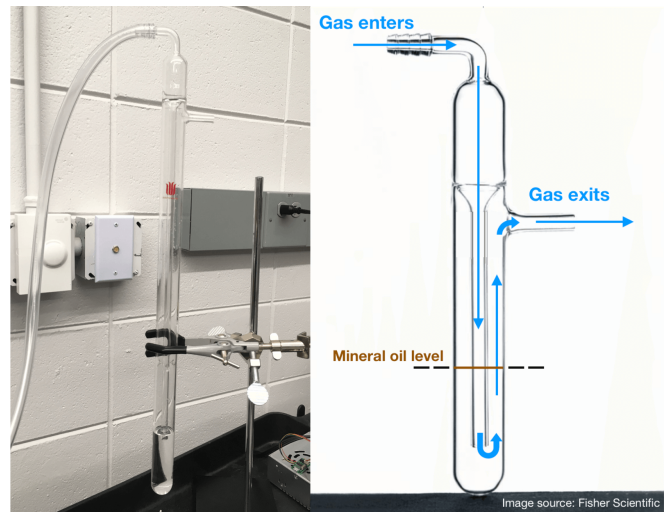

FIG. 38: Our mineral oil bubbler and a diagram of the gas flow. Gas exiting the chamber will enter the bubbler through the tube attached to the top and travels down the thin inner glass chute which is submerged in mineral oil at the bottom. Gas exits the chute at the bottom in bubbles which rise through the oil and exit through the horizontal exit tube near the top. The flow rate can be approximated by the rate of bubbles in the oil. This system prevents air from entering the chamber since any outside gas would not be heavy enough to sink through the heavy oil and enter the inner tube.

FIG. 39: The Raspberry Pi.

B. Installation

The project is currently working to install the chamber outside of the Maria Goeppert-Mayer lecture hall. The spark chamber will be unveiled there on November 8, 2018.

VIII.   CONCLUSIONS

From test box design to HV circuit struggles to gas flow characterization, we hope that this document provides a suitable reference for any hoping to construct their own tracking particle detector. Our shelf-style construction faced initial flaws regarding gas flow and edge sparking that are addressed in the in-progress, full-sized chamber. We are currently in the midst of assembly and will then migrate the constructed HV pulse circuit to the new chamber before testing and installing it. Testing of the gas flow and understanding the impact of gas composition on chamber lifetime will continue in parallel.

One of the most rewarding accomplishments of this project was the successful delivery of a high-voltage pulse to an ionized muon track in the test chamber within the 500 ns recombination limit. The majority of our time was devoted to understanding, assembling, debugging, remaking, and re-understanding the coincidence unit, HV pulse circuit, and spark gap. We operated in the regime of microseconds for most of the project, far from the speed necessary to spark on muon tracks. Troubleshooting the amplifier circuit and selecting the proper transformer were crucial in finally optimizing the pulse speed.

We are incredibly excited to have put together a functional cosmic ray spark chamber prototype in the short span of one summer. Each of us entered with our own minimal background in electronics, circuitry, programming, and machining, and emerged with a significantly better understanding of both the theory and methods behind particle physics projects. We can now proclaim with confidence our identity as the proud parents of a fully home-grown spark chamber.

IX.   ACKNOWLEDGEMENTS

[1] J. Collins, *Construction of a Prototype Spark Chamber*, University of Cambridge, October 2009.
[2] M. Oreglia, Personal correspondence, 2018.
[3] M. Collier and L. Wolfley, *Assembly Manual for the Berkeley Lab Cosmic Ray Detector*, Lawrence Berkeley National Laboratory, October 2006.
[4] "Ltspice," Analog Devices. [Online]. Available: http://www.analog.com/en/design-center/design-tools-and-calculators/ltspice-simulator.html
[5] P. de Grouchy, *Construction and Evaluation of a Fast Switching Trigger Circuit for a Cosmic Ray Detection Spark Chamber*, University of Cambridge.
[6] P. Horowitz and W. Hill, *The Art of Electronics*, 3rd ed. Cambridge University Press, 2015.
[7] J. W. Cronin, *Spark Chambers*, Princeton University, 1967.
[8] M. A. Lieberman and A. J. Lichtenberg, *Principles of Plasma Discharges and Materials Processing*, 2nd ed. John Wiley & Sons, Inc., 2005.
[9] C. Wadhwa, *High Voltage Engineering.* New Age International, 2007.

