



FACULTÉ
DES SCIENCES



UNIVERSITÉ LIBRE DE BRUXELLES

Detection and Mitigation of Propagating Electrical Discharges Within the Gas Electron Multiplier Detectors of the CMS Muon System for the CERN HL-LHC

Thesis submitted by Elizabeth Rose STARLING

in fulfilment of the requirements of the PhD Degree in Physics (“Docteur en Physique”)

Academic year 2020-2021

Supervisor: Professor Gilles DE LENTDECKER

Thesis jury:

Barbara CLERBEAUX (Université libre de Bruxelles, Chair)

Laurent FAVART (Université libre de Bruxelles, Secretary)

Gilles DE LENTDECKER (Université libre de Bruxelles, Promoter)

Archana SHARMA (CERN)

Eraldo OLIVERI (CERN)

Michael TYTGAT (Universiteit Gent)



The author gives permission to make this doctoral dissertation available for consultation and to copy parts of this doctoral dissertation for personal use. In all cases of other use, the copyright terms must be respected, in particular with regard to the obligation to explicitly note the source when quoting results from this doctoral dissertation.

December 12th, 2020

Title: Detection and Mitigation of Propagating Electrical Discharges Within the Gas Electron Multiplier Detectors of the CMS Muon System for the CERN HL-LHC

Author: Elizabeth Rose Starling

Doctor of Science in Physics

Academic year: 2020 – 2021

Abstract

In preparation for the High-Luminosity Large Hadron Collider (HL-LHC) at CERN, the Compact Muon Solenoid (CMS) Detector is undergoing a series of upgrades to its existing infrastructure, and is adding in several completely new subdetector systems. The first of these new systems, called GE1/1, is a series of 144 gas electron multiplier (GEM) detectors, arranged as 36 two-detector “superchambers” in each of the muon endcaps of CMS. These detectors are a subtype of micropattern gas detectors, and consist of three layers of “GEM foils”, thin sheets of polyimide coated with 5 μm of copper on each side and chemically etched with holes of 50 - 70 μm diameter at a pitch of 140 μm . These layers are stacked on top of a printed circuit board (PCB) readout and sealed within a gastight volume that is filled with Ar:CO₂ 70:30, and a high voltage is applied to the foils to create electric fields within the GEM detectors. When a muon enters the detector and ionizes the gas within, the ionized electrons encounter these fields and multiply in Townsend avalanches at each successive foil layer, until they are read out at the readout PCB at a gain of $\sim 10^4$.

In early 2017, a demonstrator system known as the “slice test” was installed into the negative endcap. Consisting of 10 GEM detectors, the two-year-long slice test served as both a proof of concept for the GE1/1 system and an invaluable learning experience that would permanently impact not only the GE1/1 project, but future GEM systems GE2/1 and ME0 as well. During the slice test, it was observed that readout channels were being lost in the course of operation to such a degree that the operational lifetime of the system was in serious jeopardy. These losses were attributed to damage to the front-end readout ASIC (VFAT) inputs, caused by propagating electrical discharges within the detectors, and a dedicated campaign to study the discharges was launched. The results of this study will be presented in this dissertation.

A campaign to mitigate these discharges and their resulting damage was launched. In order to protect the sensitive VFAT from damage, several external protection circuits were proposed and thoroughly tested. The results of these tests, which are presented herein, determined that a series of resistors totaling 470 Ω would be installed on the VFAT hybrid. When coupled with an additional 200 k Ω resistor on the HV filter, this reduced the probability of damage following a discharge from 93% to 3%.

As GE2/1 and ME0 are not due to be installed for another few years, more complex discharge-prevention measures can be put into place. As such, the following measures have been examined, and results will be discussed herein: A new, larger VFAT hybrid is being manufactured, whose larger surface area can accommodate more robust protection circuits than those considered and used for GE1/1. As well, double-segmented GEM foils, in which both the top and bottom of each foil is segmented into $\leq 100 \text{ cm}^2$ sectors that are separated by resistors, were examined for use in the detectors. These double-segmented foils were found to introduce a cross-talk signal in the detectors that results in false signals being treated as true signals, which causes a saturation of the GEM bandwidth and results in unwanted dead time. These cross-talk signals, as well as the compromises made to reduce the cross-talk while maintaining robust discharge prevention, will be discussed.

Keywords: Compact Muon Solenoid (CMS), Gas Electron Multiplier (GEM) Detectors, Electrical Discharges, Readout Electronics

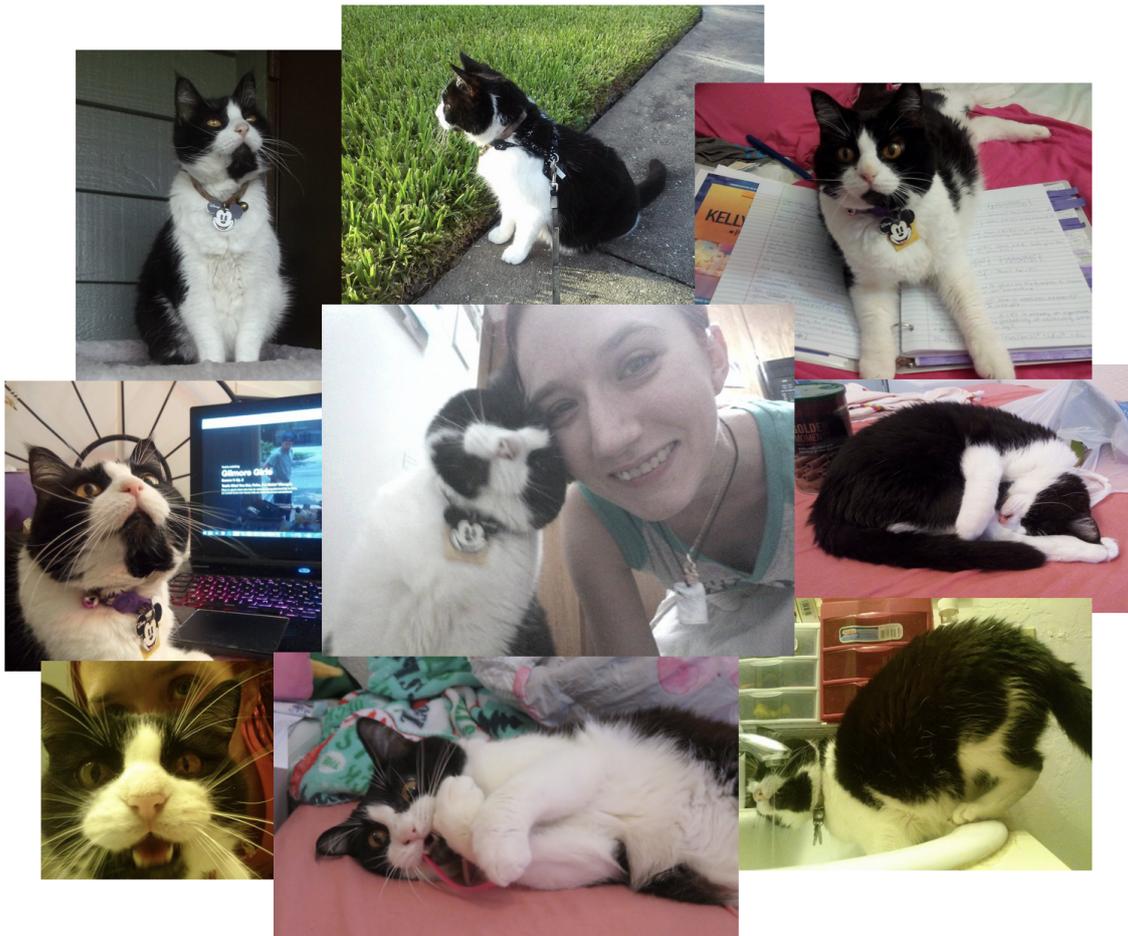
Dedication

This dissertation is dedicated to Dinah. You are my heart, my love, my life.

I wouldn't be where I am today without you. You were a friend when I had none, family when I didn't know what family meant. You gave me reason to live when I had nothing left, you gave me someone to always come back to. You gave me hope, you gave me love, you gave me a reason to keep getting up after every fall. Every step of the way, every leap of faith and every downfall, you were by my side, helping me take the next step forward. You gave me the courage I needed to fly across the world and start this new journey. This dissertation, this person, wouldn't exist without you.

You are my everything, now and forever. And when we see each other again, I promise to tell you everything you made possible just by being you.

I love you. With all my heart, always.



Acknowledgements

First and foremost, I would like to thank my supervisor, Dr. Gilles De Lentdecker, for taking a chance on an American girl he had never met before, who was living on the other side of the ocean. Who believed in my potential so much that he was willing to fly me to Belgium for a weekend on the off-chance I would succeed in getting a highly prestigious grant not normally given to foreigners. I hope one day to have that kind of blind optimism, and to help someone else succeed like he helped me.

I would also like to thank my unofficial adoptive supervisor, Dr. Archana Sharma, who was the steady rock of support I needed throughout my PhD. She believed in me on a personal and professional level more than anybody else, and was always willing to listen and offer assistance, no matter how busy her schedule. She has encouraged me to take every opportunity, to branch out into science communication and work on my strengths beyond just physics. She, more than anyone else, made me feel truly needed in the GEM community and the CERN community at large. I could not have gotten through this PhD without her, and for that I will be forever grateful. I hold her as a prime example for how to lead with kindness and caring.

It would be impossible to get through these acknowledgements without mentioning a different sort of adopter - Felicia S. Vreeland, my mom in all senses but the biological. Her kindness is so vast, she saw a near-stranger on Facebook and adopted her into her family as though she was her own flesh and blood without question. It has filled my heart with a joy I never knew, to be able to share my accomplishments with her and have her be proud of me, to have her tell everyone how smart and successful her daughter is. Her pride in me helps instill a pride in myself I never had before, and her love fills me with a love I never had before. I can truly say, as I have never been able to say before I met her, "I love you, mom".

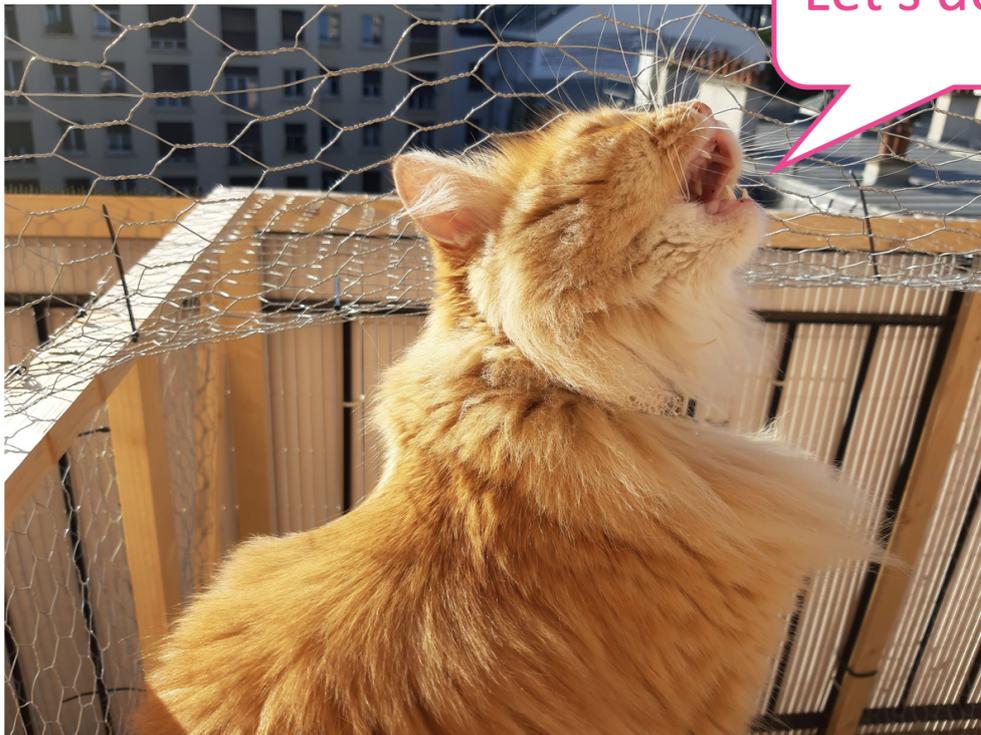
And for a love of another kind, there is of course Matt Edwards. The past three and a half years have been the things dreams are made of. I will never forget the adventures we've gone on together, how many new and incredible places we've seen. From the streets of London to the sights of Paris, to the ancient temples of Athens and fairytale castles in Munich, we've made happy memories all across Europe. We've donned costumes and run through the backlots of Disneyland Paris and the streets of Old Geneva. When I decided, on a whim, to walk across Liechtenstein, you were right there beside me every step of the way. And thanks to your knowledge and planning, I finally got to live out my dream of going to Japan, spending an amazing 30th birthday at Tokyo Disneyland and DisneySea, feeding the deer at Nara, exploring Kyoto, even getting to experience my first earthquake! And in between all of the adventures has been bliss of another kind - coming home to you every day, knowing I can always rely on you for comfort, for support, for a trip to Bois de la Batie when I need a cheer-up or a game of Pandemic when I need a distraction. You've cheered me on and kept me going and have truly brightened up my life, and I am so grateful to have you.

Experimental high energy particle physics is all about collaboration, so I'd like to thank all of the members of the GEM collaboration who have helped me along the way - Andrey Marinov, Anna Stamerra, Cameron Bravo, Federica Simone, Felipe Ramirez, Francesco Fallavollita, Francesco Licciulli, Giovanni Mocellin, Ilaria Vai, Jared Sturdy,

Jeremie Merlin, Johny Jaramillo, Louis Moreaux, Martina Ressegotti, Misha Dalchenko, YongHo Jeong, Zoltan Szillasi, and anyone else I'm forgetting to mention - I couldn't have done it without you!

Of course, I couldn't have done this without my #1 collaborator - Migros MBudget pomegranate energy drink. To the hundreds of cans who stayed by my side day in and day out, thank you for your service.

Finally, this wouldn't be my dissertation if I didn't take the time to thank my cat, Sophie. You're a fluffy little assbutt with a murder streak a mile wide, but nothing turns a bad day good like coming home and listening to you purr, or playing a crazy game of mousie fetch. Life is certainly never boring with you around.



This research was funded in its entirety by the F.R.S.-FNRS through its FRIA doctoral grant scheme.

Table of Contents

Abstracts	I
Abstract	I
Dedication	II
Table of Contents	VII

Detection and Mitigation of Propagating Electrical Discharges Within the Gas Electron Multiplier Detectors of the CMS Muon System for the CERN HL-LHC	1
0 Introduction	1
1 Particle Detection at the Large Hadron Collider & the Compact Muon Solenoid	2
1.1 The Large Hadron Collider	2
1.2 The Compact Muon Solenoid	4
1.2.1 The Tracker	5
1.2.2 The Electromagnetic Calorimeter	8
1.2.3 The Hadron(ic) Calorimeter	9
1.2.4 The Solenoid Magnet	11
1.2.5 The Muon System	12
1.2.6 The Trigger	17
2 Gas Electron Multiplier Detectors	23
2.1 GEM Foils	23
2.2 Townsend Avalanches	24
2.3 Gain and Affecting Factors	28
2.3.1 Collection and Extraction Efficiency	28
2.3.2 Single-, Double-, and Triple-GEM Detectors	29
2.3.3 Drift Region	30
2.3.4 Induction Region	31
2.3.5 Transfer Region(s)	33
2.3.6 GEM Foil Voltages	35
2.3.7 GEM Foil Parameters	36
2.4 Time Resolution	39
2.5 Application to GE1/1 parameters	40
3 GE1/1 Readout Electronics	44
3.1 Introduction	44
3.2 Front-End Electronics	45
3.2.1 Readout Board	45
3.2.2 GEM Electronics Board (GEB)	47
3.2.3 VFATs	52
3.2.4 Optohybrid (OH)	68
3.2.5 FEASTs	72

3.3	Back-End Electronics	73
3.3.1	μ TCA Crate	73
3.3.2	Calorimeter Trigger Processing (CTP7) Card	74
3.3.3	AMC13 Card	77
3.4	Scan Types	79
3.4.1	S-Curves	79
3.4.2	Latency Scans	87
3.4.3	S-bit Threshold Scans	90
4	GE1/1 and the Slice Test	92
4.1	The Ancestors of GE1/1	92
4.1.1	COMPASS	92
4.1.2	TOTEM	94
4.1.3	LHCb	97
4.1.4	Summarizing History	98
4.2	Motivation for GE1/1	99
4.3	GE1/1 Detector System	102
4.4	The Slice Test	107
4.4.1	Slice Test Operations	112
4.4.2	Slice Test Results	118
5	Discharge Damage within Gaseous Detectors	125
5.1	Theory of Discharges within Gaseous Detectors	125
5.1.1	Slow Breakdowns	126
5.1.2	Fast Breakdowns	126
5.2	Propagating Discharge Mechanism in GEM Detectors	129
6	Sustained Operations Investigations	138
6.1	Introduction	138
6.2	Environmental Effects on Detector Health	140
6.2.1	Beam Luminosity	140
6.2.2	Detector Operational Conditions	141
6.2.3	Magnetic Field	141
6.3	GEMINIm01L1 _{v2}	142
6.4	GEMINIm01L2 _{v2}	145
6.4.1	Channel Loss	145
6.4.2	Detector Operation Periods	146
6.4.3	Potential Damage-Causing Events	146
6.4.4	Beam Luminosity	153
6.4.5	Magnetic Field	154
6.4.6	Environmental Factors	155
6.4.7	Conclusions	156
6.5	Channel Loss in the HV Divider Chambers	156
6.5.1	Time Series S-curves	156
6.5.2	Summary Tables	159
6.5.3	Summary Plot	162
6.5.4	High Voltage Instabilities as Discharge Indicators	163
6.5.5	Beam Luminosity	165
6.5.6	Magnetic Field	167

6.5.7	Quality Control (and a Lack Thereof)	170
6.5.8	Chamber Operation	173
6.6	Plan of Action Moving Forward	178
6.6.1	HV Training	179
7	Discharge Damage Mitigation Studies	182
7.1	Channel Loss Probability	182
7.2	VFAT3 Protection Circuits	182
7.2.1	HV3b_v2	184
7.2.2	HV3b_v3	186
7.2.3	HV3b_v4	192
7.2.4	On-Detector Testing	193
7.3	HV Filter Protection	195
7.4	GE1/1 Final Protection Configuration	196
8	Future CMS GEM Systems: GE2/1 and ME0	198
8.1	GE2/1	198
8.1.1	Motivation	198
8.1.2	Detector Design Specifications	200
8.1.3	Readout Sectors and VFAT3s	202
8.1.4	GE2/1 Optohybrid	204
8.1.5	Back-End Electronics	207
8.1.6	Protection Circuits for the GE2/1 VFAT3 Plug-In Cards	208
8.2	ME0	214
8.2.1	Motivation	214
8.2.2	Detector Design Specifications	215
8.2.3	ME0 DAQ Electronics System Overview	218
8.2.4	ME0 “ASIAGO” Optohybrid	219
8.2.5	ATCA Back-End	222
9	Foil Double Segmentation: a Discharge Mitigation Strategy for GE2/1 and ME0	223
9.1	Discharge Rate in a Double-Segmented GE1/1 Chamber	224
9.2	“Dual-polarity” Cross-Talk Signals	228
9.3	GE1/1 Segmentation Size Studies	233
9.3.1	GE2/1 Studies	240
9.3.2	ME0 Studies	248
9.3.3	“Final” Configuration Choice	255
10	Summary and Conclusions	258
10.1	The CMS GE1/1 System	258
10.2	The Slice Test and Resulting Changes	259
10.3	Future CMS GEM Systems: GE2/1 and ME0	260
10.4	Double Segmentation of Future CMS GEM Detectors	260
10.5	Final Remarks	261
	Bibliography	271

Chapter 0

Introduction

“The current state of knowledge can be summarized thus: In the beginning, there was nothing, which exploded.”

(Sir Terry Pratchett — “Lords and Ladies”)

Since the moment it was first switched on, on September 10th, 2008, the Large Hadron Collider (LHC) at CERN has been an object of fascination for millions around the globe. Designed to be the most powerful circular particle accelerator on Earth, the LHC would be capable of accelerating protons to just $3.1 \frac{m}{s}$ less than the speed of light, and its collisions would be capable of producing everything from the earliest states of matter in the universe to particles that had only previously been theorized but never seen. The discovery of one such particle, the Higgs Boson, was announced to the world on July 4th, 2012, expanding our understanding of the most fundamental laws of our universe and further fueling the interest surrounding the LHC.

The discovery of that elusive boson was made possible in part by the Compact Muon Solenoid (CMS) experiment, one of two general purpose detectors at the LHC. Located in Cessy, France, the CMS detector is a monument to the great accomplishments of collaborative particle physics, the result of years of hard work by scientists and engineers from countless nations. And although the CMS experiment has made great advances in the field, it nevertheless is open to improvement, and it is in the context of these improvements that this dissertation takes place.

Beginning in Run 3 in 2022, the CMS detector will contain a brand-new muon subsystem, called GE1/1. GE1/1 adds an entirely new technology to the CMS detector in the form of gas electron multiplier (GEM) detectors, a type of micropattern gas detectors that were invented in 1997 by Fabio Sauli. These compact gaseous detectors, which are capable of handling the immense particle rates at the most forward regions of the CMS muon endcaps, will greatly expand the capabilities of the CMS experiment with regards to muons.

However, the CMS detector is unlike any other detector on Earth, and the LHC environment is among the harshest on the planet. The conditions created within the CMS detector during LHC collisions simply cannot be replicated in any other laboratory. This necessitated a demonstrator system for GE1/1 called the “slice test” - a series of 10 prototype detectors which would be installed into the CMS endcap in January 2017 and remain there for two years. These detectors would provide the crucial experience of running within the CMS environment, and serve as a proof of concept for the final GE1/1 system.

But as any good scientist knows, Nature always has a trick or two up her sleeves. The slice test produced results unlike any we could have predicted, and the consequences of those unexpected results would fundamentally change how GEM detectors would be built and operated within CMS. It is these fundamental changes, and the results that necessitated them, that we will cover in this dissertation.

Ready? Then let's begin...

Chapter 1

Particle Detection at the Large Hadron Collider & the Compact Muon Solenoid

1.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is a circular particle accelerator located on the border of Switzerland and France. At 8.5 km in diameter, it is the largest circular particle accelerator in the world to date. When constructed, it reused the 12 m-diameter concrete tunnels of the Large Electron-Positron Collider (LEP) after the experiment was shut down in 2000 after 11 years of use, replacing the most powerful lepton accelerator in the world with the most powerful hadron accelerator in the world. These tunnels lie approximately 100 m underground, with the glacial moraine on top of them providing an additional layer of shielding from the vast amounts of radiation that is produced within the accelerator complex. The footprint of the LHC can be seen superimposed atop a map of the area in Figure 1.1.

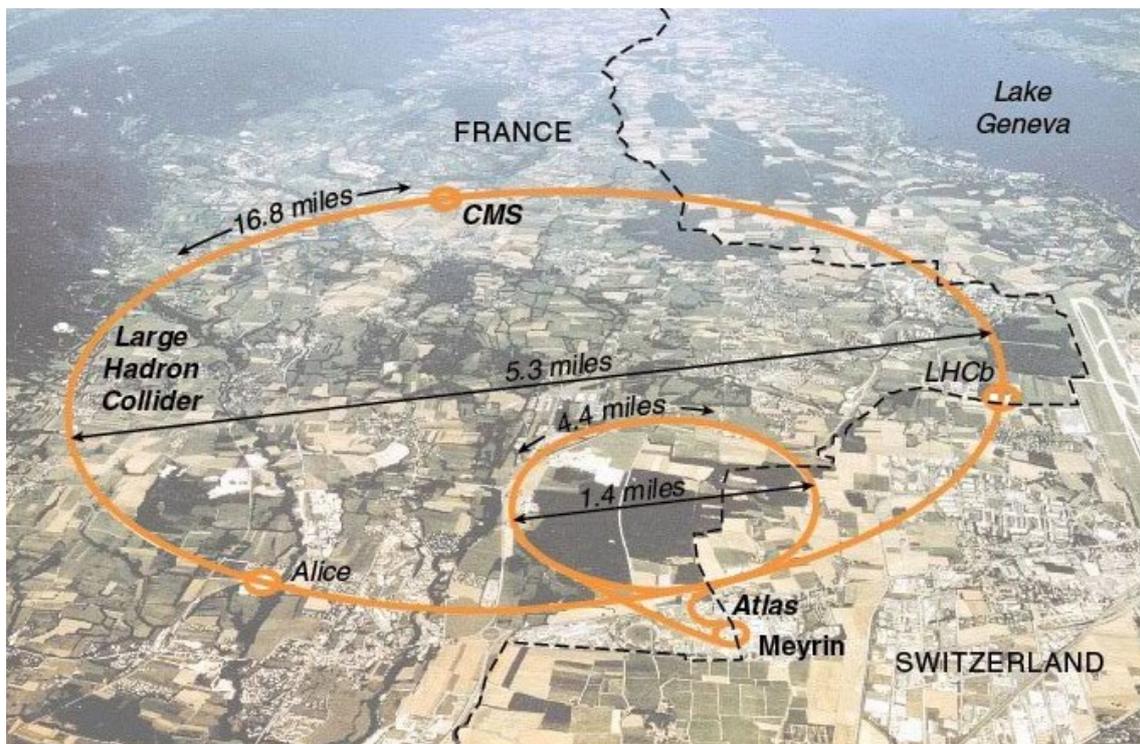


Figure 1.1 – Aerial view of the Large Hadron Collider’s location within Switzerland and France, with major experiments and distances¹ labeled

The LHC is primarily a proton-proton accelerator, though for approximately one month out of the year it accelerates heavy ions, specifically $^{208}_{82}\text{Pb}$. These protons begin as hydro-

¹in miles, because despite having lived in Switzerland for over three years now, the author remains stubbornly American and will fight to the death to preserve her use of miles and degrees Fahrenheit.

gen gas in a small bottle located at the linear accelerator complex LINAC2². This gas is fed into the excellently-named Duoplasmatron, which strips the electrons from the hydrogen atoms and accelerates the protons away, resulting in a 92 keV proton beam travelling at 0.014c. This proton beam is then fed into the LINAC, where it is further accelerated until they reach 50 MeV, travelling at 0.314c.

From here, the protons leave the linear accelerator and enter their first circular accelerator, the Proton Synchrotron Booster, labeled PSB on Figure 1.2. In the PSB, the protons are accelerated to 0.916c, with an energy of 1.4 GeV. Once the beam reaches this operational energy steadily, it is fed into the next largest circular accelerator, the Proton Synchrotron (PS). Here, the beam is further accelerated to 0.9993c, with an energy of 25 GeV.

Although at this point, the beam can be split off to many other experimental areas, as shown in Figure 1.2, the main interest of this dissertation is the LHC itself, so we follow the protons as they are funneled into the Super Proton Synchrotron (SPS), where they reach a speed of 0.999998c and an energy of 450 GeV.

Finally, from the SPS, the proton beam is fed into the LHC in two opposing directions, where it could theoretically (though it has not, as of yet) reach its design energy of 7 TeV and a speed of 0.99999999c, just $3.1 \frac{m}{s}$ slower than the speed of light, c .