Appendix A: Drawings & Designs

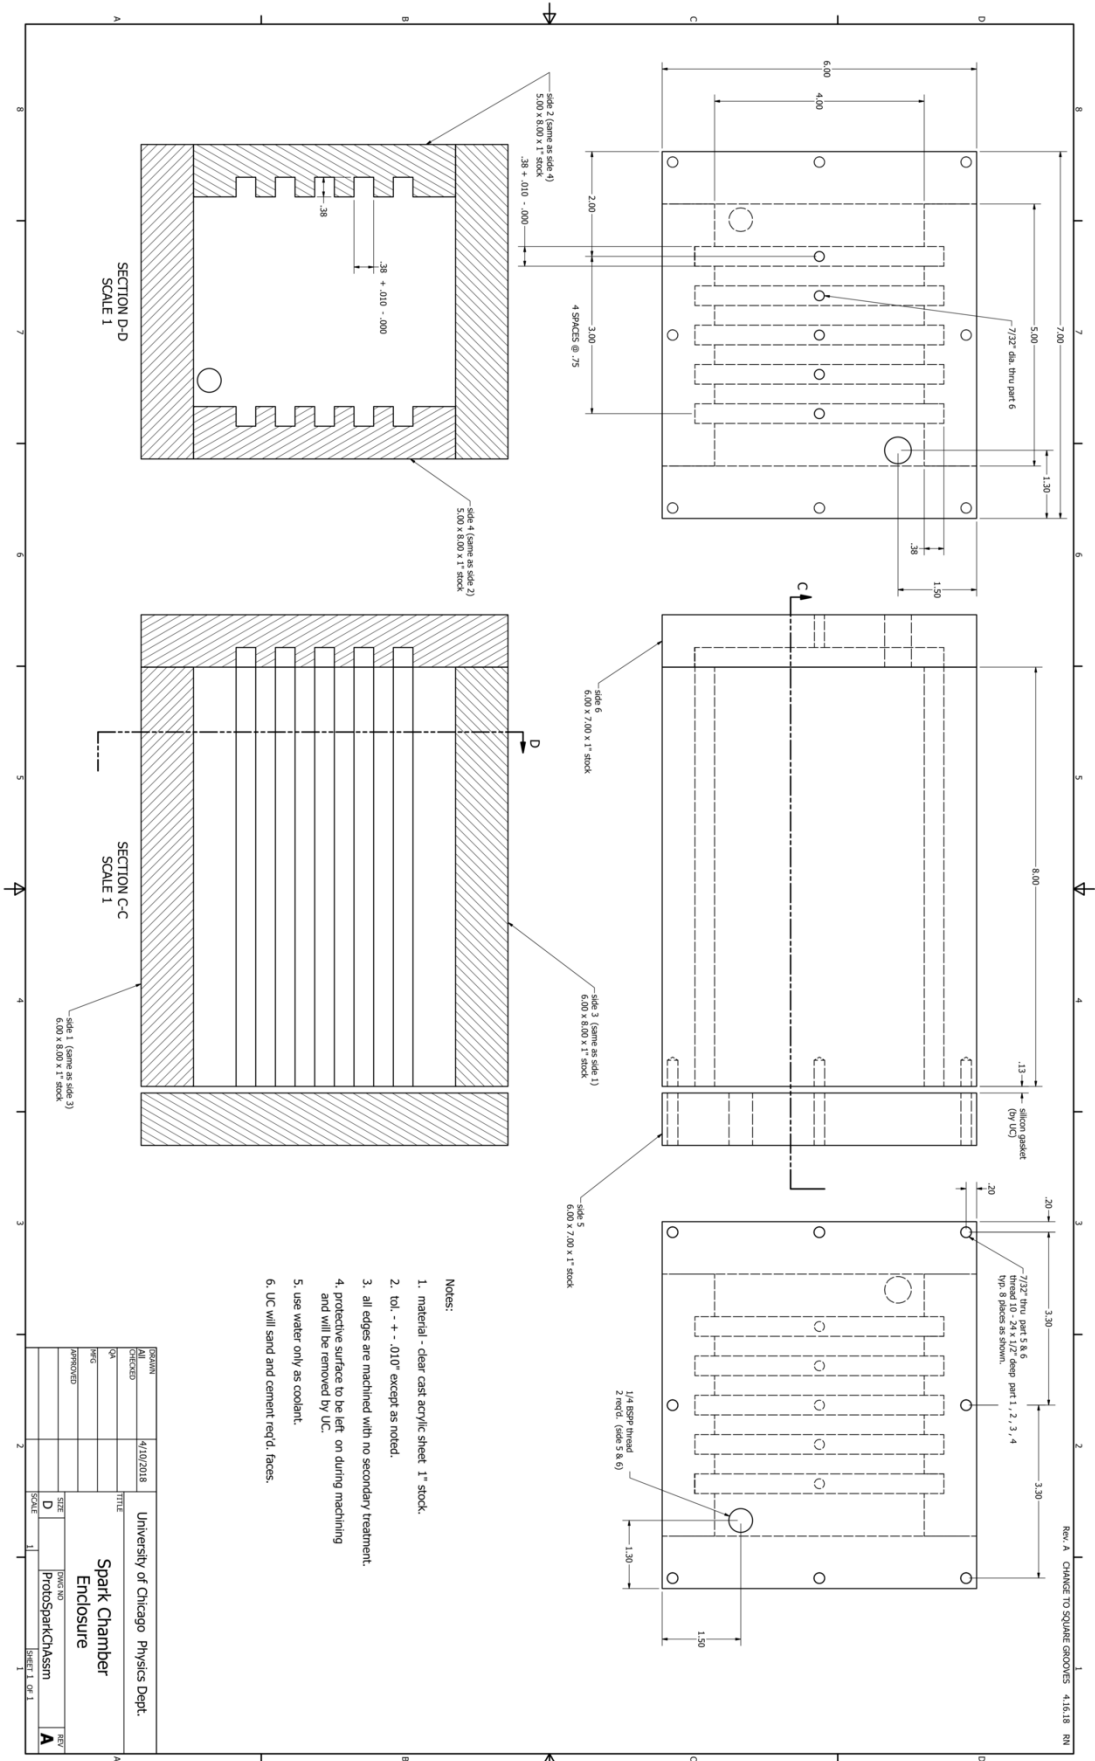