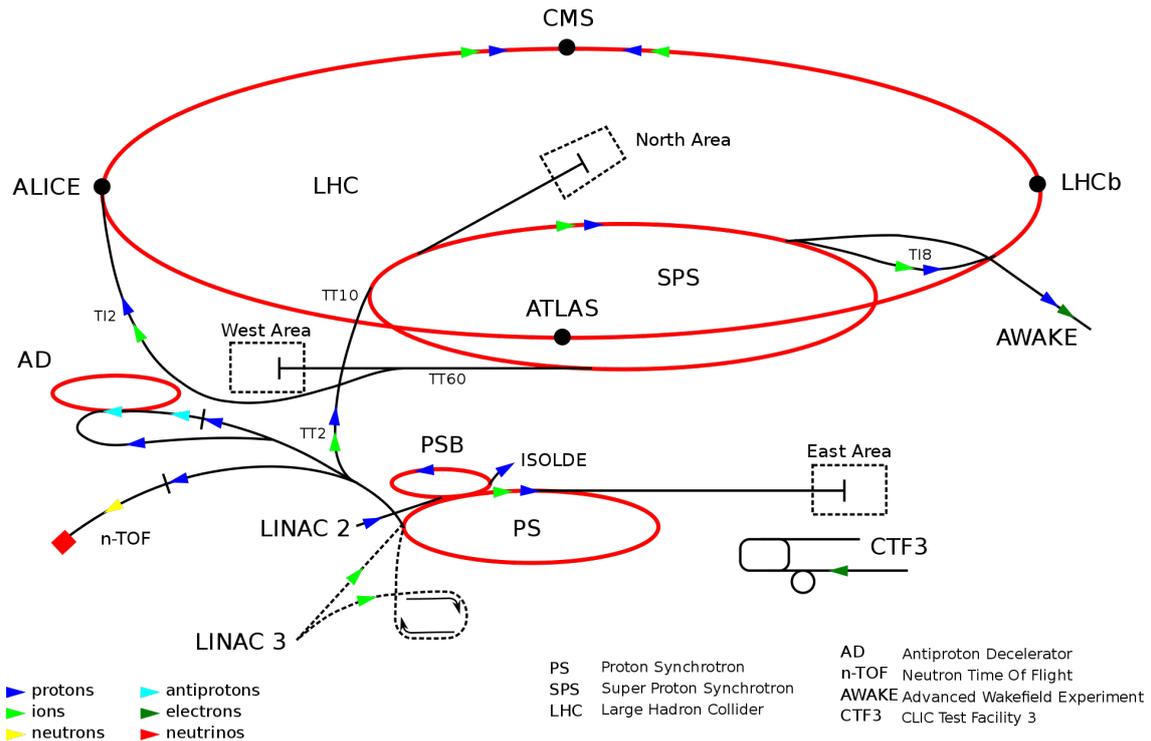


Figure 1.2 – Diagram of the LHC accelerator complex

Once stable beams are reached at this point, the real magic can begin. The two opposing beams, which are steered around the LHC ring by a series of 1,232 11 kA dipole magnets, are crossed at four approximately-equidistant collision points. Located at each one of these points is a major LHC experiment which utilizes the collision data to probe the secrets of our universe. These experiments are:

²Beginning in Run 3, this will be replaced by LINAC4. The author touched the actual bottle at LINAC2 during Open Days 2019 and it was awesome.

- Large Hadron Collider beauty (LHCb), the smallest of the four main experiments both in terms of detector and collaboration size. As a specialized b-physics experiment, LHCb is primarily interested in measuring the parameters of CP violation in the interactions of b-hadrons.
- A Large Ion Collider Experiment (ALICE), on the opposite end of the ring from LHCb. Unlike the other three experiments, ALICE is optimized to study the heavy ion collisions that occur for approximately one month out of the year. ALICE's studies focus primarily on the physics of strongly interacting matter at extreme energy densities, the so-called "quark-gluon plasma" which is theorized to have existed a fraction of a second after the Big Bang, before quarks and gluons were bound together to form hadrons and other heavier particles.
- A Toroidal LHC ApparatuS (ATLAS), one of the two general purpose experiments at the LHC which was involved in the discovery of the Higgs Boson in 2012.
- The Compact Muon Solenoid (CMS), the star of our show, and the other (and, in this author's very biased opinion, superior) of the two general purpose experiments which contributed to the discovery of the Higgs Boson in 2012. It is in this experiment where our story takes place.

1.2 The Compact Muon Solenoid

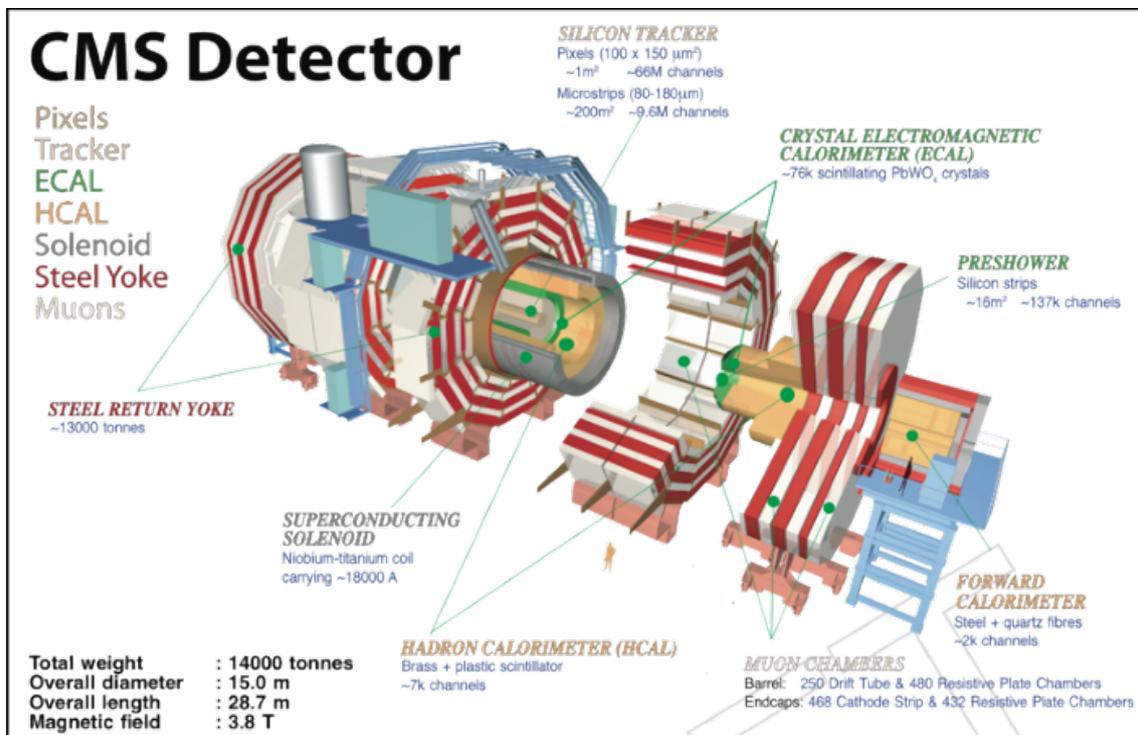


Figure 1.3 – Schematic view of the Compact Muon Solenoid experiment

Despite its name, the CMS experiment is still impressively large; at 15 m in diameter and 28.7 m in length, it weighs approximately 14,000 tonnes. Perhaps a more fitting name would be the Complex Muon Solenoid, as the CMS experiment is composed of a multitude of different layers of detector systems which work together in pursuit of the experiment's

goals. These layers, shown in Figure 1.3, are (very briefly) described in the following sections.

1.2.1 The Tracker

The tracker is the innermost detector system of CMS, being the closest to the interaction point, the point in the center of CMS where the proton-proton collisions take place. The main purpose of the tracker system is to record the paths, or tracks, of charged particles as they pass through. This is crucial to the workings of CMS, as tracking the paths of the particles as they are bent by the strong magnetic field of the solenoid magnet allows for the calculation of the particle's momentum, which is a critical measurement used in physics analysis. The tracker is designed to be extremely lightweight so as to disturb the particle as little as possible and therefore not interfere with the physics interactions it helps to measure, with each position measurement accurate to $10\ \mu\text{m}$ such that particle tracks can be reliably reconstructed by only a few measurement points. It is also designed to be extremely radiation resistant (or “rad-hard”), as, being the closest detector system to the interaction point, it receives the highest volume of particles and therefore the greatest amount of radiation. The decision was therefore made to build the tracker entirely out of silicon. The tracker is, itself, composed of two subdetector systems, which are described in the following sub-subsections.

1.2.1.1 Pixels

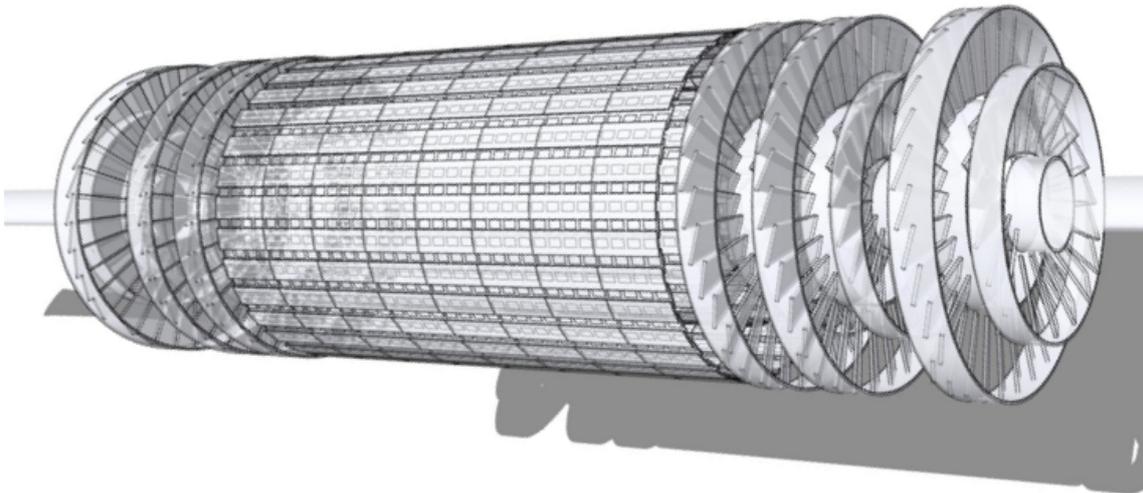


Figure 1.4 – Diagram showing the barrel (BPIX) and endcap (FPIX) sections of the pixel tracker detector as they were following the Phase I upgrades [1]

The pixel subdetector is truly the innermost detector system, with the closest of the four barrel (BPIX) layers being just 30 mm away from the beam pipe (and the other three layers being just 68, 109, and 160 mm away, respectively) following the Phase I upgrade - previously there had been only three barrel layers, with the closest layer 40 mm from the beam pipe, as shown in Figure 1.5. This proximity makes the pixel subdetector a valuable asset in measuring the shortest-lived of particles.

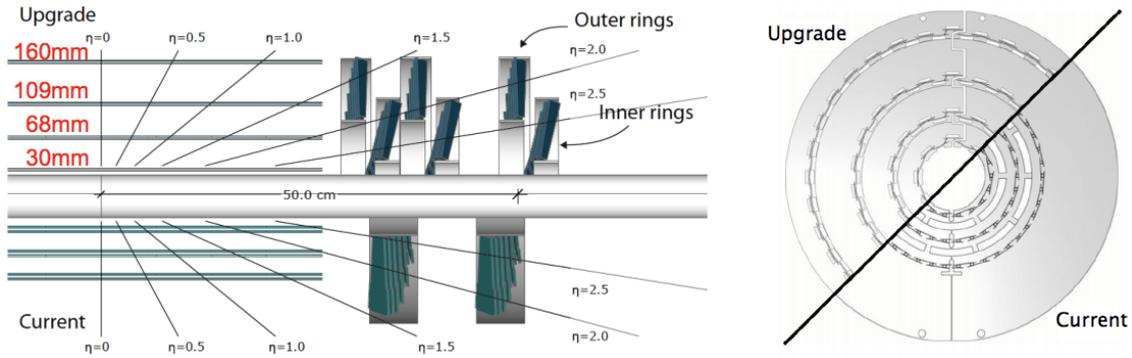


Figure 1.5 – Longitudinal (left) and transverse (right) views of the layout of the Phase I upgrade of the PIXEL detector in comparison to its original design layout [2]

In addition to these four barrel layers, the pixel detector contains six endcap disk (FPIX) layers, three on each end of the barrel section. Each of these disks is composed of an inner and outer ring, as seen in Figure 1.6.

Each layer is split into $100 \mu\text{m} \times 150 \mu\text{m}$ segments (the eponymous pixels), with 1184 in the barrel layers and 672 in the endcap disks, for a total of 124 million channels. When a charged particle passes through one of these pixels, electrons are ejected from the silicon atoms within them, creating electron-hole pairs. An electric current is then used to collect these charge pairs on the surface of the pixel in the form of an electrical signal. An electronic readout chip is connected to the silicon via bump-bonding, using a spot of solder in the center of each pixel. This chip then amplifying the electrical signal, and allows us to know where a particle has passed through. As a result of the two-dimensional pixels and the existence of multiple layers, a three-dimensional image of the particle's path through the detector can be reconstructed.

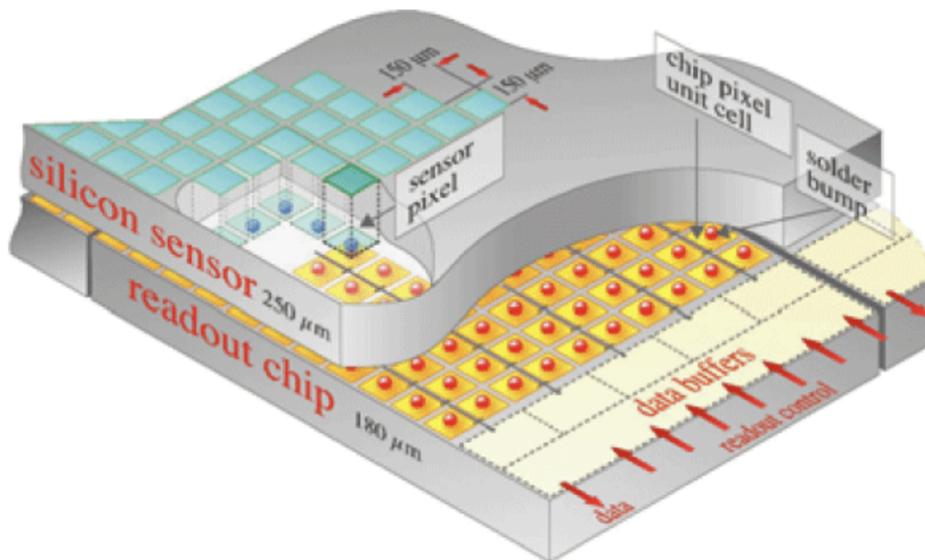


Figure 1.6 – Cross-section of the CMS pixel subdetector, showing the various components [3]

1.2.1.2 Strips

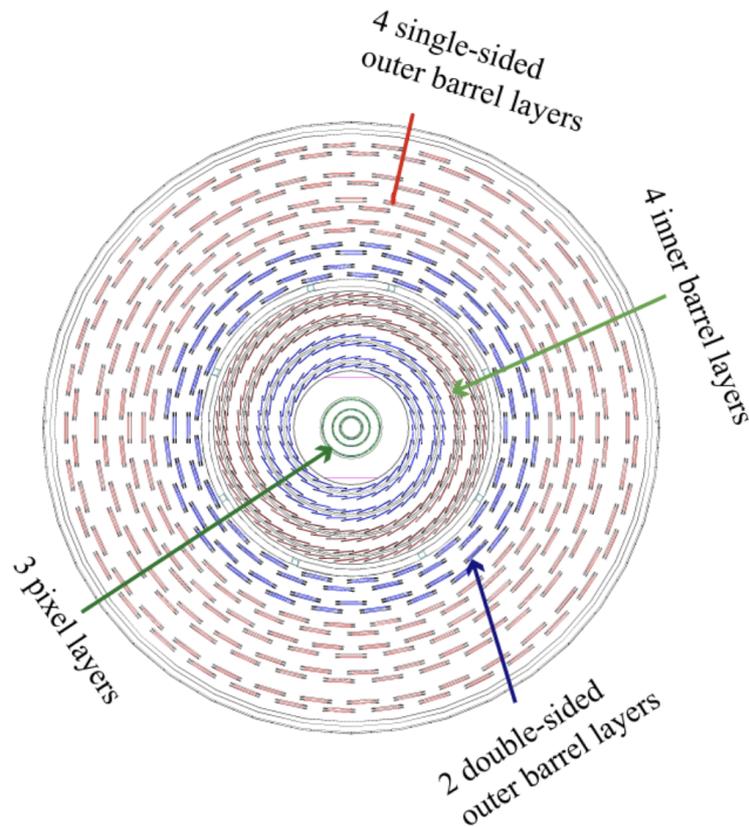


Figure 1.7 – Cross-section of the CMS strips subdetector, showing the various layers as they appear perpendicular to the beam line [3]

Beyond the innermost pixels lie ten silicon strip layers, reaching a total radius of 1.3 m. These layers are subdivided into four inner barrel (TIB) layers arranged in shells, and six outer barrel (TOB) layers, the inner two of which are double-sided and the outer four of which are single-sided. The inner barrel layers are closed off with two inner endcaps (TID) which are each composed of three small discs, and the subdetector as a whole is closed off by two outer endcaps (TEC). These different layers can be seen in Figure 1.7.

The strips subdetector contains 15,200 highly-sensitive modules which consists of three main elements: a set of strips, the readout electronics for those strips, and the module's mechanical support structure. In total, this adds up to 10 million strips read out by 80,000 microelectronic chips.

Much like the pixels, when a particle passes through one of the silicon strip modules, electrons are knocked out of the silicon atoms within the strip modules, and their movement within the applied electric field of the detector creates a small electrical current on the order of a few ns. In the case of the strips, the current is amplified by an Analogue Pipeline Voltage³ (APV25) chip. Four or six APV25 chips are housed within a hybrid which also contains the electronics required to measure key sensor information such as temperature, as well as to provide the timing information necessary to match hits with the collisions from

³Having worked with GEM detectors since 2011, The author is very familiar with APV chips, the predecessor front-end readout chips to the VFATs we use today. It is not until this moment, 14h15 on May 23rd, 2020, that she finally knew what APV stood for. It was an ongoing mystery in her previous research laboratory.

which they originated. The APV25 stores each signal in memory for several microseconds before processing them and sending them to a laser where they are converted into infrared pulses. These pulses are then transmitted over one of forty thousand 100 m fiber optic cables to the radiation-free environment of the CMS service cavern (USC55), also known as the “counting house”, for analysis.

1.2.2 The Electromagnetic Calorimeter

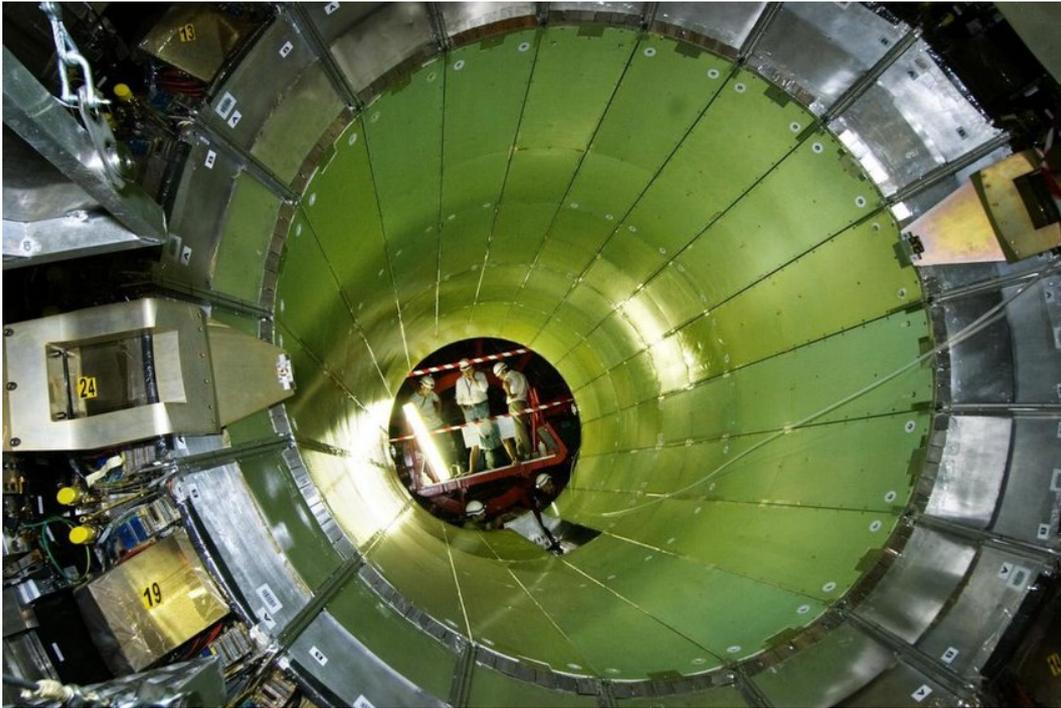


Figure 1.8 – Photo of the barrel section of the CMS electromagnetic calorimeter [4]

The next subdetector out from the beamline is the electromagnetic calorimeter (ECAL), which measures the energies of photons and electrons produced in the collisions of the LHC. ECAL is subdivided into a barrel section and two endcaps, the former of which can be seen in Figure 1.8.

In order to measure the energy of the photons and electrons to a high degree of precision despite the harsh magnetic field and radiation conditions of this area of the CMS detector, ECAL was primarily constructed out of lead tungstate (PbWO_4) crystals. These transparent metallic crystals scintillate when photons and electrons collide with their heavy nuclei, producing fast, short, well-defined bursts of blue light in proportion to the particle’s energy. These light bursts are detected by photodetectors at the back of each crystal, which then convert them into electrical signal that is amplified and transmitted to the electronics for analysis.

1.2.2.1 ECAL Barrel

The cylindrical barrel of ECAL is composed of 61,200 lead tungstate crystals which are arranged as 36 supermodules, or groupings of 1,700 crystals in glass fiber or carbon fiber pockets. Each barrel crystal measures $2.2 \text{ cm} \times 2.2 \text{ cm} \times 23 \text{ cm}$ and weighs 0.9 kg. Each crystal is outfitted with an avalanche photodiode (APD), a type of silicon semiconductor

photodetector. A strong electric field is applied to each APD, such that when a photon strikes the silicon and ionizes an atom, the ionized electron is accelerated within the electric field and further ionizes other silicon atoms, creating an electron avalanche. This produces a high current from a relatively low amount of light per incident particle in the crystals.

1.2.2.2 ECAL Endcaps

The two endcaps of ECAL are composed of 14,648 additional lead tungstate crystals, also arranged as supermodules. These crystals are larger than their barrel counterparts, measuring $3\text{ cm} \times 3\text{ cm} \times 22\text{ cm}$ and weighing 1.6 kg. Each crystal is outfitted with a vacuum phototriode (VPT) rather than an APD, as the radiation in the endcaps is far higher than in the barrel and exceeds the limits for silicon photodiode use. As the name suggests, VPTs contain three electrodes within a vacuum. When photons strike the atoms in the first electrode, called the cathode, electrons are released. A voltage difference is maintained between the three electrodes, which accelerates these electrons into the second electrode, called the anode. This produces additional electrons which are accelerated into the third electrode, called the dynode, in the same way. This process also produces a high current from a relatively low amount of light per incident particle in the crystals.

1.2.2.3 Preshower Detectors

While one of ECAL's main jobs is to detect high-energy photons, this can at times be difficult. Neutral pions, which can decay into two lower energy photons, can mimic the desired high-energy photons if these two lower energy pions are closely spaced and picked up together by ECAL's crystals. This is a particular issue in the endcaps, where the angle between the two photons from the neutral pion decay is particularly small. To prevent this misidentification, a preshower detector is in place in front of each endcap.

The preshower detector is composed of two lead planes, each followed by silicon sensors similar to those used in the tracker. The lead planes create an electromagnetic shower when a photon passes through them. This shower contains electron positron pairs which are detectable by the silicon sensors. A single shower allows for the measurement of the photon's energy, whereas the two layers of lead followed by silicon detectors allow for a measurement of the photon's position as well.

Each silicon sensor measures approximately $6.3\text{ cm} \times 6.3\text{ cm} \times 0.3\text{ mm}$ and is divided into thirty-two 2 mm strips, which are arranged in a grid pattern to form a rough circle in front of the ECAL endcap detector. These 2 mm strips provide a much finer granularity compared to the 3 cm which aids in distinguishing the two individual pion decay photons.

1.2.3 The Hadron(ic) Calorimeter

Whereas the electromagnetic calorimeter measures the energy of electromagnetic particles (photons and electrons), the hadronic calorimeter (HCAL) measures the energy of hadrons, as the name suggests. It also provides indirect measurement of the presence of weakly interacting neutral particles such as neutrinos.

HCAL is split into barrel (HB), outer barrel (HO), endcap (HE), and forward (HF) sections. The barrel is split into 36 wedges, each weighing 26,000 kg. These HB wedges form the outermost detector within the solenoid magnet, and can be seen in part during installation into CMS in 2007 in Figure 1.9. The outer barrel, on the other hand, sits just

outside of the solenoid's coil, as a way of ensuring that no energy is able to leak undetected out of the HB. The endcaps, like the barrel, are also split into 36 wedges.

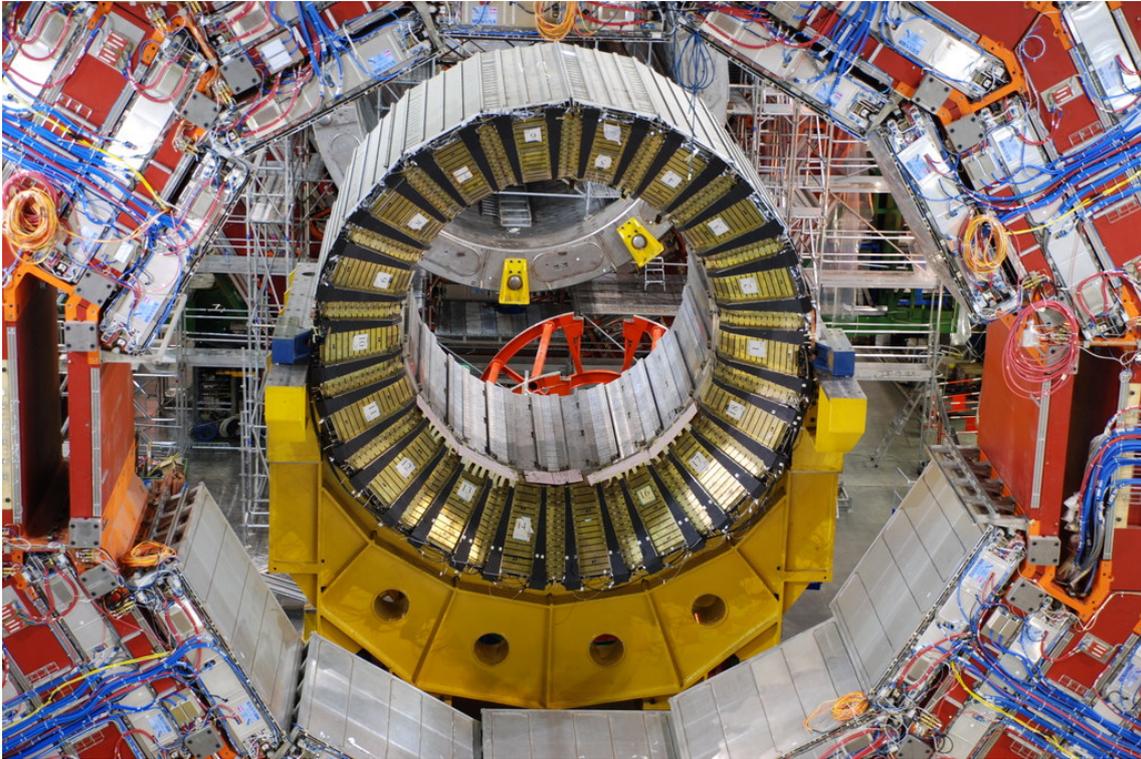


Figure 1.9 – Photo of a section of the barrel portion of HCAL being installed into CMS in 2007 [5]

HCAL is a sampling calorimeter, in that it is made of alternating layers of dense absorber material (brass⁴) and tiles of fluorescent plastic scintillator material. When a hadron hits an absorber plate, it interacts and creates a burst of secondary particles, which then reach the next layer of absorber material and create a further burst of secondary particles, and so on in a shower of particles. The shower is large, requiring approximately one meter of absorber material to contain them, limiting the minimum size of the calorimeter.

As the shower develops, the particles within pass through the scintillator tiles, which emit a blue-violet light in response. Each tile contains wavelength-shifting optical fibers, each with a diameter of less than 1 mm. These optical fibers absorb the scintillation light and shift its wavelength into the green region of the visible spectrum. This green light is then transmitted through clear optical cables to readout boxes which are strategically situated at different points within HCAL. Prior to any manipulation, this light is optically totaled over a stack of tiles, called a tower, resulting in a measurement of the original particle's energy and aiding in the particle's identification.

The readout boxes contain photodetectors (hybrid photodiodes, or HPDs) which then amplify the received signal by approximately 2,000 times. The HPD accelerates electrons across a narrow (few mm) gap onto a silicone diode target which is divided into 19 pixels. These pixels are each capable of generating their own amplified electronic signal when

⁴The majority of this brass originates from Russian artillery shells, which had been stored by the Russian navy since WWII. The author has always been particularly fond of this fact and frequently shares it with visitors she guides on CMS tours, as it perfectly embodies the mission of CERN to do science solely for peaceful reasons.

struck, which allows for the detection and amplification of up to 19 individual calorimetry signals with a single HPD - thus requiring only 420 HPDs for the entirety of CMS.

For each collision, the signals are digitized using charge integration and encode (QIE) chips before being sent to the trigger system and DAQ system for analysis.

1.2.4 The Solenoid Magnet

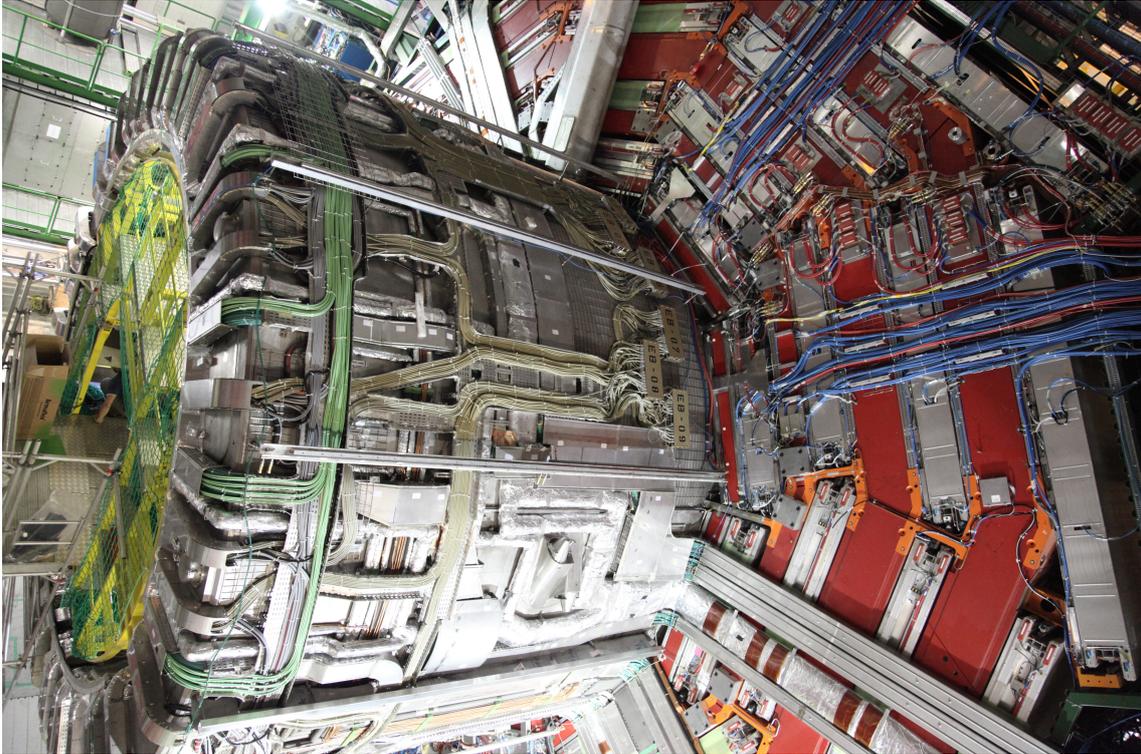


Figure 1.10 – Photo of a portion of the CMS solenoid magnet taken from above during Long Shutdown I [6]

The superconducting solenoid magnet of CMS, though not a detection system itself, is perhaps one of the most crucial components of the CMS detector - a fact which is reflected in its position within the name *Compact Muon Solenoid*. This superconducting solenoid coil, seen in Figure 1.10, is six meters in diameter and 13 meters long, and weighs over 10 kilotonnes. It stores 2.5 GJ of energy, which is often given in terms of being enough energy to melt 18 tonnes of gold, and at peak operation produces a field of 4 T, making it the most powerful solenoid magnet in the world. In order to allow enough energy to flow to create this powerful electric field, the magnet is cooled to 4.65 K, or -268.5°C . In comparison, outer space is just 2.7 K, or -270.45°C .

The purpose of the solenoid magnet is to bend the path of particles created in each LHC collision, in order to aid in their identification during physics analysis. Neutral particles have straight paths within the detector, as they are not affected by the magnetic field of the solenoid, whereas charged particles' paths are curved (with positively-charged particles curving in the opposite direction as negatively-charged particles). Furthermore, the more momentum a given particle has, the less its path is curved by the presence of the magnetic field, so the shape of the curve can be a very useful tool in momentum measurements. The stronger the magnetic field, the greater the degree of bending, so the 4 T value was chosen so as to, in combination with the high-precision measurements of the tracker and muon

detectors, allow for the accurate measurement of the momentum of even the highest-energy particles.

1.2.5 The Muon System

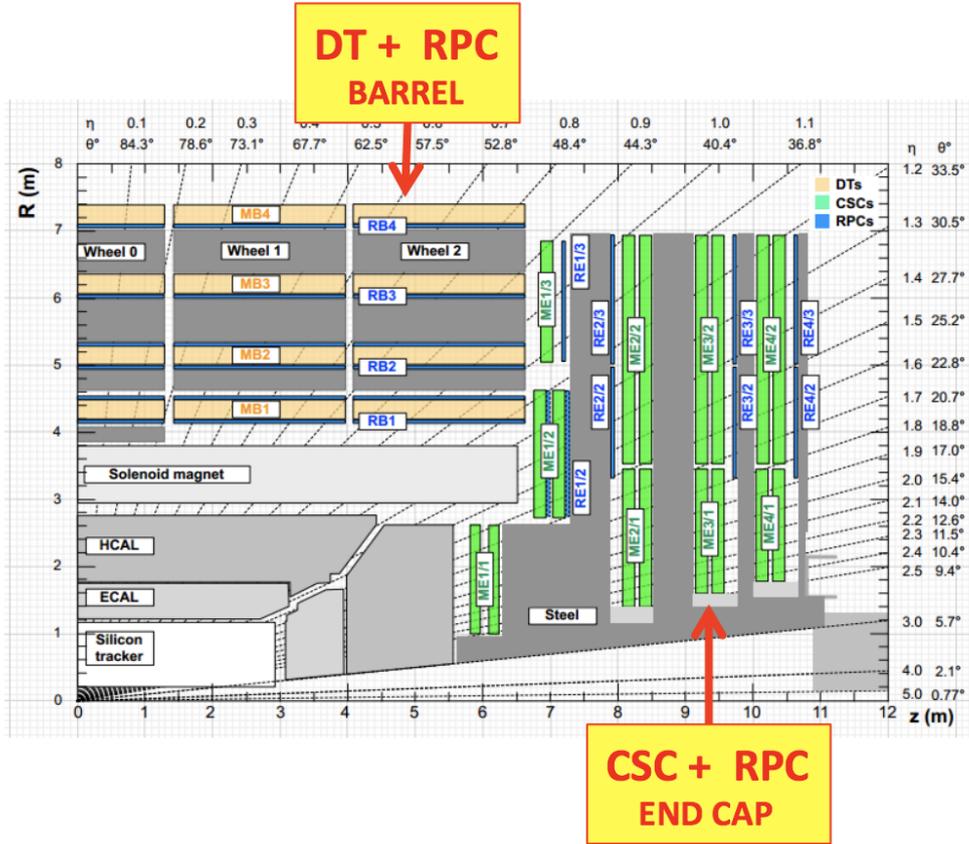


Figure 1.11 – Cross-section of CMS, detailing the subdetector locations [7]

The muon system also has great importance within the CMS detector, as reflected in its position within the name *Compact Muon Solenoid*. It also holds particular importance to the author, as it is the system in which this dissertation protect takes place.

The CMS muon system is composed of two distinct sections, the barrel and the endcap, which feature different subdetectors, as seen in Figure 1.11.

The top and rightmost y-axes of Figure 1.11 are given both in terms of the angle, θ , where $\theta = 90^\circ$ is perpendicular to the beam axis and $\theta = 0^\circ$ is parallel to the beam axis, and the pseudorapidity, η , a commonly-used spatial coordinate which describes the angle of a particle relative to the beam axis, as defined in Equation 1.1.

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right) \quad (1.1)$$

Precision measurements and Level 1 (L1) triggering (the latter of which is discussed in greater detail in Section 1.2.6) are provided in the barrel by five wheels of drift tubes (DT), which cover acceptances up to $|\eta| < 1.2$, and in the two endcaps by cathode strip chambers (CSCs), which cover $1.0 < |\eta| < 2.4$. In both the barrel and endcap regions, resistive plate chambers (RPC) provide redundant trigger and coarse position measurements. Further information on these muon subdetector systems can be found in Table 1.1 and in their

respective subsections.

Muon Subsystem	DTs	CSCs	RPCs
$ \eta $ range	0.0 - 1.2	0.9 - 2.4	0.0 - 1.9
Number of chambers	250	540	Barrel: 480 Endcap: 576
Number of layers per chamber	r- ϕ : 8 z:4	6	RB1 and RB2: 2 Elsewhere: 1
Surface area of all layers	18,000 m ²	7,000 m ²	Barrel: 2,300 m ² Endcap: 900 m ²
Number of channels	172,000	Strips: 266,112 Anodes: 210,816	Barrel: 68,136 Endcap: 55,296
Spatial resolution	100 μ m	50 - 140 μ m	0.8 - 1.3 cm
Time resolution	2 ns	3 ns	1.5 ns (intrinsic)
Fraction of working channels	98.4%	99.0%	98.3%

Table 1.1 – Properties of the CMS Muon System in 2016 [8]

1.2.5.1 Drift Tubes

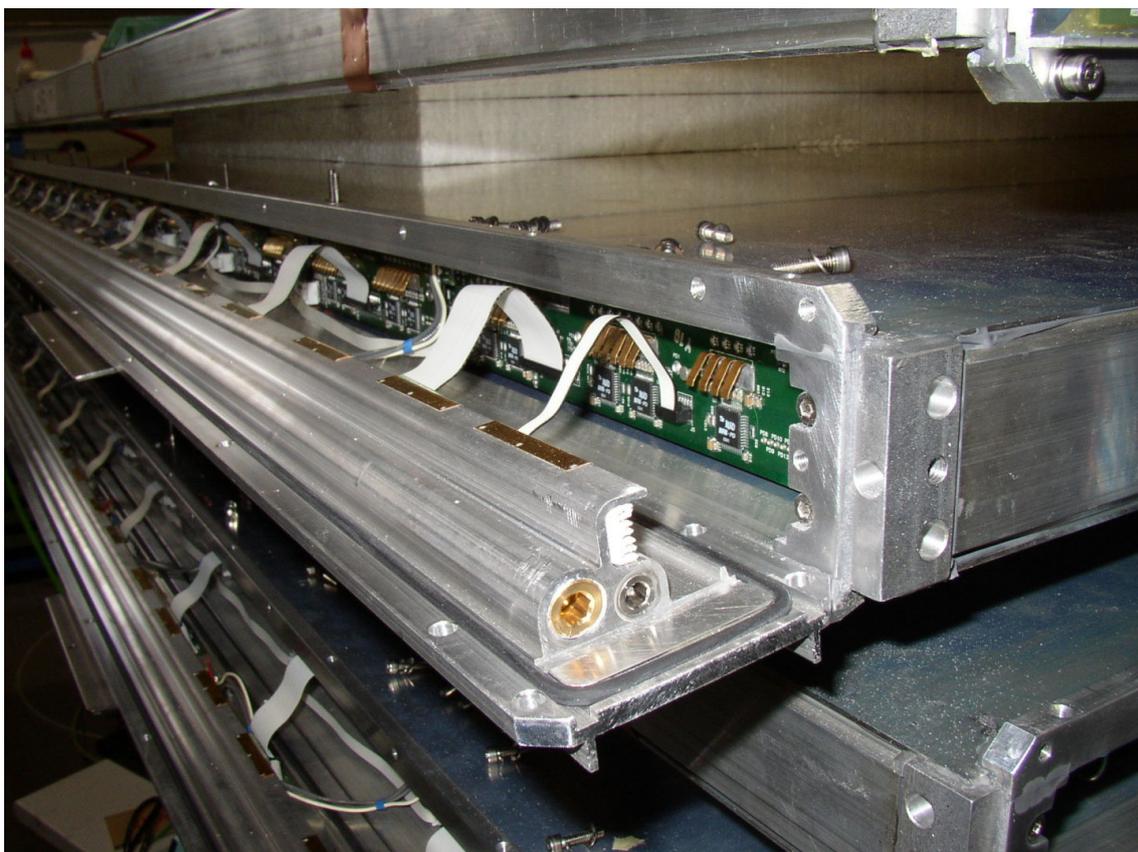


Figure 1.12 – A photo of a DT chamber during the production process [9]

CMS contains a total of 250 DT chambers, one of which can be seen in the production process in Figure 1.12. Each DT chamber is, on average, 2 m \times 2.5 m in size, and consists of 12 aluminum layers arranged in three groups of four. Each group contains up to 60

drift tubes, 4 cm wide tubes which contain a stretch wire sealed inside of a gas volume. In the middle of these three groups, the wires are orthogonal to the beamline and measure the z -position along the beam. The wires of the two outer groups, on the other hand, are parallel to the beamline and provide a track measurement in the magnetic bending plane (r, Φ).

The DT chambers are installed in the five wheels of the return yoke of the CMS magnet, named YB-2, YB-1, YB0, YB+1 and YB+2. Each wheel is further divided into 12 sectors which each cover $\sim 30^\circ$ around the interaction point. Each sector is organized into four stations, named MB1, MB2, MB3, and MB4 from inside to outside. Interestingly, the MB4 chambers do not contain the central group of drift tubes, and thus only make measurements in the Φ coordinate, not the z coordinate [9].

When a muon passes through the gas volume of the drift tubes, it ionizes the gas atoms within. These ionized electrons then follow the electric field towards the positively-charged stretched wire at the center of the tube. These electrons are then measured in two ways in order to obtain two coordinates for the muon's position when it encountered the detector: the first is measured by registering the position along the wire where the electrons hit, and the second is measured by calculating the muon's original distance from the wire by multiplying the speed of an electron in the tube by the time taken to travel in it.

1.2.5.2 Resistive Plate Chambers

The RPCs of CMS use a Freon-based gas mixture which contains $C_2H_2F_4$ and SF_6 , both of which are greenhouse gases with a very high global warming potential at 1,430 and 23,900 respectively. And despite the horrific effects these gases have on the planet, the RPC detectors within CMS have a significant amount of leaks. As a result, they release the equivalent of 6 tonnes of CO_2 per *hour*. In comparison, the US per-capita carbon emissions, which is one of the highest in the world, is 16.1 tonnes per *year*. Although they have a plan for reducing their emissions, this plan is to reduce their emissions by 70% by Run 5 in 2033, and they will continue to release the equivalent of 1.9 tonnes of CO_2 per hour during Run 3. This is literally all you need to know about the CMS RPCs. [10] [11] [12] [13]

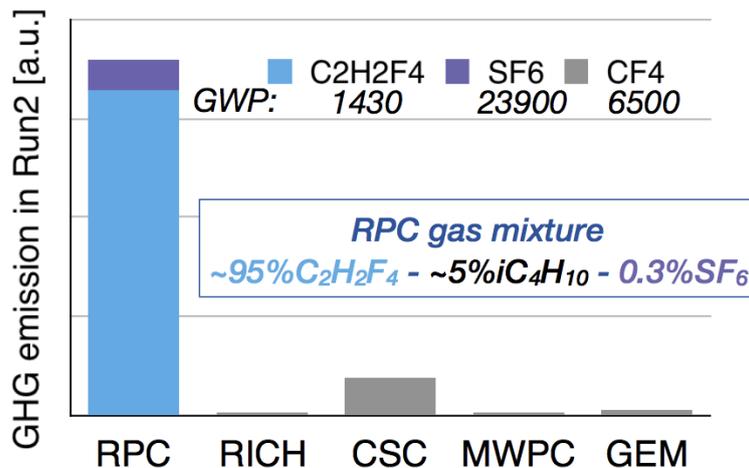


Figure 1.13 – Bonus figure showing the relative greenhouse gas emissions for the various muon systems of CMS, ATLAS, and ALICE, with the y-axis in arbitrary units. Note that the y-axis is in arbitrary units, and these statistics do not include CMS GEMs as they do not use any of the three greenhouse gases cited [12]

Alas, sometimes we must learn about things we don't like in order to broaden the reaches of our knowledge, or, in this case, to provide a fair comparison between the various CMS muon systems. So let us briefly describe the technical aspects of these mega-polluters.

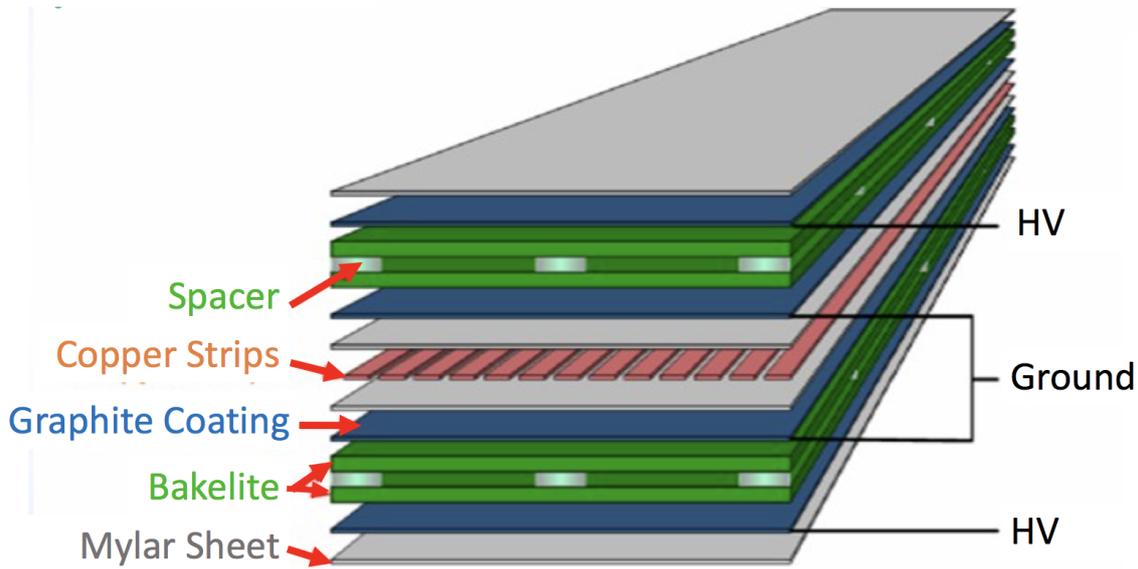


Figure 1.14 – Diagram of the layers of a CMS RPC

Figure 1.14 shows the layers of a CMS RPC, a double-gap RPC. Here, we see that, as the name might suggest, there are two gas-filled gaps, which sandwich a common strip-based readout layer between them. On either side of the gas gaps is an insulating Bakelite (polyoxybenzylmethylenglycolanhydride) panel, which is covered in a conductive layer of graphite. During detector operation, a high voltage is applied to this graphite layer.

RPCs have an excellent time resolution, $\mathcal{O}(1ns)$, the CMS RPCs in particular utilize their unique layout in order to operate each gap at a lower HV setpoint and thus reduce the readout time required, thereby improving the time resolution.

However, in addition to their heavy usage of potent greenhouse gases and their propensity to leak heavily and frequently, the RPCs come with a distinct disadvantage. That is, their basic working principle is that of a spark chamber. As a charge particle travels through the detector, it ionizes the gas within the gaps. A trigger system is then used to apply high voltage to those plates, creating an electric field immediately after the particle has passed through, producing sparks which follow the trajectory of the particle [14]. So while the detector has a low time resolution for a given particle, the presence of the sparks results in significant deadtime between events. This, above all else, is what disqualified RPC technology from being installed into the slots which would later be filled by gas electron multiplier (GEM) detectors in the project now known as GE1/1, as this deadtime made them unsuitable for handling the high particle rates found there.

Despite this disadvantage, RPCs are located in both the barrel and endcap regions of CMS. The barrel region contains 480 chambers, which vary size and are located in the range of $|\eta| < 0.8$ and $|z| < 7$ m. These detectors are split between five wheels, which are subdivided into 12 azimuthal sectors and equipped with six radial layers of RPCs. These layers are referred to as RB1, RB2, RB3 and RB4, where RB1 and RB2 consist of two layers each. As a result, RB1 and RB2 allow for RPCs to be placed on either side of a DT chamber, whereas RB3 and RB4 have an RPC mounted only in front of each DT.

The endcap regions, on the other hand, cover the range of $|\eta| < 1.6$. The detectors are spread out over eight disks (four each on the positive and negative sides of CMS) which are divided into two rings, called R2 and R3 [15]. Here, they are interspersed with CSCs rather than DTs.

1.2.5.3 Cathode Strip Chambers

The CSC system is composed of rings of trapezoidal CSC chambers mounted on eight disks - four in each endcap - for a total of 450 chambers. On each of these disks is two to three rings of chambers, which are designated by $ME\pm S/R$, where ME stands for “muon endcap”, the \pm sign indicates which endcap they reside in, S indicated the disk or station, and R is the ring number. The chambers in the outer rings (ME_x/2 or ME_x/3) are considerably larger than the chambers in the rings which are closer to the beamline (ME_x/1). The ME1/x station can be seen in Figure 1.15.

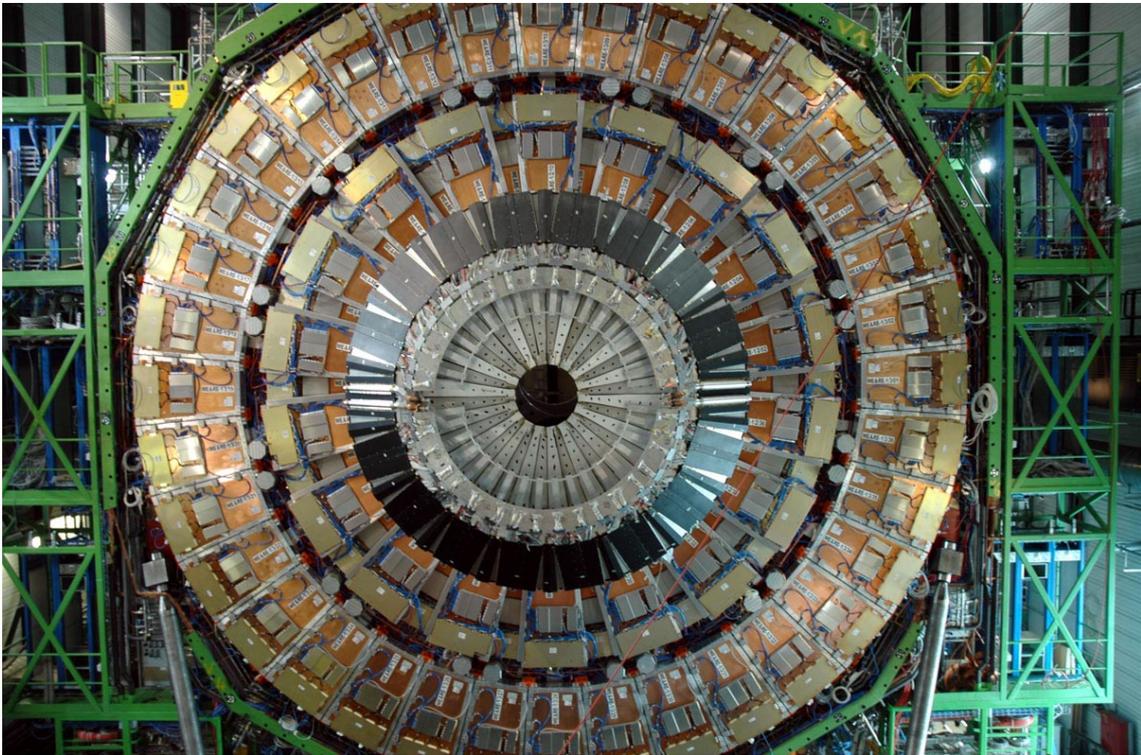


Figure 1.15 – The complete ME1 station as it was in 2008 [16]

Despite their size differences, each chamber contains six detection layers sealed within a gas volume (filled with 50% Argon, 40% CO₂, and 10% CF₄), which are each composed of a positively-charged anode wire plane which is stretched between two negatively-charged copper cathode planes. One of these planes is continuous, whereas the other is segmented into strips in order to provide a measurement of the particles’ positions. The anode planes are separated by a gas gap of width of 2.54 cm, with the exception of the forward-most set of detectors, ME1/1, where this distance is 2.2 cm. With ME1/1 as the exception due to its placement within the CMS detector, the strips of adjacent layers within a chamber are staggered for better imaging.

When muons pass through the CSC chambers, they ionize the gas within. The ionized electrons then flow to the anode wires, creating an electron avalanche. The positive ions, on the other hand, flow away from the anode wires and towards the copper cathode, inducing

a charge pulse in the strips there. Because the strips and the wires are perpendicular to each other, this provides two position coordinates for each particle.

The wires are read out in groups, the width of which varies between different chambers from 1.5 cm to 5 cm. These wire groups are supplied with HV in accordance to the size of the chamber, with the largest chambers containing five HV segments. The strips, on the other hand, are read out individually, with their widths varying between 5 and 12 mm between chambers [17]. This resolution, along with the presence of six detection layers per chamber, make the CSCs highly capable in accurately identifying muons and matching their tracks to those found in the tracker.

1.2.5.4 GE1/1: Gas Electron Multiplier Detectors

Finally we come to the star of our show, GE1/1: the first of the CMS Phase II upgrades, the first time an entirely new technology has been introduced into CMS, and the first time GEM detectors have been used in CMS. And, like any star, it demands an appropriate introduction. So we will turn, in a moment, to Chapter 2 for a more in-depth look at the technology driving CMS' newest subdetector system.

1.2.6 The Trigger

Before we do, however, let us take a moment to discuss an important feature of CMS in which the CMS GEMs will play a key role in the future: the trigger.

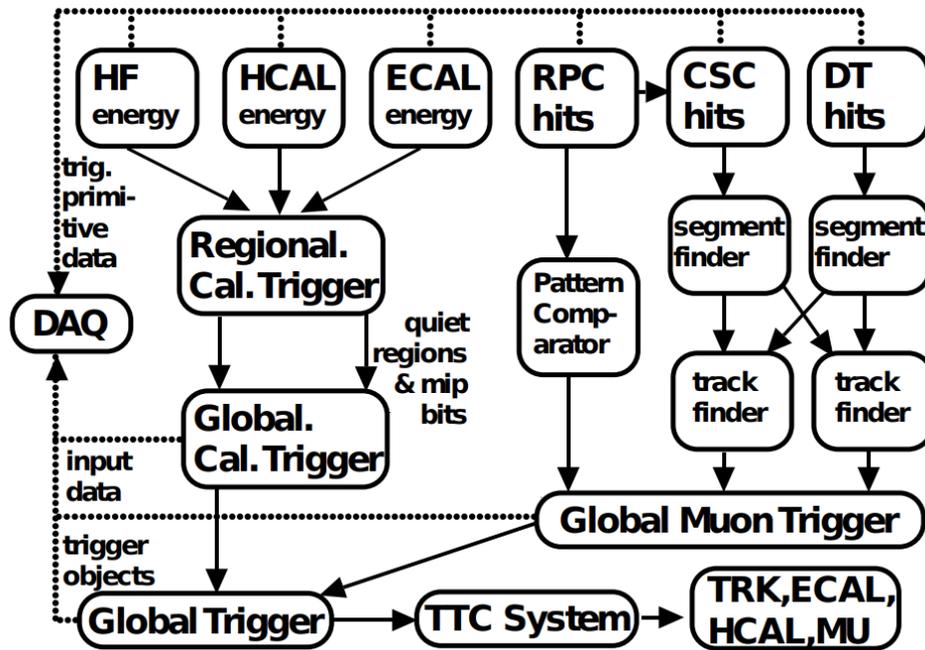


Figure 1.16 – Block diagram of the CMS L1 trigger system, prior to the installation of the GE1/1 system [18]

Each bunch crossing within the CMS detector produces approximately 1 MB of data. And with collisions in the CMS detector occurring every 25 ns, this amounts to approximately 40 TBps of data - a number which, while incredibly large, will only increase with the coming High Luminosity Large Hadron Collider (HL-LHC) upgrade. As the computing resources of CMS are finite in terms of both money and storage space, real-time data selec-

tion is necessary in order to filter out bad or uninteresting events from good or interesting events, such that only the latter will be saved and the former can be permanently deleted. This is where the trigger comes into play.

The trigger itself is broken into two steps: the Level 1 (L1) trigger and the High-Level trigger (HLT), where the L1 trigger is what we are referring to when we discuss real-time data selection.

The L1 trigger system has a fixed latency, where it must decide if an event should be accepted or rejected within $3.2 \mu\text{s}$ of a given collision. It does so using information from the calorimeter and muon detector systems, as shown in Figure 1.16. Here, we see that data from the calorimeters (ECAL as well as the forward and barrel HCAL) are processed first into the regional calorimeter trigger (RCT) and then into the global calorimeter trigger (GCT). On the muon side of things, energy deposits from the RPCs are processed via a pattern comparator, while hits from the CSCs and DTs are processed via a system of segment-finders and track-finders. These are then sent onwards to the global muon trigger (GMT).