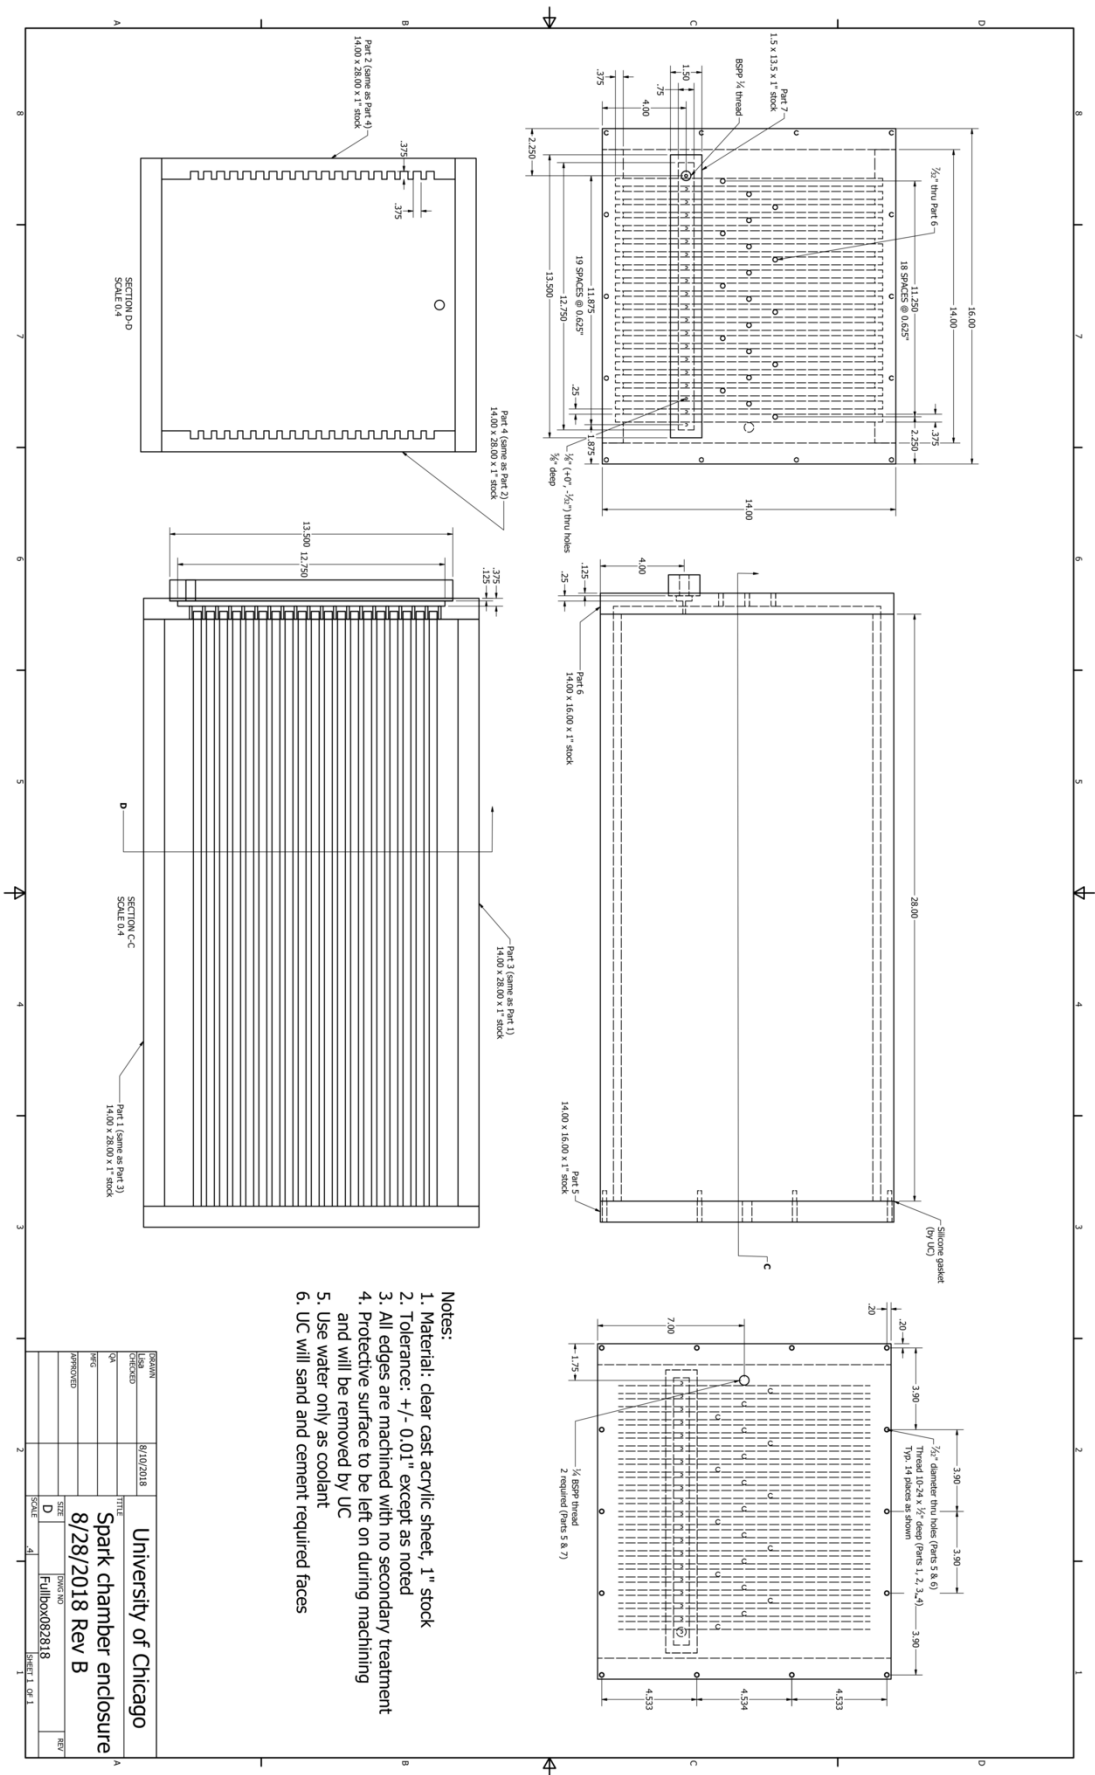Full assembly CADs of the test box and full chamber begin on the following page.