The trigger objects received from the GCT and GMT are then synchronized to both each other and to the LHC clock, before being sent via the crate backplane to the global trigger logic (GTL) module. It is here that the trigger algorithm calculations are performed (with more detailed information on these algorithms and the versatility of the CMS L1 trigger found in Reference [18]) and a final trigger decision is made to accept or reject an event.

If the trigger decision is made to accept an event, a Level-1 accept (L1A) signal is sent via the trigger, timing, and control (TTC) system to the tracker, calorimeter, and muon detector systems. This signal triggers a readout of the entire CMS detector system, such that the data can be forwarded to the HLT system.

While the L1 trigger is designed to operate in near-realtime, the HLT system does not need to operate so quickly, and as such it can operate in a similar way to offline processing. For each event, all charged particle trajectories (including electrons, muons, and jets) are reconstructed, and all of the energy depositions in the calorimeters are measured with a fine granularity. The tracks are then associated with their respective energy depositions, and as multiple p-p collisions can occur in each bunch crossing, they must also be associated with the correct collision. Identification criteria are applied, and from this information, events which may be interesting for further data analysis are selected by the HLT. While the L1 trigger must make its decision within $4 \mu\text{s}$, the processing time for the HLT is on the order of tens of ms, dependent on the event in question.

1.2.6.1 The Muon Trigger

Of course, Figure 1.16 is an outdated one, as it is missing one very crucial component of the muon trigger: the CMS GEMs! Nevertheless, it is important to understand the muon trigger system in greater depth, so as to better understand the role that the CMS GEMs will play in it.

1.2.6.1.1 The RPC Trigger We begin with the RPCs, as, unlike the CSCs and DTs, the main purpose of the RPCs is to serve as dedicated trigger detectors rather than dual-purpose tracking and trigger detectors. Thus, the trigger can be built from information which comes directly from the RPCs rather than first building a muon track within them, and there is no local processing other than synchronization and cluster reduction. To do

this, the RPCs use the pattern comparator trigger (PACT) algorithm. The signals coming from the RPC chambers are compared to a set of pre-defined muon track patterns, shown in grey in Figure 1.17. This fit pattern then gives the transverse momentum (p_T) and sign of the muon. In the case of the barrel RPCs, the required coincidence to determine the correct pattern is 6/6, 5/6, 4/6, or, in the case of the inner barrel layers, 3/4. In the endcap, this coincidence must be either 4/4 or 3/4 layers. Of course, a higher coincidence is ideal, with the level of coincidence defining the reconstruction quality of the muon. More details on how this algorithm works can be found in Reference [19].

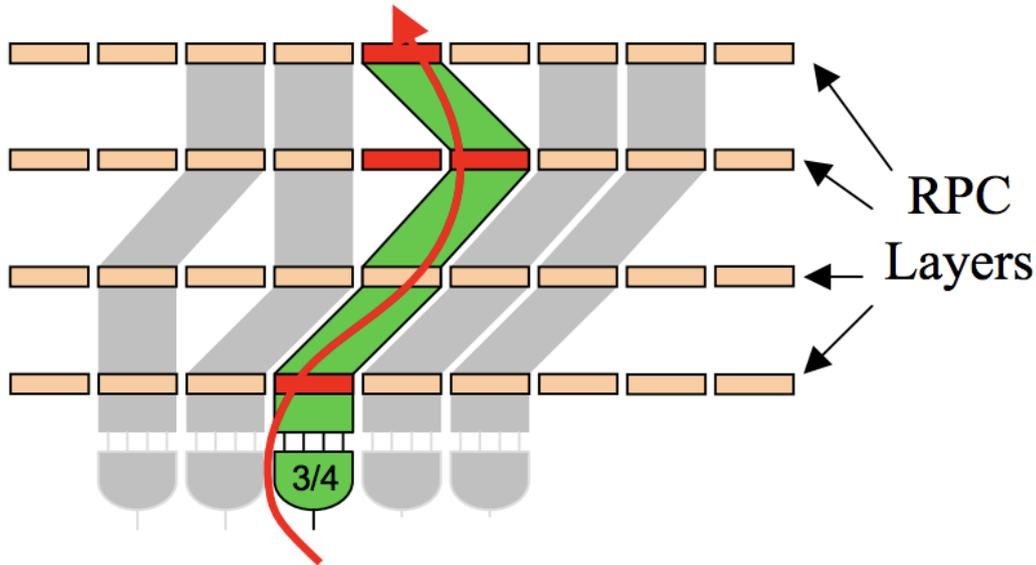


Figure 1.17 – Figure illustrating the PAC algorithm choosing the most-likely muon track from a set of pre-defined patterns [19]

This algorithm is applied to each muon hit in the RPC system, before sending the data to the RPC muon sorter. This sorter then selects the four muons with the highest p_T in the barrel, and the four muons with the highest p_T in the endcaps, and sends this information to the GMT [20].

1.2.6.1.2 The DT Trigger Next, we move to the DTs. Unlike the RPCs, the DTs first create a local trigger (DTLT) prior to forming their portion of the global muon trigger, which contains information about the radial position of the muon, the bending angle, and the reconstruction quality as determined by the number of DT layers which were used to build a given muon track segment. These DTLT track segments are reconstructed by on-detector electronics, shown in Figure 1.18. Each DT is equipped with a mini crate hosting these readout and trigger electronics, and a single chamber can generate up to two DTLT in the transverse plane per bunch crossing.

The creation of a DTLT begins with the Bunch and Track Identifier (BTI) application-specific integrated circuit (ASIC). There is one such BTI for every four wires within the DTs, for a total of approximately 55,000 BTIs for the entire subsystem. The BTI considers uses any three hits to define a straight line and a time stamp. If the four layers for that BTI agree, this is a high quality track reconstruction. Otherwise, if only three layers agree, it is a low quality reconstruction. From each chamber, the two best muon track segments are sent to the Track Correlator (TRACO), which then attempts to correlate

these segments within the inner and outer superlayers within a given azimuthal (Φ) sector and, if successful, outputs a single track segment for the two superlayers. This track segment is then forwarded to the trigger sector collector (TSC) system for that ϕ sector. Here, the segments are synchronized and transmitted over 6 Gbps optical links to USC55, where the data is deserialized and sent to the DT track finder (DTTF) system [21].

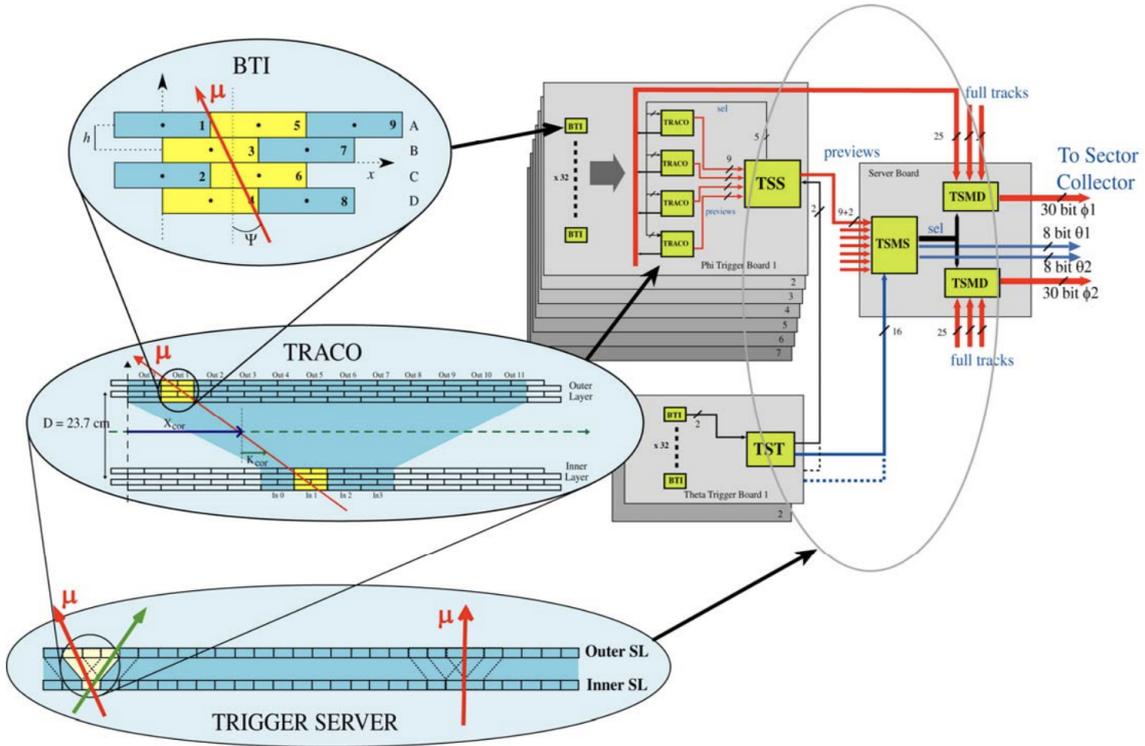


Figure 1.18 – Schematic of the steps and electronics used to form the DTLT [22]

The DTTF processes the DTLT information in order to reconstruct muon tracks which are measured in several stations and then assigns a p_T value to each track. It does so using a sector processor (SP), which is capable of finding up to two muon tracks per sector. It first uses the position and bending angle of each DTLT track to compute the expected position at the outer stations (or, in the case of the fourth station layer, the expected position at the third station). This is then compared to the position of the actual DTLTs, and either accepted or rejected via a programmable tolerance window. If accepted, the difference in the Φ positions of the two inner segments is translated into a p_T value. The Φ and θ coordinates of the track are also computed, and the quality of the track is noted based upon the position and number of stations used to generate the track [23].

In order to take into account tracks which might cross detectors in multiple sectors, each SP also receives a subset of DTLT information from their neighboring SPs via parallel electrical connections. As well, the SPs in the external wheels also receive track segments from the CSC trigger, which we will discuss in Section 1.2.6.1.3.

Then, a wedge sorter (WS) collects up to 12 of these tracks from the 6 SPs of a given wedge, where a wedge consists of five DT sectors at the same Φ position. From these, it selects the two best tracks based on the reconstruction quality and p_T . These two tracks from the 12 total WSs are forwarded to the final sorting module, the barrel sorter (BS). The BS chooses from among these the four final muon candidates, which are delivered to the GMT [18].

1.2.6.1.3 The CSC Trigger The overall flow of the CSC trigger system is shown in Figure 1.19.

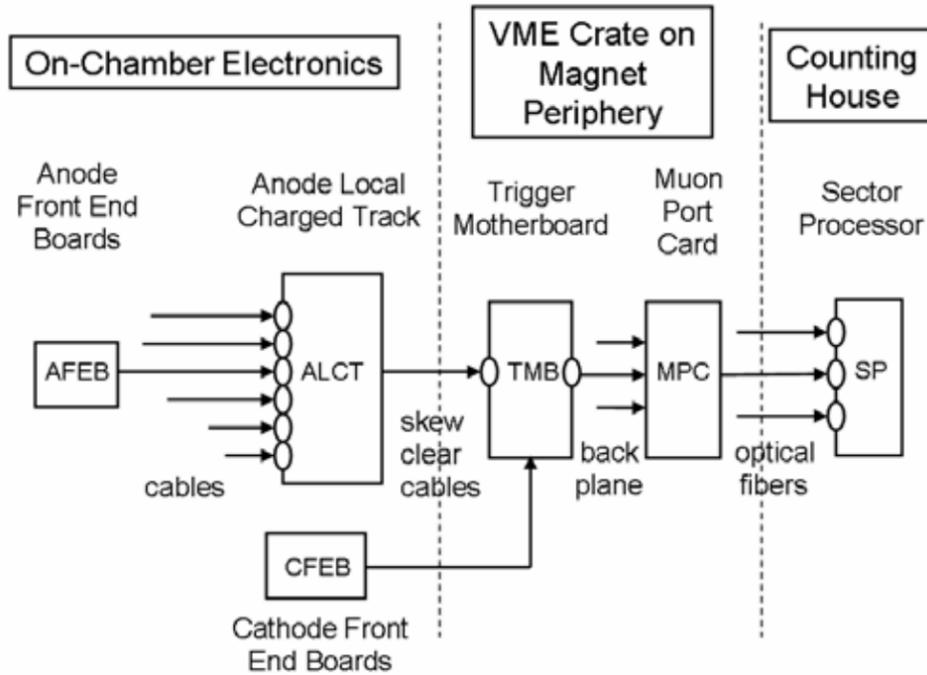


Figure 1.19 – Schematic of the steps and electronics used to form the muon trigger from the CSCs [24]

As with the RPCs, the trigger primitives for the CSCs, called local charged track (LCT) segments, are determined through a comparison of the data received from the detectors and a set of pre-defined patterns that are expected of a muon candidate traversing at least four layers. In order to reduce the number of noise triggers, and because the full developmental time of a cathode signal is on the order of 75 ns, the hits are compared to the patterns twice - before and after a “drift delay” [24].

The anode LCT (ALCT) is formed from the wire hits by a match to an envelope which points towards the interaction region. The cathode LCT (CLCT), on the other hand, is formed from the strip hits by a match to a set of patterns whose bending gives a coarse measurement of the muon momentum early on in the trigger path. In order to specify their local position, both the CLCT and the ALCT are tagged by their strip and wire groups [24]. When both a CLCT and ALCT exist in a single chamber, they are correlated within the CSC trigger motherboard (TMB) into the LCT. In the LCT, the CLCT provides information on the position of the track segment in Φ , while the ALCT provides information on the radial distance of the track segment from the beamline in addition to precise timing information. Each detector can send up to two LCTs per bunch crossing [18].

Up to 18 LCT segments from nine detectors are then collected by a muon port card (MPC), which is located in the same VERSAmodule Eurocard (VME) crate as the TMBs. The best three of these 18 are then transmitted via optical fiber to the CSC track finder (CSCTF) which, like the DTTF, is broken into multiple SPs which each correspond to 60° in Φ . The CSCTF then makes comparisons of track segments within different detector stations, testing for their compatibility in Φ and η with a muon coming from the collision vertex. For a good track, the p_T of the muon candidate is then calculated and recorded. Each SP forwards up to three muon tracks to a sorter card, which makes the final decision

on the four best muons and forwards these four to the GMT [18].

1.2.6.1.4 The Global Muon Trigger The GMT can now fulfill the following functions: synchronizing incoming regional muon candidates from the RPC trigger system, DTTF, and CSCTF (and, beginning in Run 3, the GEM trigger system); merging or deleting duplicate candidates; performing an optimization of the p_T assignment for merged candidates; sorting muon candidates; assigning a quality to the candidates; and storing information about the candidates in the event data.

In the GMT, the muon candidates from the various regional triggers are assigned a “sort rank” according to their p_T values (adjusted to a common scale), and matched geometrically according to their η and Φ values, with some degree of tolerance built in to account for the differences in resolutions between the systems. The matching candidates from the DT + barrel RPCs, and from the CSCs + endcap RPCs, are then merged, and each measured parameter (η , charge, and sort rank from the DTs or CSCs, as well as Φ , p_T , etc...) is merged as well. The candidate can then be assigned a three-bit quality value. In the event the regional triggers cannot be merged, this quality value will be lower, and depends on the quality as defined by its respective regional trigger system and the η value of the regional trigger.

From this information, the GMT sorts each merged (or unmerged) trigger according to its sort rank. This sorting process occurring independently in the barrel and the endcap regions, with the four highest ranking candidates from each being forwarded to a final sorter. From these eight candidates, the four highest ranking are forwarded to the global trigger (GT), where they are joined with the trigger information from the calorimeters, as shown in Figure 1.16.

With this information in mind, we can now turn to Chapter 2 to get our first look at gas electron multiplier detectors.

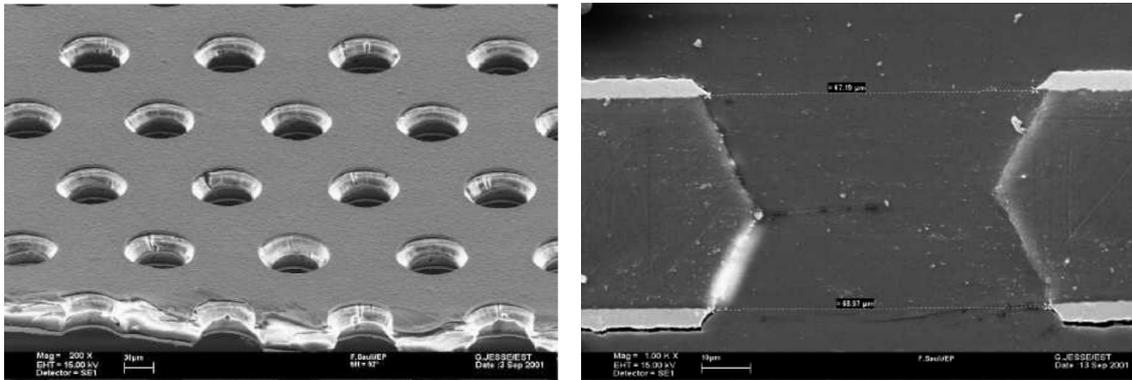
Chapter 2

Gas Electron Multiplier Detectors

The focus of this dissertation, gas electron multiplier detectors, are a species of compact, gaseous particle detectors, belonging to a class of particle detectors known as micropattern gas detectors (MPGDs). They were first created in 1997 by physicist Fabio Sauli within the Gas Detector Development Group at CERN [25], and have since been used in particle physics experiments around the globe. At CERN, GEMs have found use in experiments such as the COMPASS tracker, TOTEM, the inner region of the M1 muon station of LHCb, the upgrade of the ALICE time projection chamber (TPC), and the upgrade of the CMS endcap muon system, the latter upgrade being the subject of this dissertation. They have also been used in smaller laboratories for muon tomography, for the imaging of shielded nuclear materials for defense purposes, and have been studied for use in the medical field.

2.1 GEM Foils

The primary feature of a GEM detector is the presence of GEM foils, 50 μm -thick polyimide sheets that are coated with 5 μm of copper on each side and etched with tapered holes using a technique of masked photolithography. These holes have an external diameter of 70 μm , an internal diameter of 50 μm , and are spaced with a pitch of 140 μm . A microscopic view of one such GEM foil can be seen in Figure 2.1.



(a) Top-down view of a GEM foil

(b) Cross-section of a biconical GEM foil hole

Figure 2.1 – Microscopic views of GEM foils

It should be noted that these measurements and images are specific to a standard GEM foil, and that other variants such as thick GEMs (THGEMs) or GEMs with differing hole diameters and pitches, do exist, and will be discussed in brief in later sections of this dissertation. However, the main focus of this dissertation is on detectors which uses the standard GEM foil described above.

Several manufacturing techniques have been developed for the etching of GEM foils, including double- and single-masked photolithography, plasma etching, and laser drilling. However, it is thanks to the refinement of a single-mask photolithography technique that makes it now possible to manufacture large-area GEM foils, up to and above a square

meter in area. The stages of this technique are shown in Figure 2.2, with the first step on the top of the figure and the final step on the bottom.

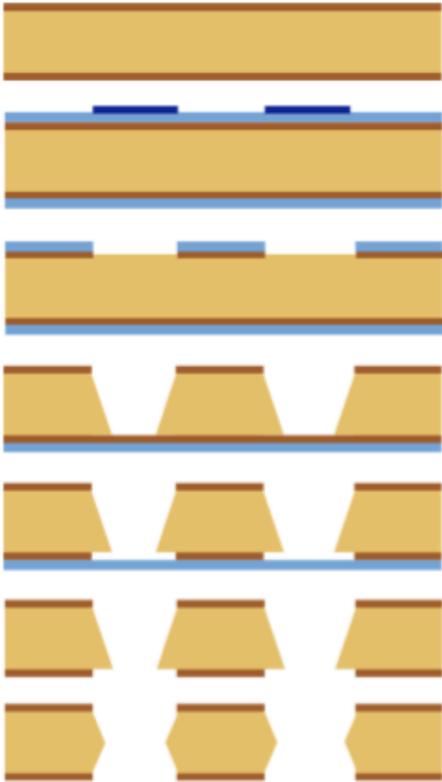


Figure 2.2 – Single-mask photolithography technique for manufacturing biconical-holed foils [26]

First, we begin with a bare polyimide sheet, and sputter $5\ \mu\text{m}$ of copper onto each side via vacuum deposition.

Second, a layer of photosensitive resin is applied on top of the copper cladding, and a reusable lithographic mask is placed on the top layer of photosensitive resin. This mask acts as a stencil for the placement of the holes, as the resin is exposed to ultraviolet light through transparent mask, and its photoresistive properties result in a protective layer everywhere except for the opaque spots where the holes will be present.

Third, the unexposed resin and the copper cladding beneath it are removed via the process of electro-etching.

Fourth, the top layer of photoresist is removed, and the polyimide below the holes in the copper is chemically etched away with a solution of ethanol and ethylenediamine [27], resulting in a wide conical hole. The process is often referred to as “wet etching” due to the fact that this portion of the process is achieved within a large liquid acid bath, as opposed to “dry etching” which would involve plasmas rather than liquid solvents.

Fifth, the bottom layer of copper is etched away in the same diameter as the top layer of copper.

Sixth, the bottom layer of photoresist is removed.

Finally, the remaining lip of polyimide is chemically etched away, transforming the hole geometry from conical to biconical. This hole geometry plays an important role in the gain of the detector, as will be discussed in Section 2.3.

Although the process of single-mask photolithography is a more expensive one, it also comes with many benefits. Notably, in double-mask photolithography, in which both photoresistive layers are treated with UV light simultaneously and the holes are simultaneously etched from both sides of the foil, great care must be taken to ensure the alignment of the masks such that the holes will be etched correctly. Previously, this onerous task of manual alignment made it essentially impossible to produce GEM foils of an area greater than $40\ \text{cm}^2$, thus limiting the possible size of the corresponding GEM detectors.

2.2 Townsend Avalanches

The GEM foils of the GEM detector are sealed within a gas-tight volume and flooded with a gas mixture. In the case of the CMS GEMs, this gas mixture is composed of $\text{Ar}:\text{CO}_2$ 70:30, in order to provide the ideal balance between signal gain (via the Ar) and spark quenching (via the CO_2). Other non-standard mixtures may also be used, including different ratios of Ar and CO_2 , as the author studied at INFN LNF in 2014, or mixtures which also include

other gases such as the Ar:CO₂:CF₄ 45:15:40 mixture used in the GEM detectors of the LHCb experiment. In the latter case, great care must be taken in the construction of the gas recirculation system, as CF₄ is a strong greenhouse gas and therefore must not be released into the environment. This also applies to other such gases which may be included in a given mixture, such as SF₆, another previously-considered gas.

A high voltage is then applied to the foils, generally a potential difference of 300-500 V between the two copper surfaces, which creates an electric field of up to $100 \frac{kV}{cm}$ through the holes, as seen in Figure 2.3.

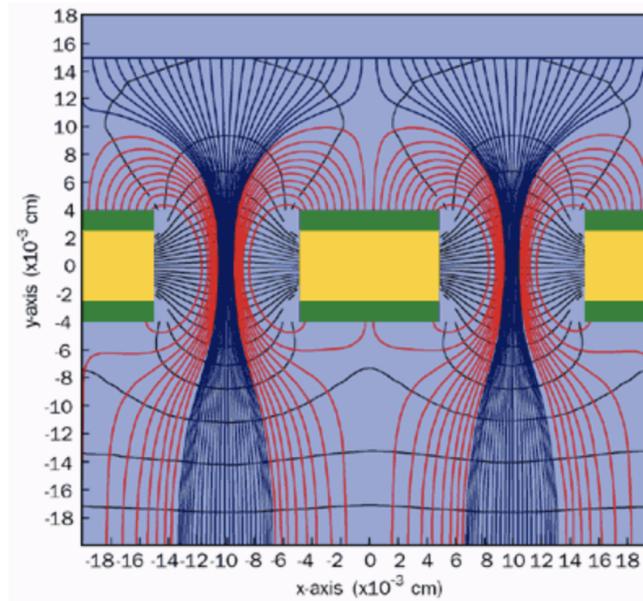


Figure 2.3 – Cross-sectional simulated view of the electric fields of two GEM foil holes

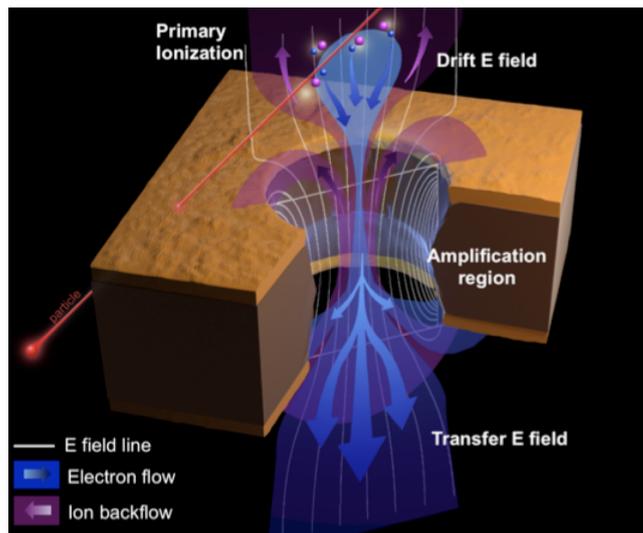


Figure 2.4 – Path of an electron through the electric fields present within a GEM foil hole

When an energetic particle such as a muon enters the gas volume, it ionizes the gas within. The number of primary electrons created in this way has a Poisson distribution whose average value is dependant on both the energy of the ionizing particle and the gas mixture used in the detector [28]. These electrons then encounter the electric field at the

top of the first foil layer, called the drift field. They follow the field lines of this drift field towards the closest hole within the GEM foil, as shown in Figure 2.4.

Once within the hole, they encounter what is known as the amplification region. It is here that the ionized electrons gain enough energy to produce their own secondary electron-ion pairs, as shown in Figure 2.5. Each of these secondary electrons can then gain enough energy to produce their own additional electron-ion pairs, and so on, in a multiplication mechanism known as a Townsend avalanche (often colloquially called an electron avalanche). In the case of this secondary ionization, we can quantify the number of secondary electrons produced by a given primary electron per unit path length via the first Townsend coefficient, α .

As each pair is produced, the negatively-charged electrons drift towards the anode, and the positively charged ions drift towards the cathode. As the ions are larger and therefore heavier than the electrons, they drift towards the cathode more slowly than the electrons towards the anode, producing the characteristic teardrop shape seen on the right of Figure 2.5. This characteristic shape was imaged in a cloud chamber by Heinz Raether in 1964, as seen in Figure 2.6.

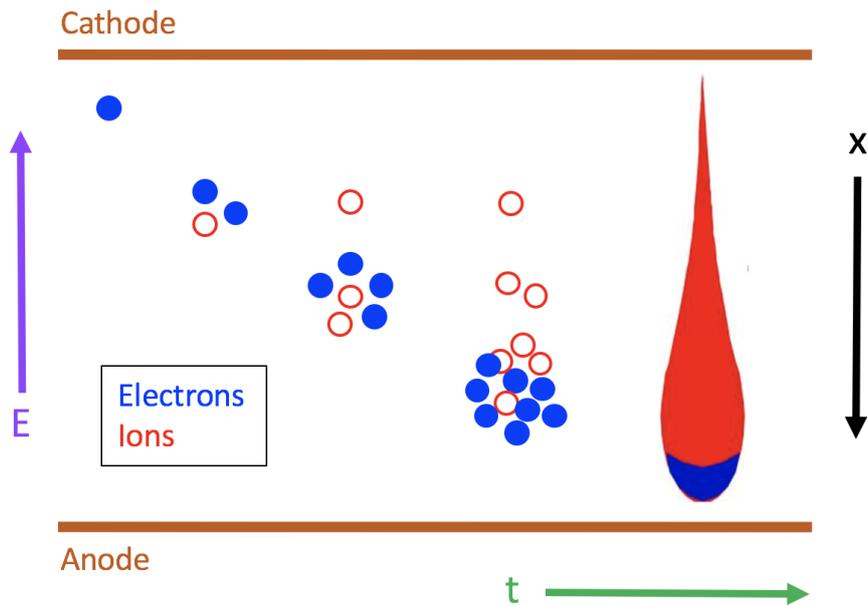


Figure 2.5 – Mockup of a Townsend avalanche over time

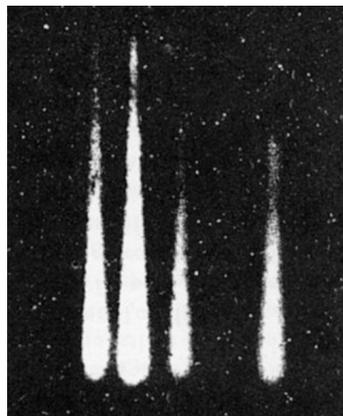


Figure 2.6 – Townsend avalanches imaged within a cloud chamber by H. Raether [29]

In GEM detectors with multiple GEM foils, such as the triple-GEM detectors used in CMS, each of the electrons produced in the Townsend avalanche then exits the hole and enters the transfer electric field between the first foil and the next. The process then begins again in much the same way for each proceeding foil, with the avalanche growing in size at each amplification region it encounters. This process is shown for a triple-GEM detector in Figure 2.7, with the primary ionizing event shown as the diagonal line across the detector and the electron avalanche branching off from it. In this case, when the avalanche exits the third and final foil, the electrons drift along the final electric field, called the induction field, until they reach the readout PCB. Here, a signal is induced on the readout strips or pad present there.

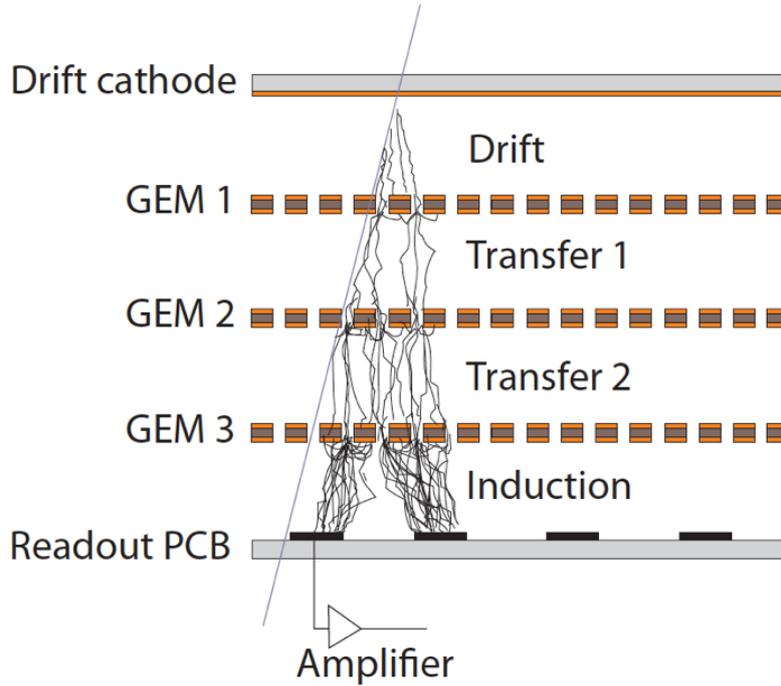


Figure 2.7 – Electron avalanche within a triple-GEM detector

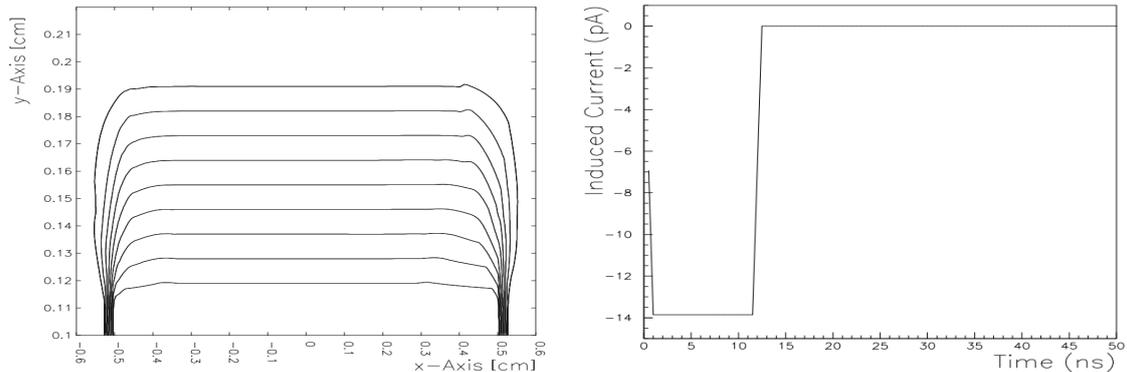
The formation of this induced signal depends entirely on the motion of the multiplied electrons within the induction region of the detector. We can therefore say that the signal formation begins the moment the first electron emerges from the lower electrode of the bottommost GEM foil and ends once the final electrons is collected on the readout electrode. Using the Shockley-Ramo Theorem, we can calculate the induced current in Equation 2.1. Here, I_k is the induced current on the k^{th} electrode due to a moving charge q at velocity $\vec{v}_d(x)$, and $\vec{E}_k(x)$ is the electric field which is created by raising the k^{th} electrode to a potential V_k .

$$I_k = -\left(\frac{q\vec{v}_d(x) \cdot \vec{E}_k(x)}{V_k}\right) \quad (2.1)$$

If we assume that $V_k = 1$ V and all of the other readout strips or pads are connected to ground, then this equation simplifies into Equation 2.2, where $\vec{E}_k^w(x)$ is the weighting field, the component of the electric field in the direction of v at the charge's instantaneous position if the charge were removed and the above assumed conditions applied.

$$I_k = -q\vec{v}_d(x) \cdot \vec{E}_k^w(x) \quad (2.2)$$

In his PhD thesis in 2002 (see Reference [30]), Davide Pinci simulated the behavior of $\vec{E}_k^w(x)$ within a triple-GEM detector and found it to be generally consistent within the induction region of the GEM foil, as shown in Figure 2.8a, implying that the electron drift velocity is uniform within this region. This therefore suggests that each electron which emerges from the bottommost GEM foil induces a rectangular signal on the nearest readout structure, as seen in Figure 2.8b, whose width is dependent on the time t required for the electron to cross the entirety of the induction region.



(a) Weighting field of a 1 cm^2 readout pad with a 1 mm induction gap (b) Signal current induced on a 1 cm^2 readout pad by one electron drifting in the induction gap

Figure 2.8 – Simulation studies performed by D. Pinci in 2002 on a triple-GEM detector with pad readout structures [30]

This is described in Equation 2.3, where x is the height of the induction region and v_d is the electron velocity in that region.

$$I = -\left(\frac{q}{t}\right) = -\left(\frac{qv_d}{x}\right) \quad (2.3)$$

2.3 Gain and Affecting Factors

The multiplication from the primary electron to the final induced signal charge is known as the gas gain, or simply the gain. In the case of the CMS GE1/1 detectors, this gain is approximately 10^4 . However, this is a variable parameter which depends on many factors, including the gas mixture, the applied voltages as they relate to the density of electric field lines within the amplification channel, the structure of the detector itself, and the parameters of the GEM foil holes including diameter, shape, and pitch.

2.3.1 Collection and Extraction Efficiency

However, before we can discuss the features which affect the gas gain, we must first clarify the type of gain that is of interest to us. That is, in most cases, we do not measure the real, intrinsic gas gain (G_{int}), but rather the effective gain (G_{eff}). This is a crucial distinction, as the effective gain, which is defined as the ratio of the detected charge to the primary ionization charge, is generally lower than the real gas gain of the amplification stage. In order to explore the reasons for this, we must define several parameters.

First, we define the collection efficiency in Equation 2.4 as the ratio between the number of primary electrons which are generated above the GEM foil and the number of electrons which actually enter a GEM foil hole:

$$\epsilon^{coll} = \frac{\text{electrons collected in the holes}}{\text{electrons produced above the holes}} \quad (2.4)$$

If conditions are such that the electric field lines are not effectively focused into the GEM foil holes, or there is insufficient electron diffusion above the foils, then the collection efficiency will be reduced [31]. This is because, as a result of these effects, some portion of the primary electrons are collected on the upper GEM foil electrode, or otherwise enter the hole but hit the polyimide surface between the two electrodes and are collected there rather than multiplying in and passing through the hole. For a gas mixture of Ar:CO₂:CF₄ 45:15:40, these effects amount to a loss of 20% and 5% respectively [28].

In addition, in the case of electronegative gas mixtures, additional losses can occur prior to multiplication as a result of recombination. Again taking the case of a gas mixture of Ar:CO₂:CF₄ 45:15:40, the high electric field near the hole ($\sim 10 \frac{kV}{cm}$) can result in the recombination of the primary electron as a result of the high electron attachment with regards to the first Townsend coefficient. This effect amounts to a loss of 10% in addition to the two effects described above, for a total of a 65% collection efficiency [28].

The second parameter we can define is the extraction efficiency, defined in Equation 2.5 as the ratio between the number of electrons which are produced within the holes during the multiplication stage and the number of electrons which are extracted from the holes.

$$f^{ext} = \frac{\text{electrons extracted from the holes}}{\text{electrons produced inside the holes}} \quad (2.5)$$

Again in this case, diffusion plays a role in the loss of efficiency. As a result of diffusion, the multiplication electrons are not confined to the center of the GEM foil hole. Thus, in cases where the electric field below the GEM foil is relatively low, some of these electrons might be collected on the bottom electrode of the foil rather than being transferred to the anode or onto the next GEM foil. In simulation studies performed by Davide Pinci [30], for an induction field of $5 \frac{kV}{cm}$, 3% of electrons are trapped at the surface of the hole due to diffusion, 10% are ion captured within close proximity of the bottom of the hole, and of the remaining electrons, 50% are lost to the lower GEM foil electrode, leading to a 35% extraction efficiency.

From these two parameters, we can further define a composite parameter called the electron transparency (T) of a GEM foil as being the combination of the collection efficiency and extraction efficiency, and thus relate the two efficiencies to the effective gain.

$$G_{eff} = G_{int} \cdot T \quad (2.6)$$

$$G_{eff} = G_{int} \cdot \epsilon^{coll} \cdot f^{ext} \quad (2.7)$$

2.3.2 Single-, Double-, and Triple-GEM Detectors

As discussed in Section 2.2, the multiplication of the primary electron takes place in the amplification region within the GEM foil holes. As a result, it is simple to extrapolate that the addition of more than one GEM foil to a detector would result in an increase in total gain. We find this to be true, as shown in Figure 2.9. Thus, we see that a double-GEM detector can reach higher gains at lower applied voltages than a single-GEM detector can, and a triple-GEM detector can reach higher gains at lower applied voltages than a double-GEM detector can.

In addition to this desirable increase in gain and decrease in applied voltage, we see in Figure 2.9 that the discharge probability is strongly reduced with the presence of additional foils, as the amplification is split across multiple cascaded stages. As electrical discharges are an inherent danger for gaseous detectors (see Chapter 7), this is an incredibly desirable trait, hence the choice to utilize them for the CMS GEM project.

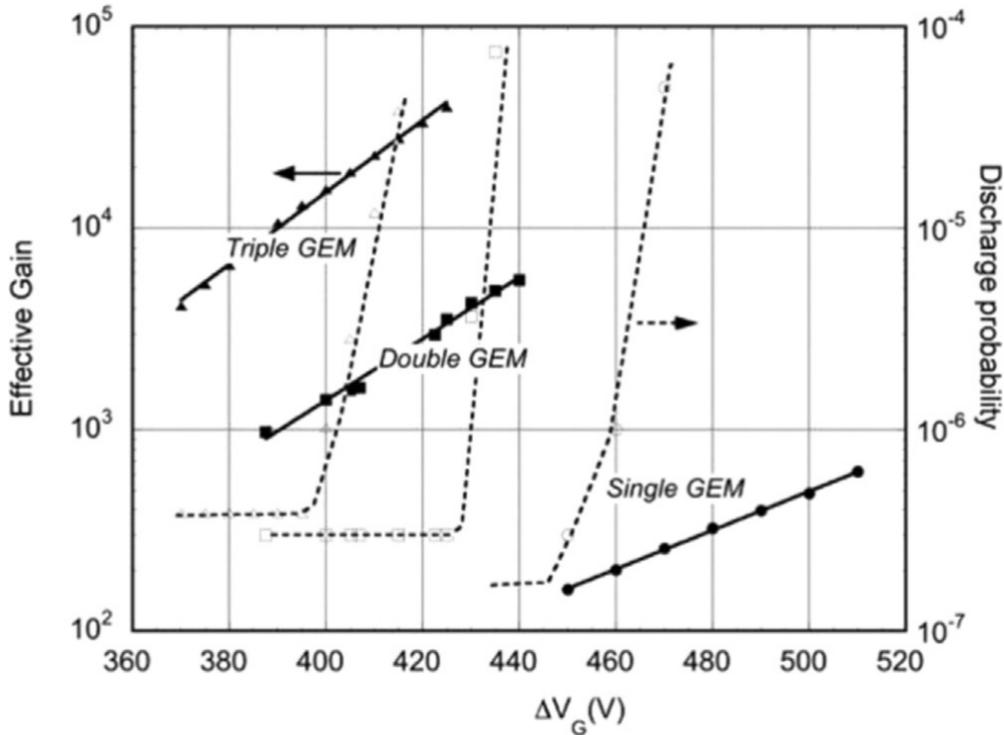


Figure 2.9 – Effective gain as a function of voltage for different multi-GEM detector configurations as well as their discharge probabilities [32]

However, not all triple-GEM geometries are made equal. Many factors other than the number of GEM foils must be taken into account, including the height or thickness of the various regions, the voltages applied to the GEM foils, and the strength of the electric fields.

2.3.3 Drift Region

The thickness of the drift region of the detector, shown in Figure 2.7 as the area between the drift cathode and the first GEM foil, must be chosen such as to ensure a high intrinsic detection efficiency. As this is the region where the primary ionization occurs, it must be large enough to ensure a sufficient number of electron-ion clusters such that minimally-ionizing particles can still be efficiently detected, but not so large as to increase either the pile-up effects that are seen at high particle rates or the rate of ageing of the detector itself. For the standard gas mixture of Ar:CO₂ 70:30, the ideal thickness is generally considered to be 3 mm.

The electric field present in this region must also be carefully chosen, taking note that its main function is to direct the primary electrons which are produced within the drift region into the holes present in the first GEM foil. Despite being the furthest region from the readout board, this electric field, called E_D , has a large effect on the signal which is induced there.

In Figure 2.10, we see the effect E_D has on the relative amplitude of the final signal (i.e. the collection efficiency) of a single-GEM detector operated in Ar:CO₂ 70:30. Here, three measurements were made at different foil voltages, shown in the key as G1 (V) / G2 (V). All three measurements were made with an induction electric field of $E_T = 5 \frac{kV}{cm}$ and a transfer electric field of $E_T = 3 \frac{kV}{cm}$.

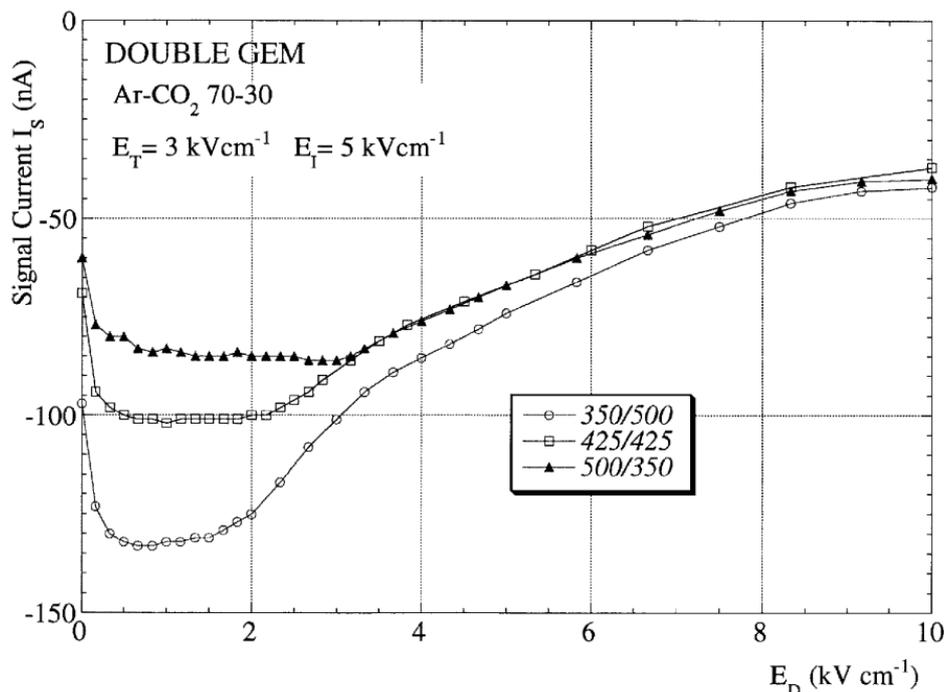


Figure 2.10 – Signal current induced on the readout board of a double-GEM detector vs. the drift electric field [33]

At low values of E_D ($<0.5 \frac{kV}{cm}$), the three curves exhibit a sharp decline in signal current. This is the result of a low electron drift velocity combined with a large diffusion, which results in a greater amount of recombination. Similarly, for high values of E_D ($>3 \frac{kV}{cm}$), the signal current again decreases. In this case, the decrease is the result of an increasingly large amount of the electric field lines terminating on the top electrode of the GEM foil rather than focusing into the holes, leading to a higher fraction of the primary electrons being collected on the electrode. Thus, in order to maximize the collection efficiency and drift velocity, the ideal drift electric field value lies somewhere on the plateau, between 0.5 - $2.5 \frac{kV}{cm}$, though this value may differ depending on the particular gas used and settings applied.

2.3.4 Induction Region

The thickness of the induction region of the detector, shown in Figure 2.7 as the area between the bottom of the final GEM foil and the readout board, must be chosen so as to optimize the fraction of collected signal. From Equation 2.3, we know that this is dependent in part on the electron drift velocity and the thickness of the induction region, which therefore suggests that the signal can be optimized by using a fast gas mixture and/or a small induction gap. However, this equation does not take into account the effects of having too small of an induction gap. For sub-millimeter gaps, a high mechanical precision would be required in order to avoid electrical discharges onto the readout board

and in order to maintain proper gain uniformity over the entirety of the detector area [28]. As this is both difficult and expensive to achieve, such a small gap is not considered. For the standard gas mixture of Ar:CO₂ 70:30, the ideal thickness is generally considered to be 1 mm.

From Equation 2.2, we see that the electric field in this region also plays a crucial role, and thus must also be carefully chosen. The main purpose of the induction electric field, called E_I , is to extract the electrons from the holes in the bottommost GEM foil and transfer them to the readout board.

In Figure 2.11, we can see the effect of E_I on the current induced on the various electrodes of a single-GEM detector with a GEM foil hole configuration of pitch / diameter = 140 μm / 70 μm . Here, the drift electric field of $E_D = 1 \frac{\text{kV}}{\text{cm}}$ has been chosen to be on the plateau for this gas mixture. I_S is the current collected on the readout pads, as is desired, I_T is the current on the top electrode of the GEM foil, I_B is the current on the bottom electrode of the GEM foil, I_D is the current on the cathode, and $I_{TOT} = I_S + I_B$ is the real GEM gain.

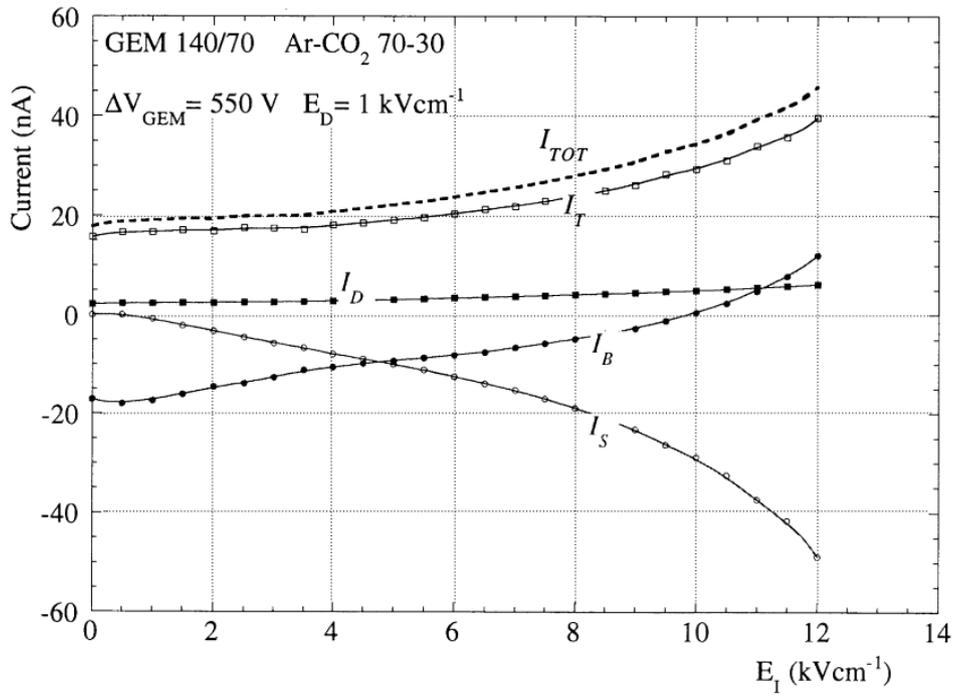


Figure 2.11 – Current on the various electrodes of a single-GEM detector vs. the induction electric field [33]

For the null case, where $E_I = 0 \frac{\text{kV}}{\text{cm}}$, we see that I_S is 0 and I_B is at a maximum. This represents the case in which the electrons within the GEM foil holes are entirely collected on the bottom electrode of the GEM foil, as the electric field lines within the holes terminate on said electrode. This is confirmed by the Shockley-Ramo Theorem (Equation 2.1), which suggests that in the absence of an electric field in the induction region, the induced current at the readout board is $I_S = 0$.

From there, increasing E_I increases I_S and decreases I_B , thus allowing more and more electrons to be transported across the induction gap and collected on the readout board rather than being collected on the bottom GEM foil electrode. This trend is shown for several voltage configurations in Figure 2.12, with the data key displaying voltages as G1 (V) / G2 (V).

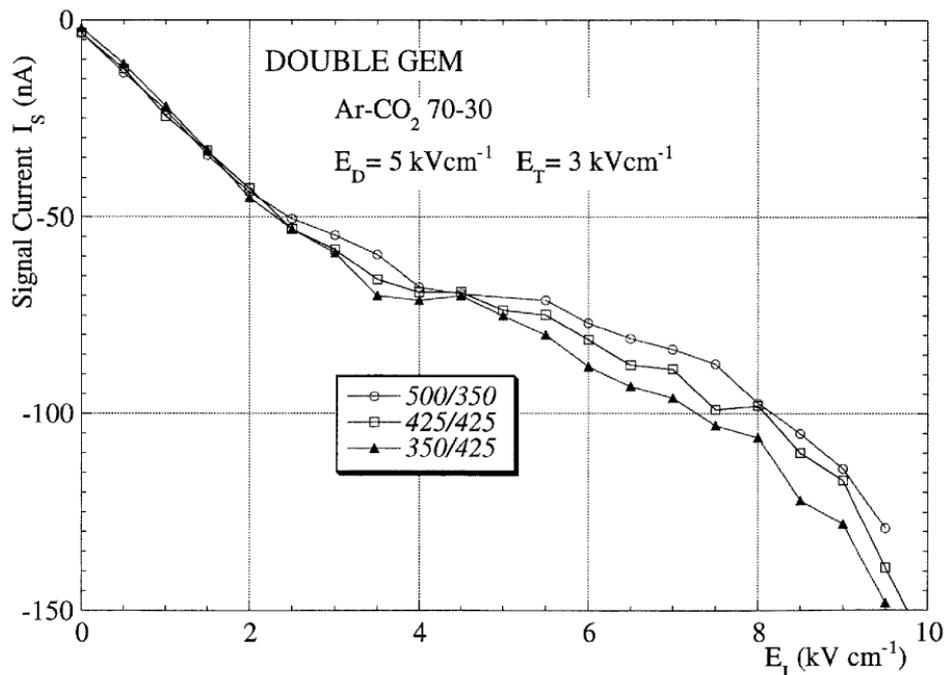


Figure 2.12 – Signal current induced on the readout board a double-GEM detector vs. the induction electric field for various voltage configurations [33]

However, as with the height of the induction gap, this does not tell the whole story. Rather, for values of E_I greater than $\sim 8 \frac{kV}{cm}$, the probability of a discharge to the anode greatly increases. As a result, a value must be chosen such as to maximize the signal collection without increasing the discharge probability to too great of a degree. In general, this value is chosen to be $\sim 5 \frac{kV}{cm}$.

2.3.5 Transfer Region(s)

For GEM detectors other than single-GEM detectors, there may be one or more transfer regions, which are regions that are sandwiched between two GEM foils rather than a GEM foil and either the anode or cathode. These are shown in Figure 2.7 as the areas between G1 and G2, and G2 and G3. The thickness of the transfer gaps is chosen so as to minimize the discharge probability while taking the time performance of the detector into account. However, they need not be of equal thickness, as each transfer region must take into account different considerations.

For the first transfer gap, the gap between G1 and G2, one particularly important consideration is the minimization of the bi-GEM effect. To understand this effect, we must recall that any electron produced within the gas volume of the detector must, by necessity, pass through all subsequent GEM foil layers if it is to be read out as signal on the anode. For a primary electron formed in the drift region by an ionizing particle, this means passing through three amplification stages in a triple-GEM detector, one at each foil layer. An electron produced in the first transfer region, on the other hand, would only need to pass through the remaining two amplification stages. If the amplification of those stages was large enough, electrons produced in this region could reach a large enough amplitude to be discriminated by the readout electronics. Considering a drift velocity v_{drift} and a transfer region thickness g_t , this signal could be read out at some time difference Δt from the original signal of the electron produced within the drift region, as given by Equation 2.8

[34].

$$\Delta t = \left(\frac{gt}{v_{drift}} \right) \quad (2.8)$$

This significantly impacts the time resolution of the detector and must therefore be minimized. Several studies by G. Bencivenni et al. (References [34], [35]) have shown that a significant reduction in this bi-GEM effect can be achieved with a transfer gap thickness of 1 mm. As with the height of the induction gap, smaller gap thicknesses, though they may provide theoretical benefits, are in practice difficult and expensive to produce and maintain uniformly over the active area of the detector, and thus are not recommended.

In the second transfer gap, however, this effect becomes negligible, since as the electrons only pass through one amplification region before reaching the anode, it is significantly unlikely that they could reach an amplitude over the electronics threshold. Thus, the focus in this region lies more on reducing the discharge probability. If the number of electron-ion pairs were to exceed the Raether limit (see Section 5.1 for more details), causing the electron avalanche to transition into a streamer within the third multiplication stage, a discharge could occur. In order to minimize this probability without relying on the addition of quenching gases, we instead focus on the thickness of the transfer gap. The transverse dimension of the electron cloud increases by the square root of the electron drift [36], which is itself dependent on the width of the gap. Thus, by increasing the width of the transfer gap, the electron diffusion within the region is increased, and the probability of overcoming the Raether limit within a given GEM foil hole is reduced and therefore so too is the discharge probability.

Although the amount of this increase is not constrained by physics, it is nevertheless constrained by the requirements of the experiment in which the GEM detector is located. In CMS, for instance, the GE1/1 detectors for the muon endcaps were restricted to the space that was originally foreseen for RPCs. As this space could not be altered without dramatically affecting the rest of CMS, the GEM detectors were placed under strict thickness limits. In order to comply with these limits while still retaining a large reduction in discharge probability, the choice was made to make the second transfer gap 2 mm wide. However, this is specific to this particular case, and as such, other GEM detectors in different applications may choose different, often larger, values.

For the electric fields within these regions, we do not have to make vastly different considerations for each region. In both cases, the main purpose of the electric field is to extract electrons from the holes of the upper GEM foil and transfer them in a focused manner into the holes of the lower GEM foil. As such, the goal of the field choice is to maximize both extraction efficiency from the upper foil and collection efficiency for the lower foil.

In Figure 2.13, we can see the effect of E_T on the signal current I_S collected from a double-GEM detector in Ar:CO₂ 70:30. Here, both the drift and induction electric fields have been chosen to be $5 \frac{kV}{cm}$. Three data sets are presented, with the data key displaying voltages as G1 (V) / G2 (V).