FIG. 40: Assembly drawing of the test chamber.

FIG. 41: Assembly drawing of the full spark chamber.

Appendix B: Instructional notes on scintillation counter construction

1. Cleaning the PMTs

The purpose of the optical cement is to better match the indices of refraction of the glass PMT lens and the acrylic scintillation material so light does not scatter as much at their interface, and scratches in the PMT lens would cause light to reflect off and not enter into the photo tube. Notwithstanding, the layer of optical cement left on the Cronin PMTs was thicker than you might think; holding the razor blade at a slight angle will make the scraping process much faster and easier, and will not do any harm as long as you guide the blade with a bit of care.

2. Polishing the scintillators

We highly recommend wearing gloves while polishing the scintillator, which makes cleaning up all the white dust from the sanding process that much easier (that stuff gets *everywhere*).

600 grit: Polish away all indents or imperfections in the scintillator with a small square of 600 grit sandpaper by holding the scintillator upright with the cut side facing upwards and rubbing a slightly dampened sheet of sandpaper back and forth laterally across the surface. Cutting out and using one square of sandpaper at a time provides greater control over the direction of sanding and helps minimize potential scratching on the sides of the scintillator by the sandpaper. Continue polishing until you are left only with even striations from the sandpaper all going in the same direction. Wipe the surface of the scintillator every so often with a slightly dampened cloth to clear the surface of the plastic scintillator dust from sanding.

800 grit: Next, polish the scintillator with a small square of 800-grit sandpaper in the crosswise direction (in the direction perpendicular to the direction of polishing for the 600-grit stage), until all striations left by the previous stage of sanding have been replaced by uniform up-and-down striations.

1200 grit: Finally, polish the scintillator with a square of 1200-grit sandpaper, rubbing from side-to-side in the same orientation as the 600-grit polishing. Continue polishing until all grooves from the 800-grit stage have been completely removed. At this point, you should be able to tilt the scintillator slightly and see the other scintillator sides through the freshly polished surface. After polishing, wash the scintillator with soap and water using a gentle sponge, and dry thoroughly.

3. Wrapping the scintillators

Take a large sheet of aluminum foil and flatten it out on a table - get a big sheet, larger than what you need. Place one scintillator on top of the foil, and lightly trace its position with a pen or pencil. Remove and cut out a 5-5.5" diameter hole in the center of this rectangle, or whatever size the circle of insulation is on your PMT. We want to prevent the foil from going anywhere near the end of the PMT that will be attached to the scintillator, as that end will be held at high voltage when in use. Replace the scintillator, and wrap it tightly with the foil, like a present. Use Scotch tape to seal the edges of the foil and hold it firmly in place.

4. Coupling the PMTs with scintillators

At this point, mix your optical cement according to manufacturer instructions, and let sit for 15-20 min for the air bubbles to rise. In the meantime, prepare your scintillation surface for application of the optical cement, again according to manufacturer instructions. For us, this meant lightly sanding the exposed scintillator with 400-grit sandpaper, cleaning the spot with isopropanol, and letting it dry, then gently cleaning off the PMT lens with a microfiber lens (glasses) cleaning cloth. Position the PMT over the sanded gap left in the foil, aligning it so that the cardboard is uniformly just shy of the edges of the scintillator (Fig. 17, right). These slivers of extra space between the cardboard and the scintillator will allow the tape used to seal the assembly together to make good contact with the scintillator edges and subsequently provide a sturdier seal. Tape down one side of the PMT's plastic insulation. This holds it in place while applying the optical cement.