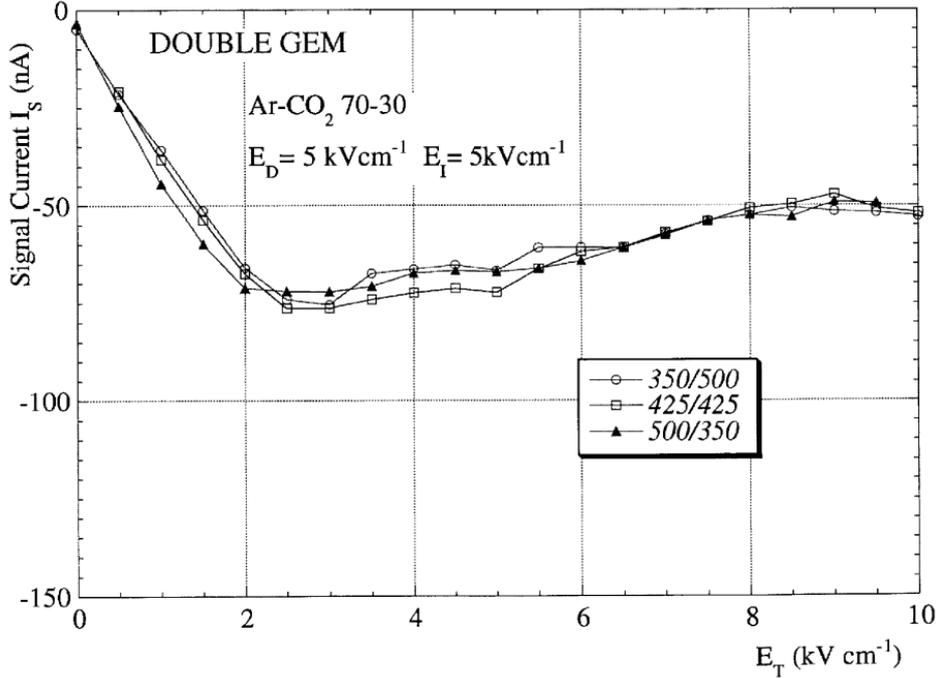


Figure 2.13 – Signal current induced on the readout board of a double-GEM detector vs. the transfer electric field [33]

At low values of E_T ($< 2 \frac{kV}{cm}$), the transfer field is not large enough to effectively transfer the electrons to the next GEM foil in the stack. As with the null case of E_I as described in Section 2.3.4, this results in the extraction of the electrons from within the holes largely by the electric fields within the holes themselves, causing them to be collected on the bottom of the G1 electrode rather than transferred to G2. By increasing E_T , this extraction efficiency improves.

However, for high values of E_T ($> 6 \frac{kV}{cm}$), we see instead a decrease in the collection efficiency, as the transferred electrons collect on the top electrode of G2 rather than being effectively focused into the holes within the foil.

Thus, in order to reach a balance between extraction and collection efficiency, for GEM detectors which use a gas mixture of Ar:CO₂ 70:30, the chosen value for the transfer electric fields generally lies in the range of 3 - 5 $\frac{kV}{cm}$.

2.3.6 GEM Foil Voltages

For single-, double-, and triple-GEM detectors alike, the intrinsic gas gain (G_{int}) of the detector is directly proportional to the voltages which are applied to the GEM foil(s). To show this, we start with Equation 2.9, where $\alpha(x)$ is the first Townsend coefficient and $\eta(x)$ is the attachment coefficient in the path δx , i.e. the probability that an electron drifting through a gas under the influence of a uniform electric field will undergo electron attachment in a unit distance of drift. Both $\alpha(x)$ and $\eta(x)$ are dependent on both the electric field and the gas mixture within the detector.

$$G_{int} = e^{\int [\alpha(x) - \eta(x)] \delta x} \quad (2.9)$$

For high values of electric field, such as those within the GEM detector holes ($\sim 100 \frac{kV}{cm}$), the attachment coefficient $\eta(x)$ becomes negligible. This allows Equation 2.9 to be

simplified into equation 2.10, where $\langle\alpha\rangle$ is the average of the first Townsend coefficient along the electron path through a given GEM foil hole.

$$G_{int} \propto e^{\langle\alpha\rangle \cdot V_{GEM}} \quad (2.10)$$

Then, in order to represent a triple-GEM detector, we expand this into Equation 2.11, with the product and sum operations going from $k=1$ to $k=3$ such that $V_{GEM}^{tot} = \sum_{k=1}^3 V_{GEM_k}$ is the sum of the voltage differences on each of the three GEM foils.

$$G_{int} \propto \prod_{k=1}^3 e^{\langle\alpha\rangle \cdot V_{GEM_k}} \propto e^{\langle\alpha\rangle \cdot \sum_{k=1}^3 V_{GEM_k}} \quad (2.11)$$

Taking into account the parameters discussed in Section 2.3.1, we can modify this equation to define the effective gain in Equation 2.12. As before, ϵ_k^{coll} and f_k^{ext} are the collection and extraction efficiencies of the k^{th} GEM foil.

$$G_{eff} = G_{int} \cdot T_{tot} = \prod_{k=1}^3 e^{\langle\alpha\rangle \cdot V_{GEM_k}} \cdot T_k = e^{\langle\alpha\rangle \cdot \sum_{k=1}^3 V_{GEM_k}} \cdot \prod_{k=1}^3 \epsilon_k^{coll} \cdot f_k^{ext} \quad (2.12)$$

The rightmost portion of Equation 2.12 is key, as it shows that the effective gain of a detector is affected by the sum total of the voltages applied to the GEM foils rather than any one applied voltage individually. Thus, for any given desired total voltage, the individual applied voltages can be engineered so as to minimize the probability of a discharge occurring within that portion of the detector and other undesirable effects.

In particular, we must consider the bi-GEM effect mentioned in Section 2.3.5. In order to minimize the probability of an electron in the transfer region of the detector reaching a high enough signal amplitude to be discriminated by the electronics, the first GEM foil should have an applied voltage that is significantly greater than the second GEM foil. In this way, the electrons generated in the transfer region will be amplified with a significantly lower gain than those generated in the drift region, decreasing the chance of the transfer electrons exceeding the signal discrimination threshold. As the bi-GEM effect does not affect the third GEM foil, we are thus left with the optical configuration described by pseudo-Equation 2.13.

$$V_{G1} \gg V_{G2} \geq V_{G3} \quad (2.13)$$

2.3.7 GEM Foil Parameters

In addition to the detector geometries, electric fields, and voltages described above, another consideration which affects the gain of the detector is the parameters of the GEM foil holes. One such parameter is the hole diameter. By reducing the diameter of the holes in the GEM foils, the electric field density in the amplification channel of the hole is increased, generally increasing the electron multiplication. This relationship is shown in Figure 2.14.

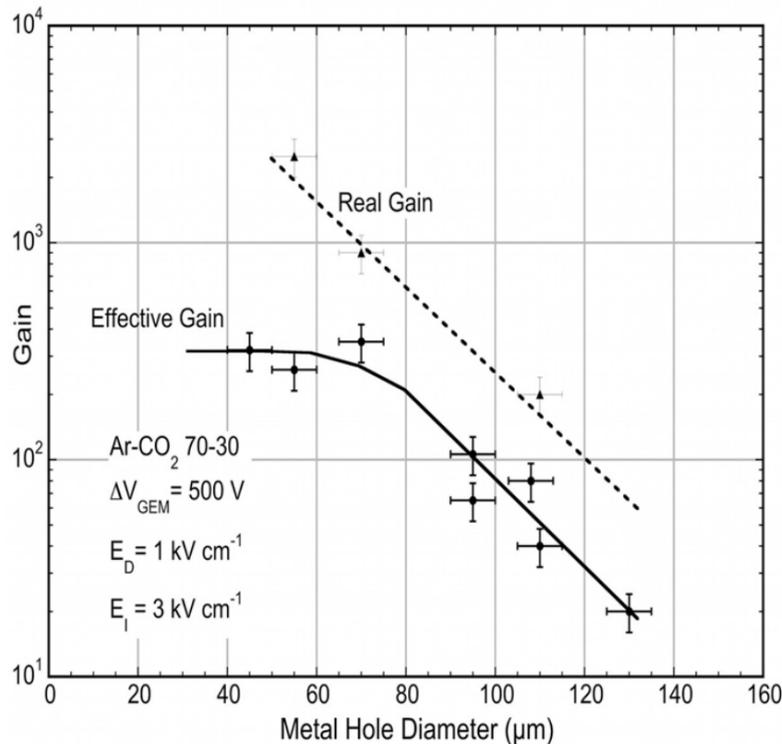


Figure 2.14 – Relationship between gas gain¹ and GEM foil hole diameter as measured in Ar:CO₂ 70:30 [32]

The effective gain curve also shows a plateau at hole diameters below $\sim 70\mu\text{m}$ in which the gas gain saturates due to the increasing losses of electrons to the lower GEM electrode. Despite limiting the possible gas gain increases available by decreasing hole width, this saturation is actually desirable, as it substantially reduces the dependence of the gas gain on the precision of the GEM foil manufacturing process, allowing for minor inconsistencies without drastically affecting detector performance. As a result, the standard choice for inner GEM foil hole diameter, $50\mu\text{m}$, falls safely within this plateau region.

Another GEM foil parameter which affects the gas gain is the shape of the holes. In gaseous detectors which utilize substrates with a high surface resistivity, a short-term “charging-up” effect can be observed. In GEM detectors, this effect manifests itself whereby some portion of the positive ions which are generated in the Townsend avalanche accumulate on the insulating polyimide of the GEM foil. As a result, the electric fields within the hole are slightly modified, and the gain increase at start-up is slowed. Thus, a hole shape which minimizes this effect would result in a faster start-up multiplication, and therefore a higher total gain.

The relative gain amplification factor is shown in Figure 2.15 for three different hole shapes - conical, biconical, and cylindrical. Here, we see that cylindrical holes offer a very stable operation, being least affected by the charging-up effect. However, this shape was also shown to be more prone to discharges at higher gains, a problem which is discussed in greater detail in Chapter 7. This shape also proved most difficult and costly to manufacture in standard GEM foils, and therefore was ruled out for the mass production of the CMS

¹Throughout this dissertation, the author has vowed to include any additions her cat, Sophie Mila Starling, wishes to include. Here, we find that the figure in question was deemed to be named “HoleDiameterGaitg6 78n.png”, with the “tg6 78” addition clearly having some critical application to the measurement process that only cats can truly understand.

GEMs.

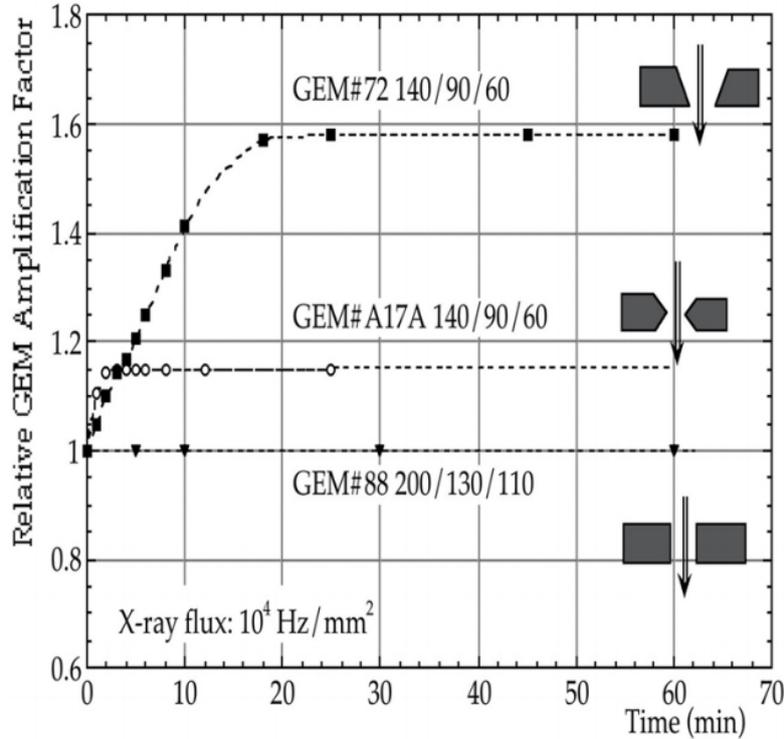


Figure 2.15 – Relationship between relative gas gain and GEM foil hole shape as measured in Ar:CO₂ 70:30, illustrating the charging-up effect as a function of time [32]

In contrast, the conical shaped holes would be the cheapest to manufacture, requiring only a single round of polyimide etching. However, in addition to the scaling problems described in Section 2.1, this shape is also most prone to the charging-up effect, as seen in Figure 2.15, making it an undesirable choice.

Thus, the standard choice for GEM hole shape is bi-conical, as this shape provides a good balance between the two, being only mildly affected by the charging-up effect, and requiring a manufacturing process that easily scales in terms of foil size without overbearing production costs.

Finally, while the pitch of the holes does not directly affect the gas gain, it does affect the collection efficiency of the electrons which are released in the upper areas of the gas volume, and therefore plays a role in the effective gas gain of the detector. Figure 2.16 shows this relationship for three different GEM foil geometries in a mixture of Ar:CO₂ 70:30, labeled on the plot as "GEM pitch/diameter" in μm . Unfortunately, no study of only hole pitch is available at this time, however, this still provides an insight into the affect hole pitch has on the electron transparency, and we can gather that higher electron transparencies, and thus higher effective gains, can be achieved with smaller GEM hole pitches.

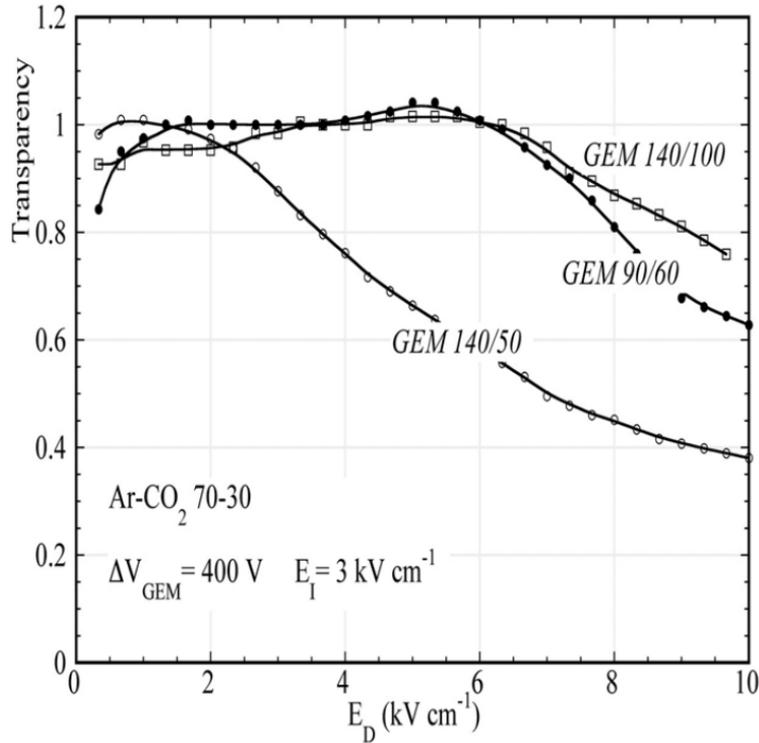


Figure 2.16 – Relationship between electron transparency and drift electric field (E_D) as measured in Ar:CO₂ 70:30 for different geometries of hole pitch / hole diameter [32]

2.4 Time Resolution

Despite the heavy focus of this chapter on the gain of GEM detectors, there is another crucial component that must be addressed: their time resolution, which must be high in order for the detectors to be of use within high rate experiments such as those of the LHC. For GEM detectors in particular, the time resolution is highly correlated with the statistical nature of the primary electron cluster which develops in the drift region after an ionizing particle interacts with the gas volume there. The number of these clusters, in turn, is dependent on the type of ionizing particle which created them, the energy of that particle, and the particular gas mixture used within the detector.

Using Poisson statistics, we can describe the spacial distribution ($A_j^{\bar{n}}(x)$) of the j^{th} electron cluster at a distance x from the first GEM foil, as in Equation 2.14. Here, \bar{n} is the average number of electron clusters per unit length x .

$$A_j^{\bar{n}}(x) = \left[\frac{x^{j-1}}{(j-1)!} \right] \bar{n} \cdot e^{-\bar{n}x} \quad (2.14)$$

Let's consider the first electrode cluster produced by a passing ionizing particle ($j = 1$). Equation 2.14 simplifies into Equation 2.15, with its root-mean-square standard deviation given by 2.16

$$A_j^{\bar{n}}(x) = \bar{n} \cdot e^{-\bar{n}x} \quad (2.15)$$

$$\sigma(x) = \left(\frac{1}{\bar{n}} \right) \quad (2.16)$$

Then, for a given drift velocity (v_d) in the drift region of the detector, we can use the previous equations to define the probability distribution (P_j) of the time of arrival (t_d) on the first GEM foil of the j^{th} and first electron clusters in Equations 2.17 and 2.18, respectively.

$$P_j(t_d) = A_j^{\bar{n}}(v_d \cdot t_d) \quad (2.17)$$

$$P_1(t_d) = \bar{n} \cdot e^{-\bar{n}v_d t_d} \quad (2.18)$$

Therefore, we arrive at the intrinsic time resolution of the detector, assuming the first electron cluster is detected 100% of the time:

$$\sigma_1(t_d) = \left(\frac{1}{\bar{n} \cdot v_d}\right) \quad (2.19)$$

From this equation, we see that the intrinsic time resolution of the detector is dependent solely on the electron drift velocity and the primary ionization that occurs within the drift region. As such, in order to maximize the detector's time resolution, a suitable gas mixture with a high drift velocity and a high primary ionization must be chosen.

Unfortunately, this simplicity assumes a 100% detection efficiency of the first primary electron cluster, which in reality is not the case. As a result of both the gain-affecting statistical fluctuations discussed in previous sections and the finite discrimination threshold of the front-end readout electronics, detection of this first electron cluster is highly unlikely, if not impossible. In order to generate a signal above this discrimination threshold, a certain amount of electron cluster pile-up is needed, which limits the time resolution of the detector.

One way of overcoming this effect is to use a gas mixture with a high drift velocity and operate the detector with a relatively low drift electric field. As discussed in Section 2.3.3, a low value of E_d permits a high collection efficiency for the amplification stage of the first GEM foil. This in turn ensures a high detection efficiency from the first electron cluster, thus improving the time resolution of the detector in addition to the previously-described benefits.

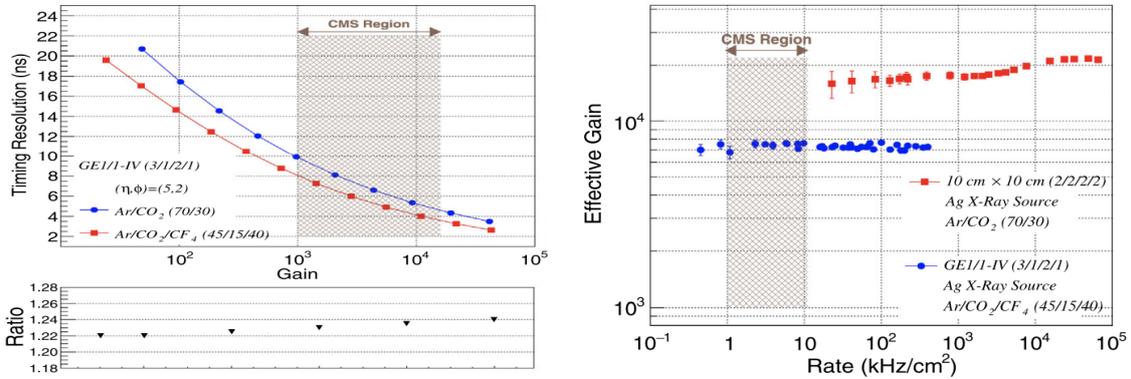
Regardless of this particular issue, the GEM detector still remains a fast detection technology compared to other currently-existing gaseous detectors, as the induced signal is reliant solely on the motion of the multiplied electrons through the stack of the detector to the readout board. The high mobility of the electrons results in a fast signal which is not affected by the ballistic deficit of detectors such as multi-wire proportional chambers (MWPCs), in which the signal pulse on the electrodes is formed by induction due to the movement of much slower ions. This inherent speed of signal formation is one characteristic which makes GEM detectors an ideal technology for use in new experiments and upgrades of current experiments.

2.5 Application to GE1/1 parameters

Applying the information presented in the previous sections, we can build up the ideal parameters for the GEM detectors used in the CMS GEM detectors.

A gas mixture of Ar:CO₂ 70:30 is used in order to provide an ideal balance between signal gain (via the Ar) and spark quenching (via the CO₂) without requiring strong

greenhouse gases such as SF_6 or CF_4 . The timing resolution and effective gain for a GE1/1-IV detector in this chosen gas mixture can be seen in Figure 2.17 in comparison with its forgone alternative of $\text{Ar}:\text{CO}_2:\text{CF}_4$ 45:15:40.



(a) Time resolution for a GE1/1-IV detector in $\text{Ar}:\text{CO}_2:\text{CF}_4$ 45:15:40 and $\text{Ar}:\text{CO}_2$ 70:30 as a function of gain (top) and their ratio (bottom) (b) The rate capabilities of a GE1/1-IV detector in $\text{Ar}:\text{CO}_2$ 70:30 for comparison in $\text{Ar}:\text{CO}_2:\text{CF}_4$ 45:15:40 and a $10 \times 10 \text{ cm}^2$ test detector in $\text{Ar}:\text{CO}_2$ 70:30 for comparison

Figure 2.17 – Comparison of operational ability for a GE1/1-IV detector in two different gas mixtures, $\text{Ar}:\text{CO}_2:\text{CF}_4$ 45:15:40 and $\text{Ar}:\text{CO}_2$ 70:30. Note that the shaded “CMS Region” spans the range of particle flux expected in CMS for the HL-LHC [37]

Triple-GEM detector technology is chosen over double-GEM or single-GEM in order to maximize the effective gain at lower applied voltages, and to reduce the probability of discharges occurring within the detector.

The GEM foils of the detector are etched using single-mask photolithography in order to ensure correct hole alignment for the large, meter-long foils. The holes are bi-conical in order to minimize the charging-up effect without greatly increasing production costs. The inner diameter of these holes is $50 \mu\text{m}$ and the outer diameter of these holes is $70 \mu\text{m}$, such that they fall within the plateau region of the effective gain of the detector, maximizing the gain while minimizing the reliance on the precision of the manufacturing process. The holes are etched at a pitch of $140 \mu\text{m}$, as this small hole pitch allows for higher electron transparency, and thus higher effective gains via a high collection efficiency.

The foils are arranged in a stack of three, with gas gaps of $x_{\text{drift}} = 3 \text{ mm}$, $x_{T1} = 1 \text{ mm}$, $x_{T2} = 2 \text{ mm}$, $x_{\text{ind}} = 1 \text{ mm}$, commonly referred to as the “3/1/2/1” configuration. x_{drift} is chosen to be 3 mm, as this is large enough to ensure a sufficient number of electron-ion clusters are produced such that even minimally-ionizing particles can still be efficiently detected, while not being so large as to increase either pile-up or ageing effects, particularly at a high particle rate. x_{T1} and x_{ind} are chosen to be 1 mm as this is the minimum height that is practically obtainable from both a financial and a production difficulty standpoint, with the former gap minimizing the bi-GEM effect and the latter gap maximizing the fraction of signal collected on the readout board as per Equation 2.3. x_{T2} is chosen to be 2 mm as the bi-GEM effect becomes negligible, and thus the focus shifts to reducing the discharge probability. By increasing the height of the gap to 2 mm, the overall thin profile of the detector is maintained while increasing the electron diffusion to such a degree that the probability of overcoming the Raether limit within a given hole is sufficiently reduced.

Finally, the applied voltages (and thus the fields) vary based on the individual detector, as well as the ambient temperature and pressure conditions. However, in general, the detectors are operated at a gain of $\sim 10^4$, requiring an equivalent divider current of 690 -

710 μA , which produces voltages on the order of those shown in Figure 2.18, which are in keeping with Equation 2.13, such that the electric fields are optimized.

	V0
Drfit	787.50
G1Top	394.10
G1Bot	306.60
G2Top	385.00
G2Bot	612.50
G3Top	367.50
G3Bot	437.50
readback setting	

Figure 2.18 – Example V0 settings for a GE1/1 detector at an equivalent divider current of 680 μA

In the most recent CMS GEM publication to the journal *Nuclear Instruments and Methods in Physics Research Section A* (see: Reference [37]), these considerations and their respective measurements are combined into one “master plot”, shown in Figure 2.19. Note that here, the “drift voltage” refers to the total voltage applied to the chamber, not the voltage applied individually to the drift layer, as this functionality was not available for the GE1/1-IV detector version. The shaded “CMS region” is then the recommended operational region for the detectors within the CMS experiment.

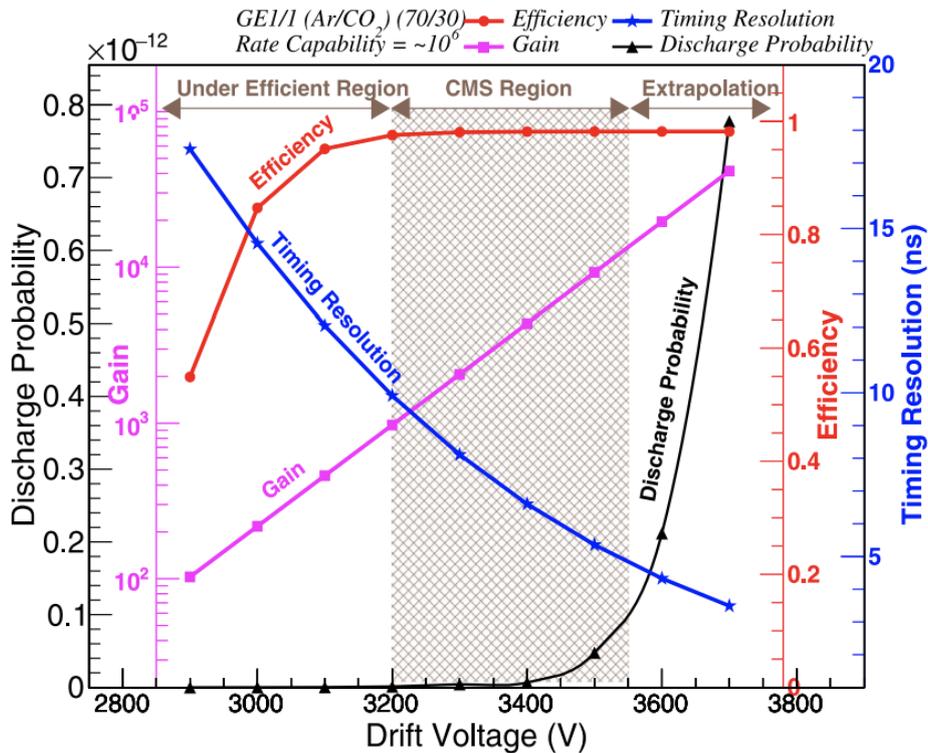


Figure 2.19 – “Master plot” showcasing the measurements of discharge probability, timing resolution, efficiency, and gain for a GE1/1-IV detector in Ar:CO₂ 70:30 [37]

In comparison, Figure 2.20 shows the same variables for the same chamber type using Ar:CO₂:CF₄ 45:15:40 instead of Ar:CO₂ 70:30, as well as the shaded recommended operational region for the detectors in this gas mixture.

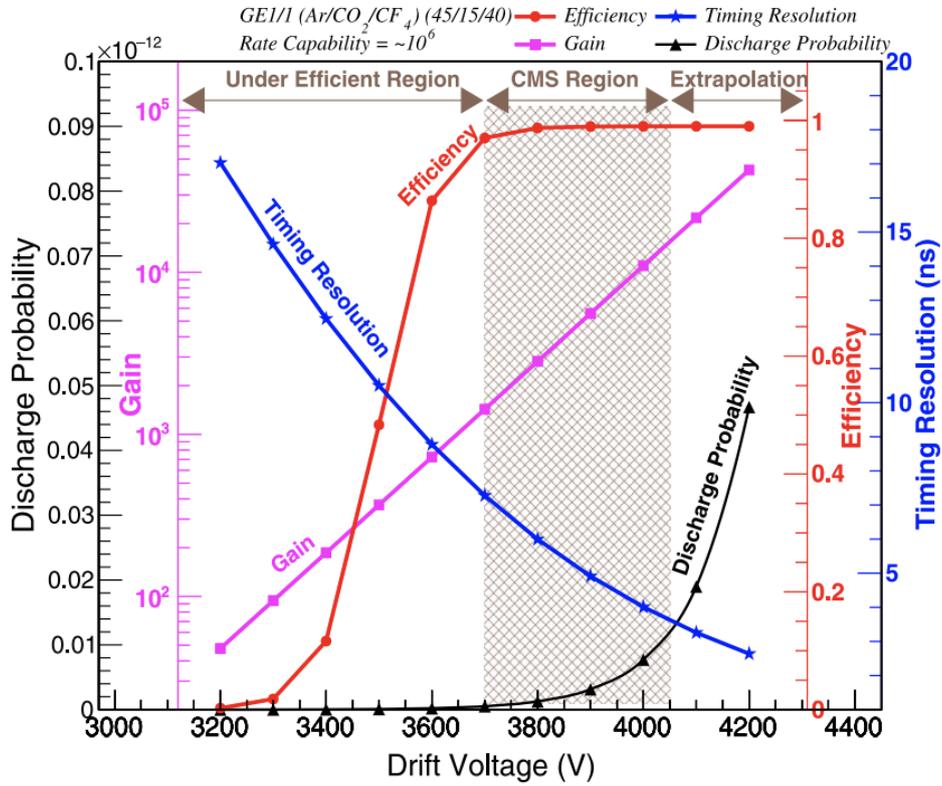


Figure 2.20 – “Master plot” showcasing the measurements of discharge probability, timing resolution, efficiency, and gain for a GE1/1-IV detector in Ar:CO₂:CF₄ 45:15:40 [37]

Note that here, a much larger drift voltage is required to be in the efficiency plateau region in comparison to that which is needed for Ar:CO₂ 70:30, which further accentuates the decision to use the more eco-friendly gas mixture.

Chapter 3

GE1/1 Readout Electronics

3.1 Introduction

Throughout the duration of this dissertation, two versions of the GEM readout electronics were in use, called v2¹ and v3 for version 2 and version 3. v2 was used solely for the ten slice test detectors, which will be discussed in Section 4.4, whereas v3 was used for the full, final GE1/1 system. The schematics for each can be seen in Figures 3.1 and 3.2, respectively.

This chapter will introduce each feature of the readout electronics, beginning first with v2, followed immediately after by v3. Note that for v3, the electronics described are the initial v3 electronics, and that certain changes that were made after the initial development are not included here - this most particularly applies to the VFAT hybrids, as the final changes made to these will be discussed in much greater detail in Section 7.2.

The GEM readout electronics can be broken into two categories: the front-end electronics, which are located on the detector, and the back-end electronics, which are located off-detector, either in the GEM racks on the side of the CMS detector within the experimental cavern (UXC55) or in the service cavern (USC55). Both will be discussed in this chapter.

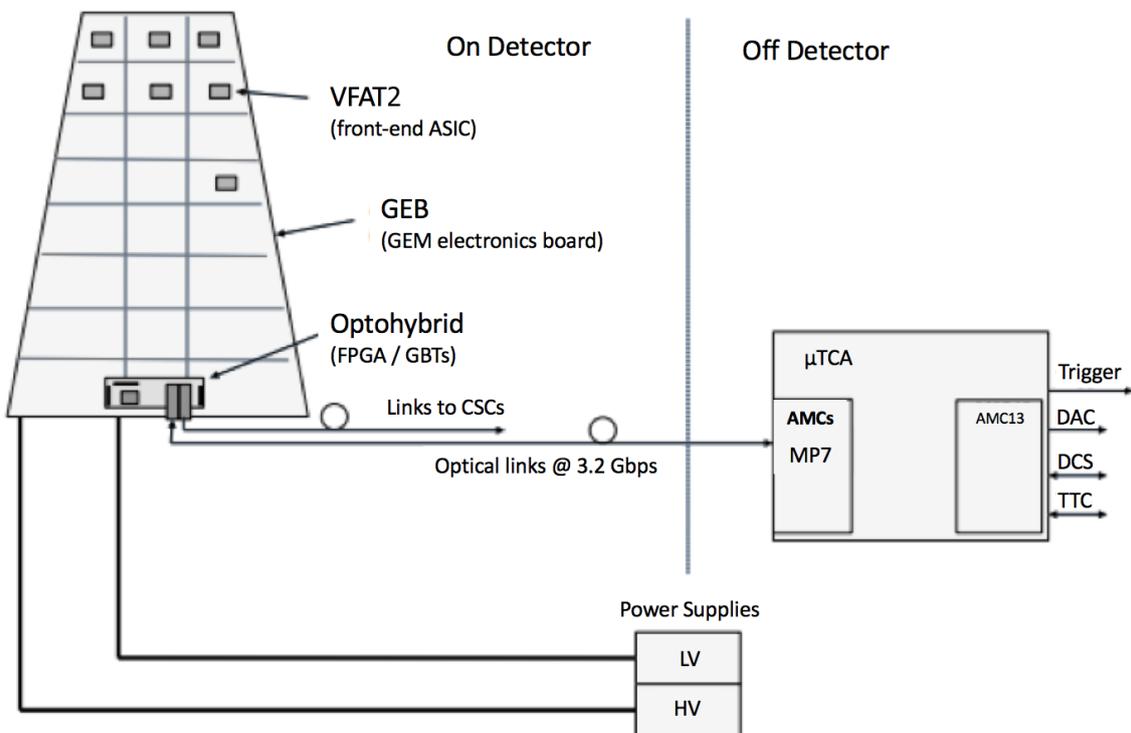


Figure 3.1 – v2 electronics schematic

¹v2 may also be called v2b, referring to a specific version of the v2 electronics - the version which was used in the slice test.

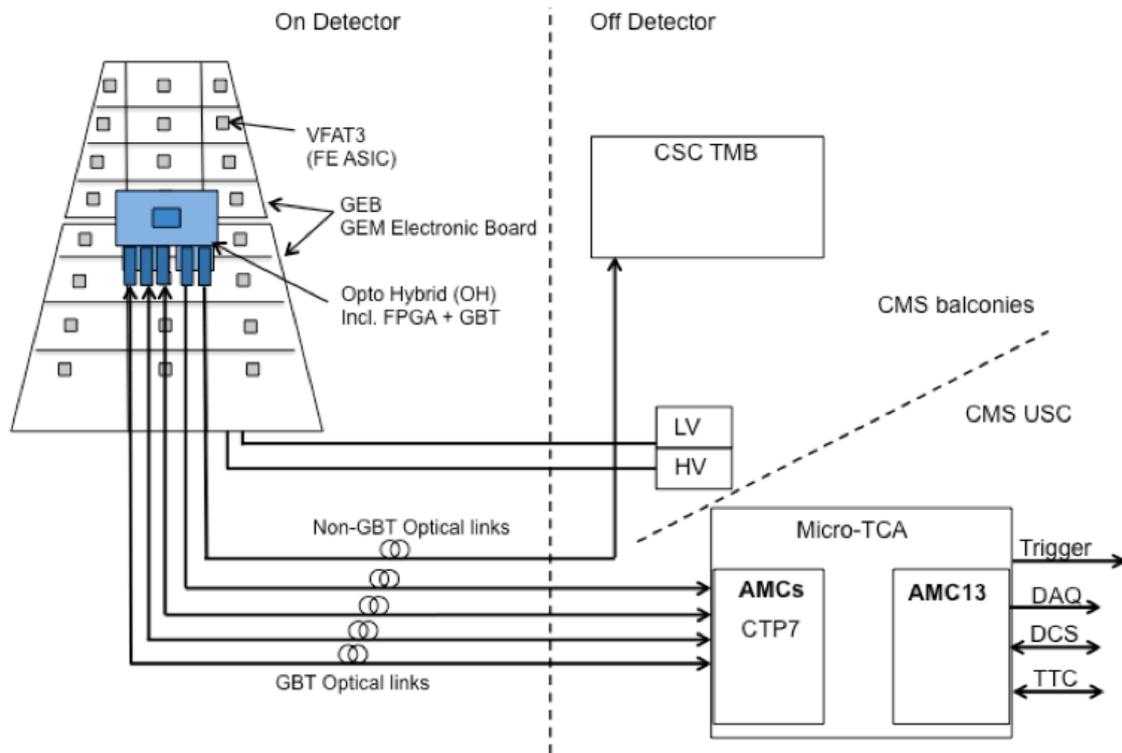


Figure 3.2 – v3 electronics schematic

3.2 Front-End Electronics

The front-end electronics have the special requirement of being radiation hard. Unlike the back-end electronics, which are mostly located within the safety of USC55, the front-end electronics must be able to reliably withstand a constant bombardment of radiation due to the collisions within the CMS detector.

3.2.1 Readout Board

The readout board is a $\sim 1\text{m}$ -long printed circuit board (PCB) which is located at the bottom of the GEM stack. As described in Section 2.1, when a Townsend avalanche exits the third and final foil, the electrons drift along the final electric field, called the induction field, until they reach the readout PCB. Here, a signal is induced on the readout strips present there, and this signal is directed to the input(s) of the corresponding VFAT chip.

The readout board is divided into 24 sections of 128 copper strips each, broken down into three sectors in ϕ (columns across the board) and eight sectors in η (rows down the board lengthwise). These sectors can be seen in Figure 3.3, where the sectors are labeled in ϕ ($i\phi$), η ($i\eta$), and by VFAT number.

Given the three sectors in ϕ , there are 384 strips across the width of the board in each η sector, and 3,072 total readout strips per detector. These copper readout strips are truly radial strips, with a width of $230\ \mu\text{rad}$ and a pitch of $463\ \mu\text{rad}$. These dimensions allow for the measurement of particle position to within the required angular resolution of $300\ \mu\text{rad}$ [38].

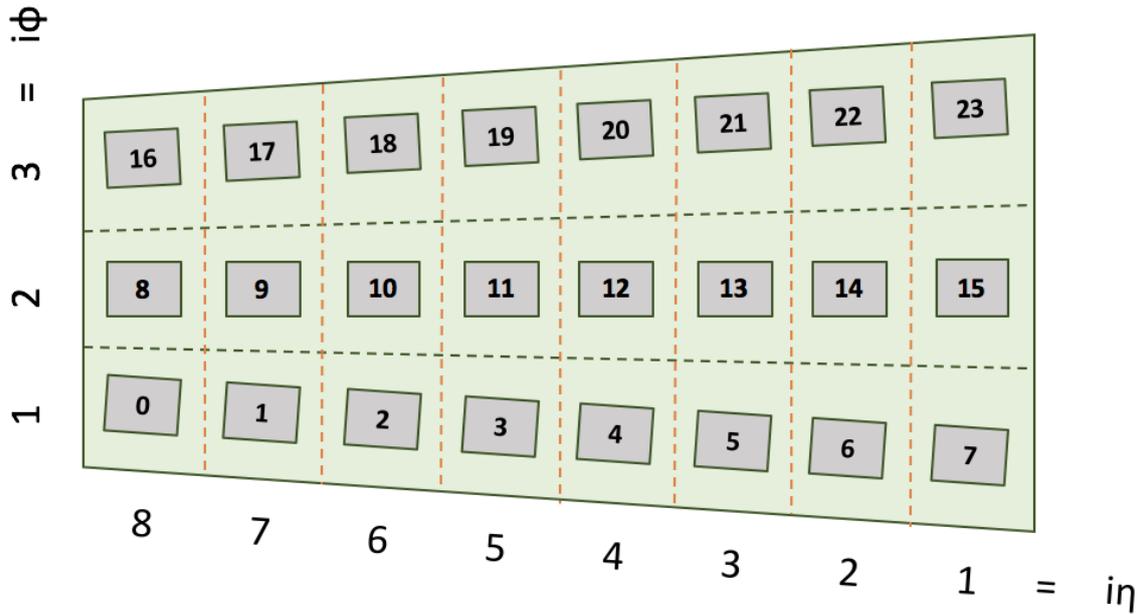


Figure 3.3 – Diagram showing the sector labeling of a GE1/1 readout board by ϕ sector number, η sector number, and VFAT number

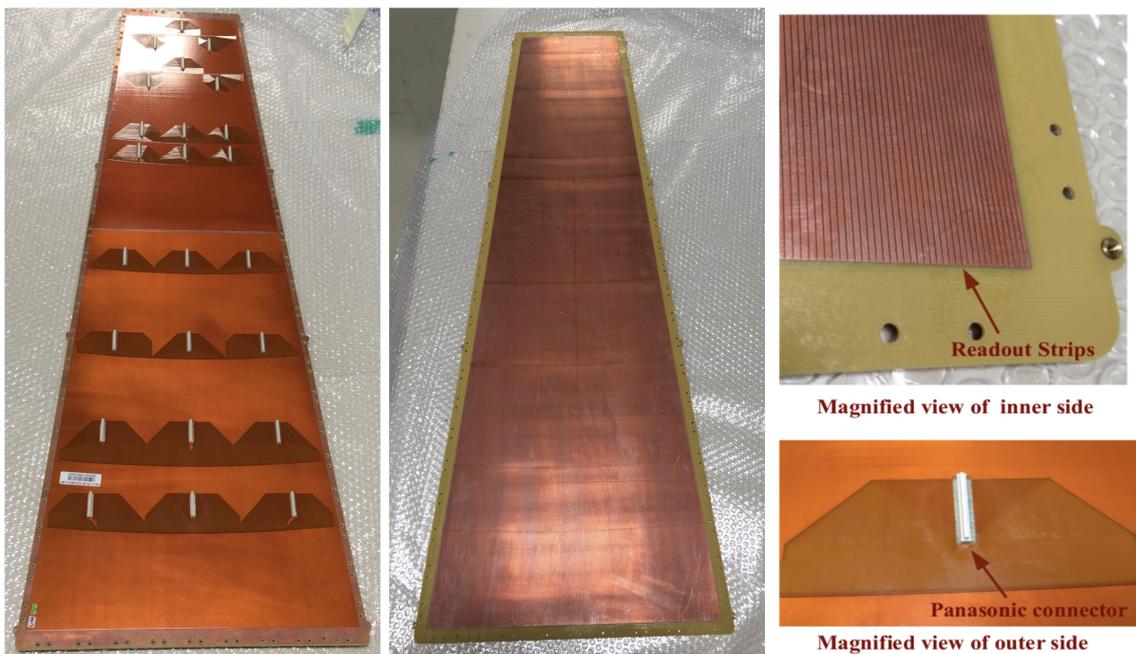


Figure 3.4 – Photos of a GE1/1 readout board, from the 2018 paper published to the journal *Nuclear Instruments and Methods in Physics Research Section A* (see: Reference [39])

Figure 3.4 shows several photos of a GE1/1 readout board. The leftmost photo shows the outer side of the readout board, the side which faces away from the active gas volume. Here we can see the 24 male Panasonic connectors, one for each VFAT / readout sector, and the traces that connect them to the readout strips on the other side. The bottom right image shows a magnified view of one of these Panasonic connectors. The middle photo shows the inner side of the readout board, the side which faces into the active gas volume. Here we see the 3,072 radial readout strips, with faint horizontal lines indicating the breaks

between η sectors. The top right image shows a magnified view of this inner side of the board, such that one can see the individual readout strips.

Unlike the remainder of the front-end electronics described in this chapter, the readout board did not undergo any noteworthy changes from v2 to v3.

3.2.2 GEM Electronics Board (GEB)

The GEM Electronics Board (GEB) is a PCB which is external to the GEM detector, and, unlike the readout PCB, it does not come into contact with the active gas volume. Its purpose is multifold: at its core, the GEB serves to host the connections and traces necessary to carry power to the FEAST, VFATs, and optohybrid, and to carry signals between the VFATs and optohybrid. It also acts as a shield, isolating the GEM detector from the noise produced by the front-end electronics.



Figure 3.5 – Photo of three stand-offs on a GEB [40]

The stand-offs, three of which are shown in Figure 3.5, are cylindrical threaded separators, which serve as connections to either ground or power. These stand-offs are soldered to a copper via in the GEB (one of which is shown in Figure 3.5, labeled H8), which allows for the necessary electrical connection. Their respective electronics, be it a FEAST or the optohybrid, is then carefully screwed into the threading.



Figure 3.6 – Application of Araldite glue to reinforce the stand-offs on a v3 GEB [41]

The soldering of these stand-offs is extremely fragile, as the exposed copper of the

via is very thin. As a result, these standoffs were prone to breakage in v2 and early v3 prototypes. In order to ensure the soldering connection is a good one, they must be soldered by hand rather than by more common mass production methods, and this hand-soldering is difficult, requiring an expert hand. To further reinforce the stand-offs, Araldite glue may be used for reinforcement, as shown in Figure 3.6. Furthermore, to protect them from breakage, a special torque screwdriver should be used during the installation of the front-end electronics, in order to ensure excessive force is not used.

Like the readout board, the GEB is divided into 24 sections. For each section, a rectangular hole is cut into the board, which allows the male Panasonic connector for the VFAT of that section to come through, allowing for the connection of the readout traces to the VFAT hybrid. Next to this hole, another, smaller male Panasonic connector is installed, which allows the data from the VFAT to be sent to the optohybrid. The VFAT hybrid is, of course, plugged into both connectors, as shown in Figures 3.7 and 3.8.



Figure 3.7 – Figure showing the connection of VFAT hybrids (grey) to the GEB. The GEB is shown in red, with holes in the GEB such as those which account for the connectors to the readout board are shown in blue. The external frame of the detector is shown in green. [38]

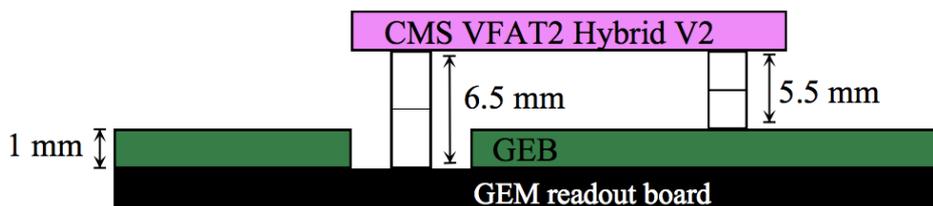


Figure 3.8 – Figure showing a side view of the GEB as it is installed on top of the GEM readout board, with a VFAT2 hybrid plugged into both Panasonic connectors, shown in white. Key measurements are included. [40]

3.2.2.1 v2

For v2, the GEB is, like the readout board, produced as a single $\sim 1\text{m}$ -long PCB. It can be seen as it is installed onto a detector in Figure 3.9, and a physical photo can be seen

in Figure 3.10. The long, light grey rectangles shown on the right of the image (bottom / wide end of the detector trapezoid) are the connections to the optohybrid, not shown.

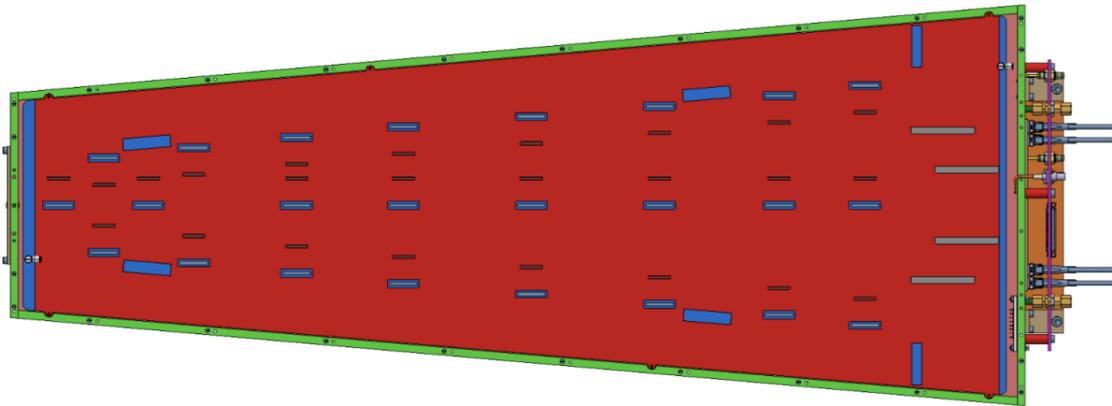


Figure 3.9 – Drawing of a v2 detector minus the front-end readout electronics. The GEB is shown in red, with holes in the GEB such as those which account for the connectors to the readout board are shown in dark grey. The VFAT Panasonic connectors are shown in blue-outlined grey, and the optohybrid connectors are shown in light grey. The external frame of the detector is shown in green. Gas, power, and other connections are shown on the right in orange. [38]

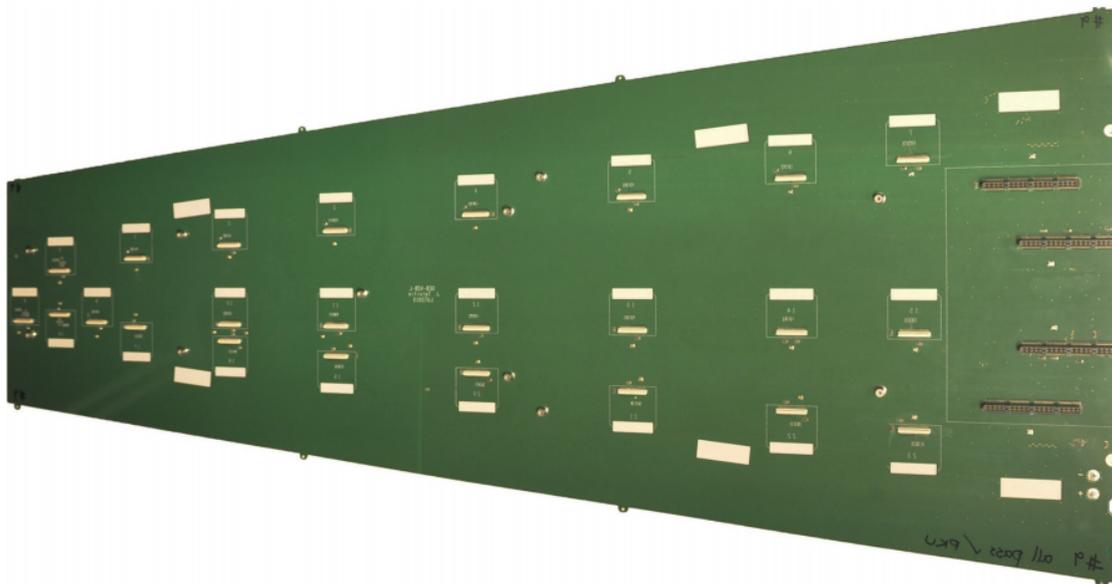


Figure 3.10 – Photo of a v2 slice test GEB [40]

However, despite the apparent simplicity of having the GEB be a single PCB, the design was not without its flaws. The GEB, despite being constrained to a thickness of only ~ 1 mm, is composed of six layers and has a large number of connections (~ 500 differential signals, 10 FEASTs, and analog and digital ground (AGND, DGND) to every VFAT) and connectors which add a great deal of mechanical stress to the PCB [40]. It also has a strict bending requirement, in which the GEB is required to be planar to within 0.75% of the measurement of its diagonal. These factors combine to make production of suitable GEBs both difficult and costly.

In addition, the placement of the optohybrid at the end of such a long PCB requires that traces such as those from the VFAT connectors at the narrow end of the GEB be very long. As the capacitance of the trace is directly proportional with its length, and the noise is proportional with the capacitance, this leads to unnecessary noise from the GEB. The length of the traces also brings concerns about signal integrity into account, especially given the VFAT3's usage of 320 MHz scalable low voltage signals (SLVS)[40].

3.2.2.2 v3

In order to solve the aforementioned problems, the v3 GEB was split into two boards of 12 VFATs each, a “narrow” board (seen on the left of Figure 3.11) and a “wide” board (seen on the right of Figure 3.11).

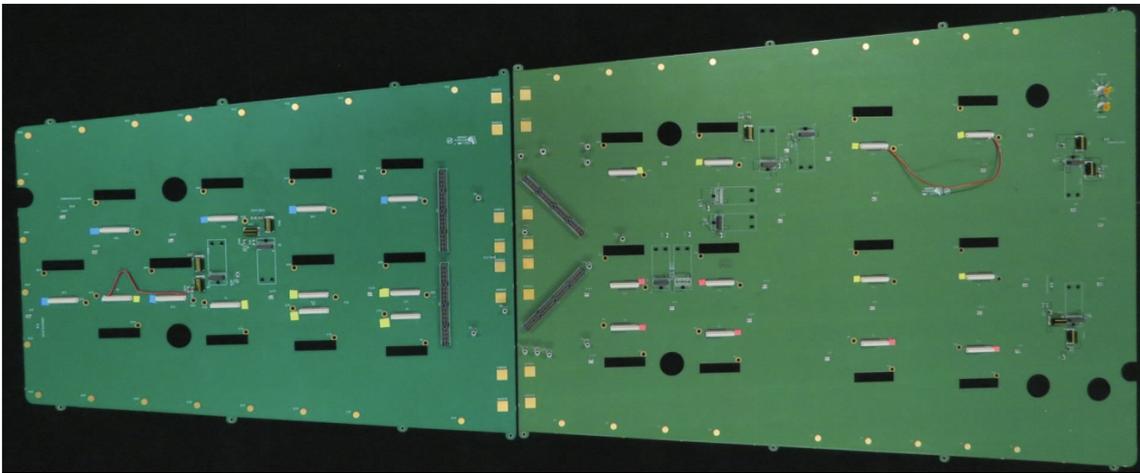


Figure 3.11 – Photo of the two portions of a v3 GEB [40]

Here, rather than connecting the optohybrid to the bottom of a single GEB, the optohybrid is connected in the middle of the detector, straddling the two portions of the GEB. The connectors which allow this are seen in the middle of Figure 3.11 in dark grey. This shortens the longest traces to approximately half of their length in the v2 GEB, significantly minimizing the noise from that particular source and improving signal integrity. Special care was taken when designing these traces, ensuring that the signal line lengths matched correctly to ensure VFAT3 synchronization, a problem which had been seen in the slice test v2 GEBs.

By splitting the GEB into two pieces and thus reducing the length of each portion, the problem of GEBs failing the planarity requirements is also reduced, though unfortunately not eliminated. Part of the reason it is not completely eliminated is that, in contrast to the six-layer v2 GEB, the v3 GEB is an eight-layer PCB, adding two additional shield layers between the GEB signal lines and the readout board. These new shield layers can be seen in Figure 3.12, which shows the layer stack of the v3 GEB PCB.

These shield layers were added as a direct result of the experience gained in the slice test (see: Section 4.4), and as a result of laboratory studies which confirmed that the GEB plays a large role in v2b electronics noise due to its strong coupling to the G3 bottom electrode. The results from this study are shown in Figure 3.13.

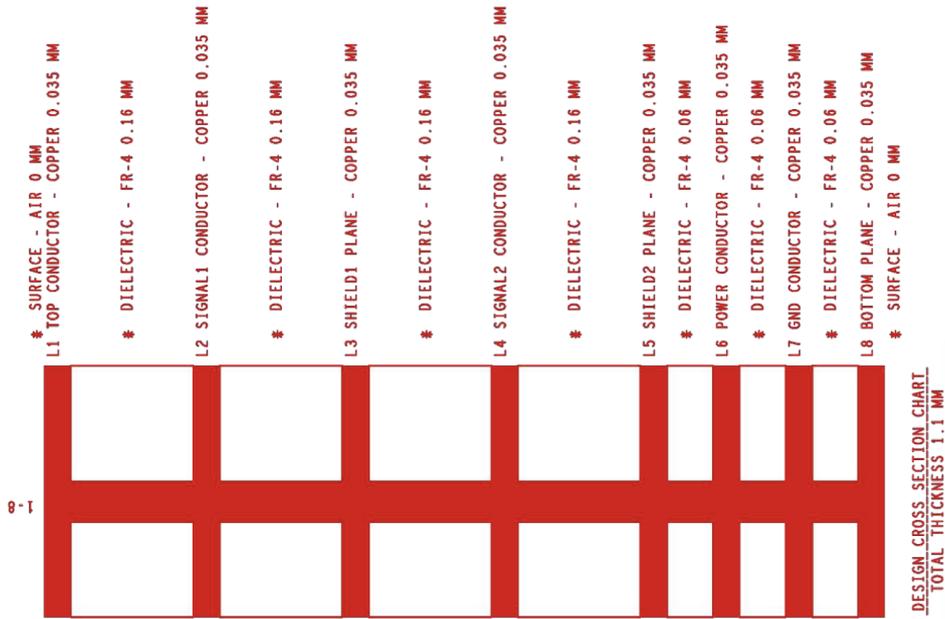
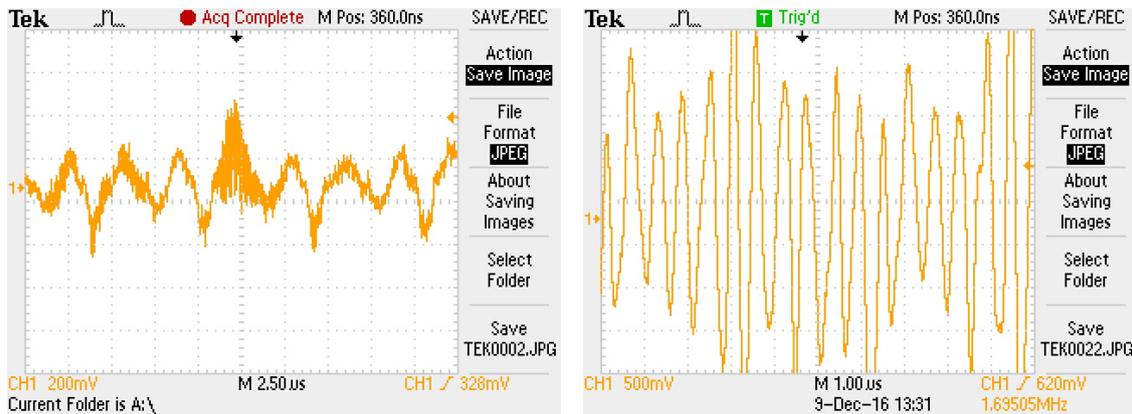
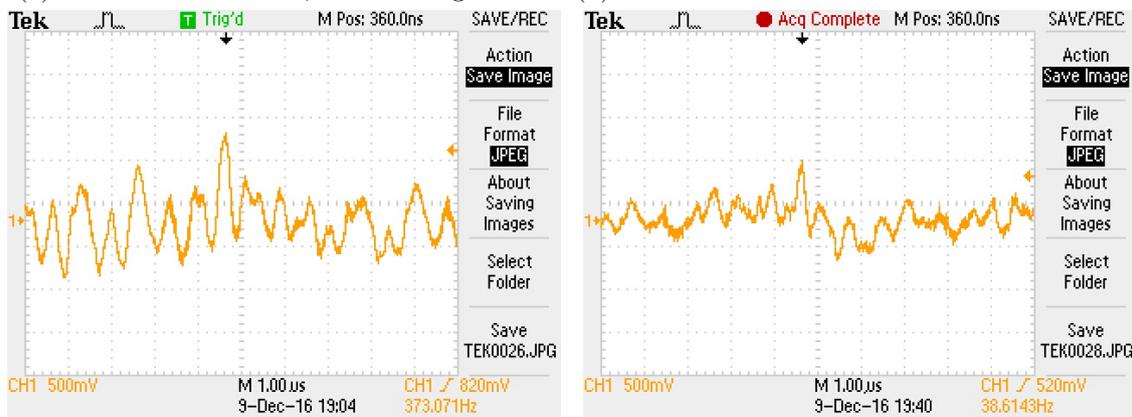


Figure 3.12 – Layer stack of a v3 GEB, showing the order the layers as well as the thickness and primary material of each. Note that for the dielectrics, the FR-4 material used is bromine-free in order to prevent outgassing. [40]



(a) Noise of the detector, initial configuration (b) Noise of the detector with the GEB connected



(c) Noise of the detector with the 24 VFATs and optohybrid connected (d) Noise of the detector with the LV cable shield connected to the GEB

Figure 3.13 – Noise measurements of v2 short detector GE1/1-VII-S-CERN-0004 in different configurations [42]

For this study, the v2 short detector GE1/1-VII-S-CERN-0004 was instrumented, and read out via the G3_bottom pad through the HV decoupling circuit on the detector, then to an ORTEC 142PC pre-amplifier, ORTEC 474 amplifier (Set to a coarse gain of 4, a fine gain of 4.5, and a differential time of 100 ns) and shaper, and finally to an oscilloscope. Initially, the star point was taken to be the SHV connector on the patch panel of the detector, to which the readout and HV return were connected. An earth ground connection was provided to the patch panel, and the pre-amplifier shield cable was connected to the readout PCB. The noise measurement given by this configuration is shown in Figure 3.13a, and has a peak-to-peak voltage of ~ 0.6 V.