When ready, apply a thin layer of optical cement to the surface of the PMT, and carefully set the PMT in place. Secure the circle of insulation surrounding the PMT in place on the scintillator with electrical tape, making sure to cover all areas of exposed scintillator. Then place strips of tape, curled back on themselves, all around the surface of the foil; these will help the cardboard adhere to the scintillator. Position the cardboard in place, and press down firmly. Let the scintillation assemblies sit undisturbed for 24-48 hours, or however much time is specified by the manufacturer of your optical cement. We moved onto the steps delineated below the morning after cementing the PMTs in place, as we were running a little short on time, but such practice is strongly disadvised.

When the PMTs have fully cured, place strips of scotch tape, curled back on themselves, all around the surface of the cardboard piece covered with black plastic (with the plastic side facing up). Align the cardboard with the bottom of the foil-wrapped scintillator such that the cardboard is again, just slightly shorter than the scintillator on all sides; press firmly to adhere.

Finally (at last! Hallelujah!), affix thick black electrical

tape on the sides of the scintillator to seal the whole deal together, like sprinkles on the sides of an ice cream sandwich (Fig. 18). If you left enough space for the edges of the scintillator to peek out, your connection should be fairly strong.

Appendix C: Explanation of Amplifier Circuit Components' Functions

1.   Resistors, Capacitors, and Diodes

R1 and C1: Together, this resistor and capacitor combination determines the RC constant for the input pulse going into the BJT. R1 also limits the amount of current going into the BJT so as to not exceed its current rating.

R2 and R3: These resistors form a voltage divider to decrease the voltage going into the gate of the IGBT.

R4: This resistor ensures that, when the pulse from the logic unit decreases the voltage at the base of the BJT, the emitter does not immediately experience the same voltage drop. Preserving this voltage difference between base and emitter causes the BJT to switch "on" and allow current to flow from emitter to collector.

R5: This resistor limits the current so as to protect the power supply that inputs V2.

R6: We observed no changes in circuit behavior between when R6 was included and when it was absent. Nonetheless, we kept it in the final circuit.

C2 and C3: These capacitors reduce noise and prevent shorting of our two power supplies.

C4: This capacitor charges and then dumps its charge into the transformer when the amplified coincidence pulse closes the switch between drain and source, which pulls the drain to ground. Thus, it delivers the pulse to the transformer.

D1: This diode prevents back emf from the transformer due to Lenz's law, which states that a change in magnetic flux causes an induced emf, and thus a current, with the opposite sign. In the case of the transformer, changing magnetic flux from the primary induces an emf and current in the secondary, which in turn changes the magnetic flux through the primary, inducing an emf and current in the opposite direction. Without D1 present in the amplifier circuit, this "backwards" current could flow back through the circuit and possibly damage power supplies.

D2: We omitted D2 because the particular IGBT we purchased has an internal diode serving the same purpose.

2.   Power Supplies

+12.5V supply: The 12.5V supply holds the base and emitter at the same voltage (12.5V) until the pulse from the logic unit drops the base voltage.

+50V supply: The 50V supply provides the voltage that drops from the drain to the source of the IGBT when the gate reaches its threshold voltage and closes the switch.

Appendix D: Procedure for assembly of capacitor-plate adapter using O-ring screws

1. Cut a 1" long section of threaded 2-56 rod with edge clippers.

2. Clean up the ends of the segment on a belt sander and thread a 2-56 nut over its entire length to redefine end threads.

3. In a lathe, shear and face off the short (6-32) end of a 6-32 to M4 adapter, then drill a 0.350"-deep hole in the adapter hex-nut. Tap the hole and clean it out with compressed air; then run the tap in and out of the length of the hole several times. Twist in the 2-56 rod by hand as far as it will go, around 0.150."

4. Screw a 2-56 nut onto the open end of the rod.

5. Drill a 0.500"-deep hole in a 10-24 threaded 1" O-ring screw; tap. Clean out the tapped hole very well, and thread in the other end of the 2-56 rod as deep as possible by hand, around 0.350." Turning the 2-56 rod must be done by hand, in order to preserve the integrity of the rod threads.

6. Screw the fast-pulse capacitor onto the end of the screw adapter; an M4 screw on the other end may be used to attach cables to capacitors.

7. Repeat steps for each high-voltage plate.