The GEB was then placed onto the detector and screwed into the detector flanges. The LV pads were then connected to the patch panel, and the LV return cable was connected to the star point. All vias on the GEB were insulated with Kapton tape. Another noise measurement was taken, the results of which can be seen in Figure 3.13b, with a peak-to-peak voltage of > 4.0 V.

The 24 VFATs were then connected to the detector and GEB, and the optohybrid was connected to the GEB as well, with the four LV wires connected to the optohybrid. Another noise measurement was taken, the results of which can be seen in Figure 3.13c, with a peak-to-peak voltage of ~ 1.6 V.

Finally, the LV cable shield was soldered to the LV return pad on the GEB, and another noise measurement was taken. Seen in Figure 3.13d, the peak-to-peak voltage is ~ 1.1 V.

From these measurements, we can deduce that the GEB strongly couples to the electrode of G3_bottom, resulting in it being a critical source of noise and lowered detector performance, and therefore a strong control of the current return paths is required. Coupled with the experience of the slice test, the additional two shield layers in the v3 GEB were deemed necessary.

3.2.3 VFATs

One of the most critical components of the front-end electronics, and to this dissertation in particular, is the Very Front-End ATLAS and Totem (VFAT) silicon ASIC. The VFAT is a digital on/off chip with 128 channels, each of which has an adjustable threshold and, in the case of VFATs used in gaseous detectors, an internal protection diode network for each channel as well. Originally designed to read out the Roman Pot silicon detectors of the Totem experiment, the VFAT has seen use in many other experiments in the past decade, including CMS GEMs and TOTEM CSCs.

In the CMS GEMs, two iterations of the VFAT chip have been used. The VFAT2 in v2 electronics, and the VFAT3 chip in v3 electronics. The differences and similarities of the two iterations will be described in the following sections.

3.2.3.1 v2

When designing the VFAT2, the basic requirement of the project were as follows:

- 128 channel low noise front-end chip for the binary readout of capacitive sensors
- 40 MHz signal sampling, resulting in zero dead time in LHC use
- Digital memory with a programmable L1A latency of up to 256 clock periods, and simultaneous storage of up to 128 triggered events

- Programmable “fast OR” trigger building outputs
- Internal calibration via internal test pulses of a programmable amplitude
- Full programming through an I2C interface
- Data packet output which includes headers, counters, flags, and cyclic redundancy check (CRC)
- Radiation-tolerant design which is suitable for use in harsh radiation environments with respect to both the capability to withstand ionizing radiation and low single event upset rates

A photo of the VFAT2 chip, as installed onto a PCB, can be seen in Figure 3.14. The chip’s corresponding block diagram can be seen in Figure 3.15. An alternative block diagram, which better illustrates the signal flow within the VFAT2 and makes the distinction between the analog/asynchronous and digital/synchronous sections can be seen in Figure 3.15.

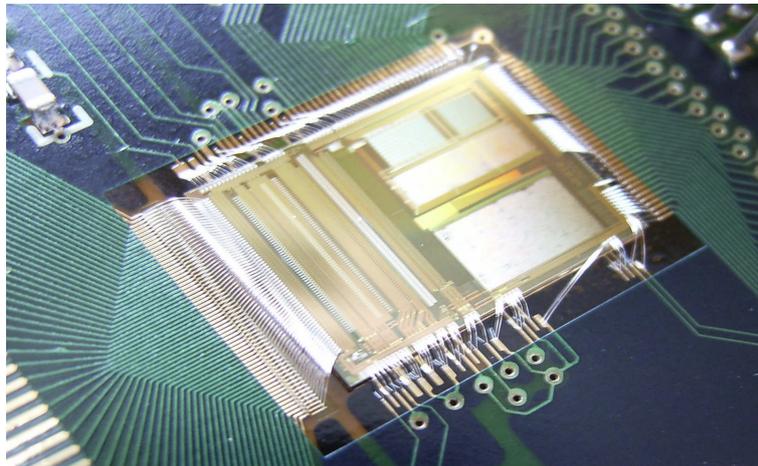


Figure 3.14 – Photo of a VFAT2 chip as installed onto a PCB [43]

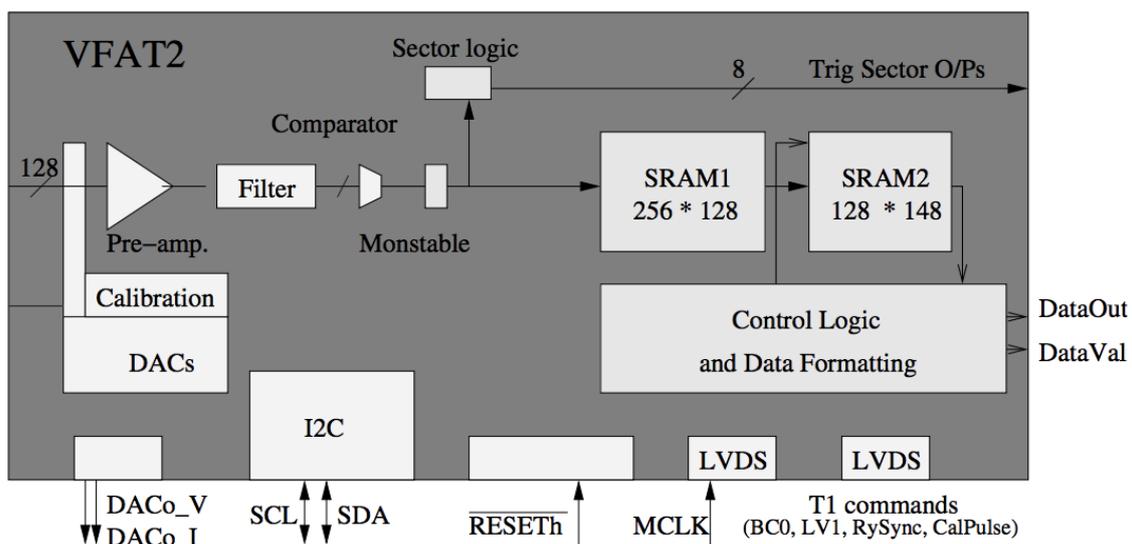


Figure 3.15 – Block diagram of a VFAT2 chip [44]

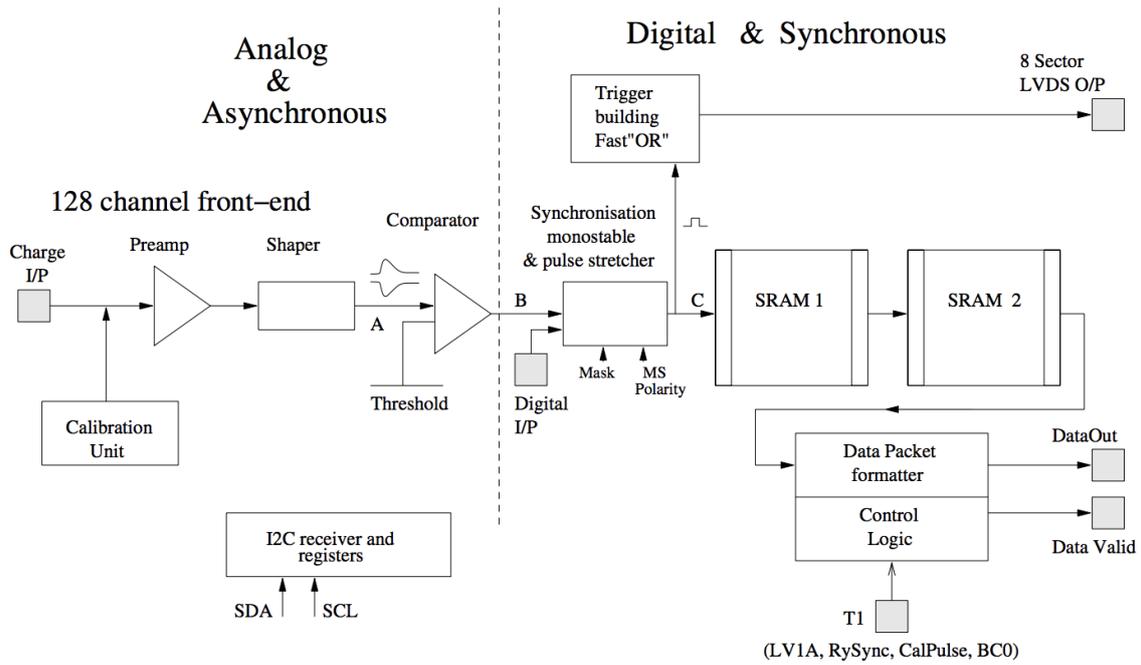


Figure 3.16 – Block diagram illustrating the data flow for a VFAT2 chip [45]

Each VFAT2 chip takes the charge which is deposited onto each of the corresponding 128 readout strips of the readout board and converts it to a binary value using a preamplifier, shaper, and comparator. If a particular channel receives a signal which is greater than the programmed threshold of the comparator, a logical 1 is produced; otherwise, a logical 0 is produced. This threshold is programmed via the ThReg DAC, and is common across all channels in the VFAT. However, as each channel may have a slightly modified gain, either by overall higher radiation exposure or by varying levels of input capacitance per channel as a result of production fluctuations, a fine adjustment of the threshold on a per-channel basis is possible via a 5-bit TrimDAC. This will be discussed in greater detail in Section 3.4.1.1.

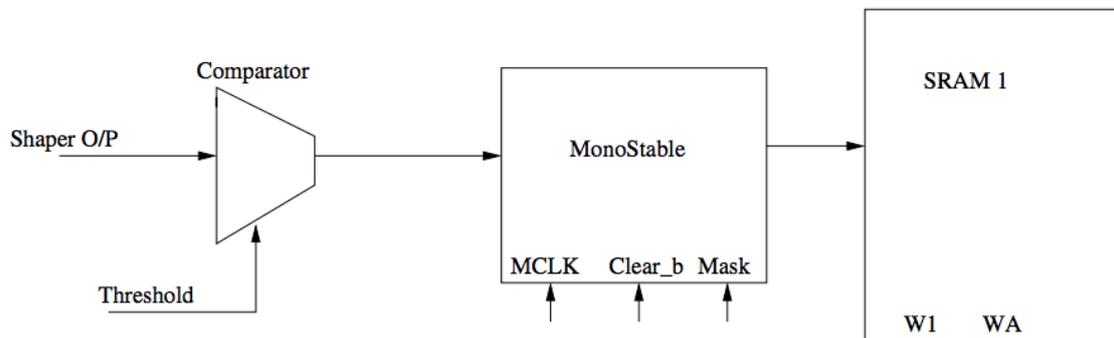


Figure 3.17 – Location of the monostable block within the block diagram of a VFAT2 chip [44]

The comparator is an asynchronous comparator without hysteresis, so when a signal passes the programmed threshold, the output of the comparator goes high and returns low again as it descends back through the threshold. For very large signals, this output may remain high for more than one clock cycle. On the other hand, for very small signals which

barely pass the programmed threshold, the output may be high for less than a single clock cycle. This is where the monostable block, shown in Figure 3.17, comes in.

The monostable block, in its default state, takes the pulse and changes its width to equal a single clock cycle, either by shrinking the pulse in the case of a large signal, or stretching it in the case of a small signal, as seen on the left and right of Figure 3.18, respectively. This one-clock-width pulse is the one which is sampled by SRAM1, not the original pulse.

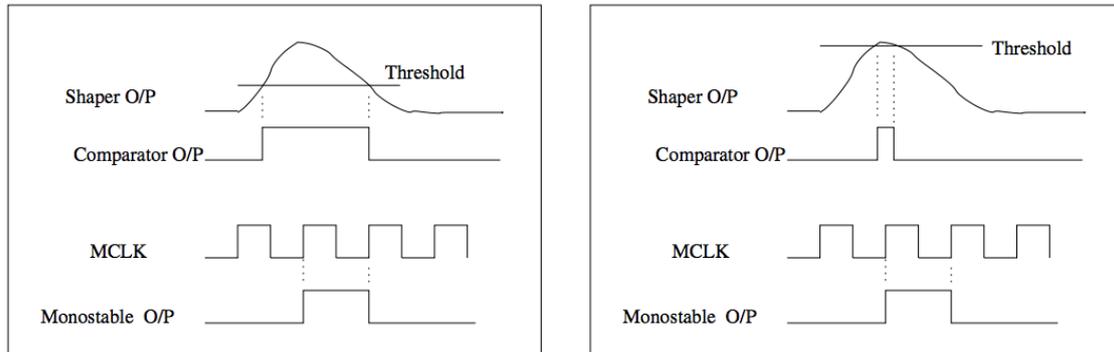


Figure 3.18 – Output of the monostable block for large (left) and small (right) signals from the comparator [44]

Although this pulse width of a single clock cycle is the default behavior, it is possible to change this behavior by choosing the monostable pulse length (MSPL) value. An MSPL value of 2, for instance, would stretch the pulse width to two clock cycles. As the MSPL is encoded in a three-bit register value, the maximum value that can be encoded is $MSPL = 8$, as shown in Table 3.1. In normal operations during the slice test, the chosen value was $MSPL = 4$.

MSPulseLength<2:0>	Function
0 0 0	The monostable pulse length = 1 clock period (default)
0 0 1	The monostable pulse length = 2 clock periods
0 1 0	The monostable pulse length = 3 clock periods
0 1 1	The monostable pulse length = 4 clock periods
1 0 0	The monostable pulse length = 5 clock periods
1 0 1	The monostable pulse length = 6 clock periods
1 1 0	The monostable pulse length = 7 clock periods
1 1 1	The monostable pulse length = 8 clock periods

Table 3.1 – Monostable pulse length register values [44]

Once the pulse width is set accordingly, the logical 1 is written into the first of two SRAM memories, called SRAM1. All channels which did not exceed the comparator's threshold have a logical 0 written to the same memory. This writing occurs in parallel at 40 MHz for all 128 channels.

The size of SRAM1 for VFAT2 is 256×128 . The latter ensures there is space for all 128 channels, whereas the former is the result of the L1A latency not being expected to exceed 256 clock cycles, or $6.4 \mu s$. More information about this latency can be found in

Section 3.4.2².

At the same time the data is being written to SRAM1, a fast OR function can be used to set a flag which can be immediately used for creating a trigger. This is shown in Figure 3.15 as “Sector logic”, which leads to eight trigger sector outputs (as, at the time of the creation of VFAT2, it was foreseen to have up to eight programmable sectors which can be flagged with the fast OR in this way), and in Figure 3.16 as the “Trigger building FAST OR” block.

When an L1A signal is received, the data which corresponds to the triggered timeslot is transferred to the second SRAM memory, called SRAM2. Unlike SRAM1, SRAM2 contains only triggered data, and is dimensioned 128×148 for the data plus headers. Thanks to these dimensions, the VFAT2 can store up to 128 triggered events of data for all 128 channels at any given time.

The headers mentioned previously come in three flavors: the 12-bit bunch crossing number (BCN), the 8-bit event number (EN), and the 16-bit chip identification number (ID). The BCN increments once every clock cycle, resetting to zero upon receiving a bunch crossing zero identifier (BC0) T1 command via low voltage differential signaling (LVDS). The EN increments once for every L1 trigger, and is similarly reset by a BC0 command or a clear signal. Both the BCN and EN counters are cyclic, and will therefore also return to zero at the end of the counter range. The chip ID, on the other hand, is always the same and is unique per VFAT.

Once the data is written to the SRAM2, the read cycle of the VFAT2 chip begins. Here, a data formatting block streams out a binary data stream to the back-end electronics, which we will describe in Section 3.3. This read/write process is a continuous operation, without associated dead time. More detailed information on the read cycle and how the VFAT2 chip works in general can be found in Reference [44].

In addition to the main signal path from readout strip to back-end electronics, the VFAT2 also contains a calibration unit in front of the preamplifier of each channel, which is indispensable for commissioning operations. This circuit is shown in Figure 3.19.

Here, the calibration circuit generates a fast voltage step by switching between two voltages - a fixed baseline of 1.067 V, and a variable voltage (between 877 mV - 1.074 V) called VCal. For a given VCal setting, the amplitude of the voltage step is simply the difference between V_{high} and V_{low} , as shown in the boxed portion of Figure 3.19. This voltage step is then applied to the input of its corresponding channel via a 100 fF capacitor.

The amount of injected charge can then be calculated by multiplying the amplitude of the voltage step by the 100 fF capacitance via Equation 3.1.

$$Q = CV \tag{3.1}$$

In addition to the amplitude of the voltage step, the polarity and phase are also programmable by the user. The injected charge takes the form of Figure 3.20 for positive signal, and takes the form of the inverse of Figure 3.20 for negative signals, always passing through zero in either case.

²Largely because the author is writing this dissertation completely out of order and it gets rather time-consuming having to constantly shift sections around.

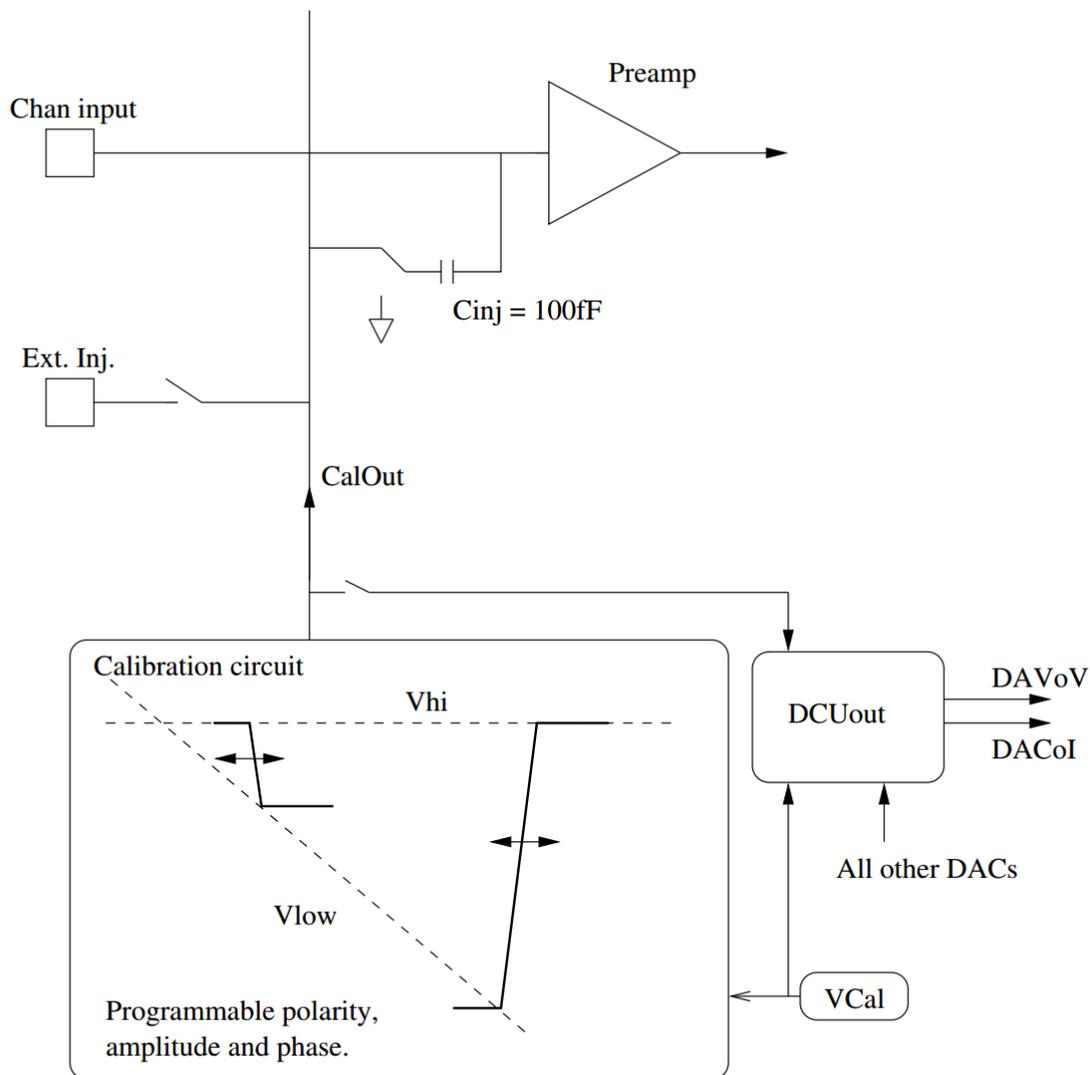


Figure 3.19 – Calibration circuit of the VFAT2 chip

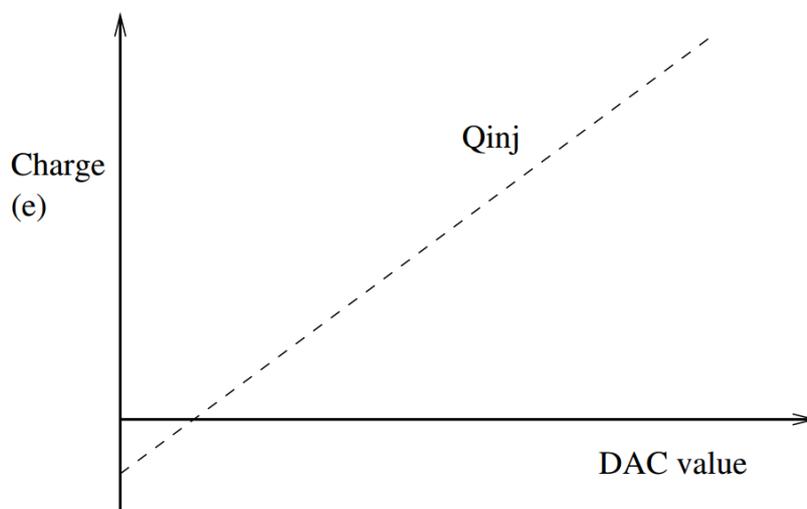


Figure 3.20 – Injected charge

In addition to this internal calibration circuit, it is also possible to apply an external

voltage pulse to a given channel for laboratory test purposes. However, this method of calibration is obviously impractical for experimental operations and thus is not described further.

Of course, this all merely describes the VFAT2 ASIC itself. However, when discussing the “VFATs” of a detector, more than likely the speaker is referring to the VFAT hybrid rather than simply the ASIC. In order to be integrated into the detector’s functionality, the ASIC is installed onto a PCB, shown in Figure 3.21.

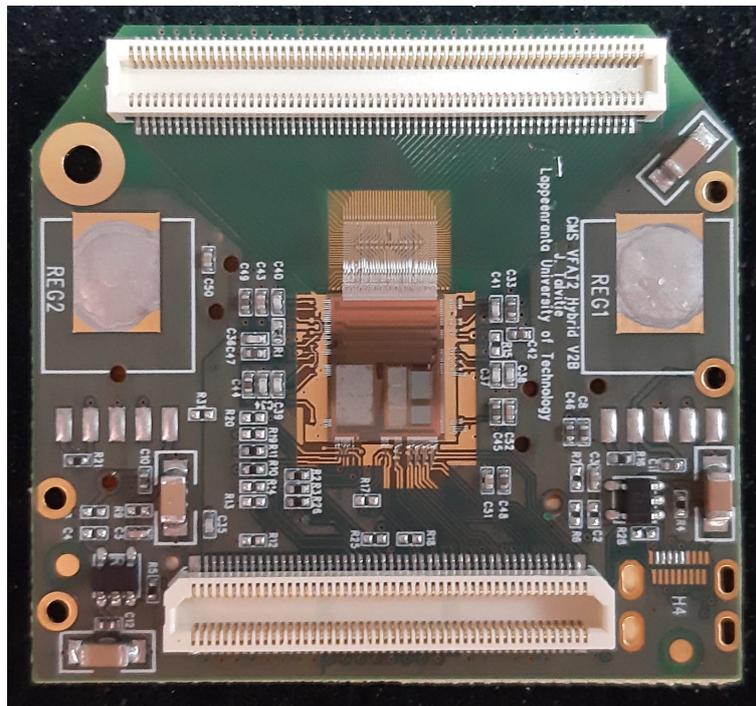


Figure 3.21 – Photo of a VFAT2 ASIC installed on a v2b hybrid



Figure 3.22 – Photo of the protective measures installed over a VFAT2 ASIC installed on a v2 hybrid. On the left, the hollow frame that surrounds the VFAT2 chip, and on the right, a solid cover is added on top of the frame to protect the bonding wires and chip from accidental damage

Here, there are two Panasonic connectors. The larger one on top connects to the readout board of the GEM detector, and is responsible for the analog signals received from the readout strips. This 128-pin connector has one input for each channel in that particular readout sector, and transmits the data to the VFAT2 chip via copper traces connected to the clear bonding wires seen below the connector. These bonding wires are extremely fragile, and therefore great care must be taken to not damage them when handling the hybrid. In order to give them some additional measure of protection, a square covering of high-density plastic was added to the hybrid, shown in Figure 3.22.

The smaller, 100-pin Panasonic connector on the bottom connects the hybrid to the GEB, and is responsible for the digital signals received from the VFAT2 chip after the processing described earlier in this section. It is also responsible for the powering of the VFAT2.

3.2.3.2 v3

When designing the VFAT3, the basic requirements of the project were as follows [46]:

- 128 channel low noise front-end chip for the binary readout of either positive or negative charges from capacitive sensors
- Trigger information at a minimum fixed latency with a granularity of two channels
- Full-granularity tracking data after L1A
- An L1A latency of up to $12.5 \mu\text{s}$
- A time resolution of $< 7.5 \text{ ns}$ with the detector included
- Integrated calibration and monitoring functions
- Interfacing to and from the GBTs at 320 Mbps
- Radiation resistance up to 100 MRad
- Robustness against single event upset (SEU) effects

Although the main purpose of the VFAT3 did not change, the above list includes a significant amount of improvements over the VFAT2. Some key improvements are listed in Table 3.2.

	VFAT2	VFAT3
Technology	250 nm CMOS	TSMC 130 nm CMOS
Signal frequency	40 MHz (LVDS)	320 MHz (SLVS)
L1A latency	3.2 μ s (Maximum programmable = 6.4 μ s)	12.5 μ s (Maximum programmable = 25.6 μ s)
Trigger granularity	16 strips	2 strips
Per-chip tracking data and slow control bandwidth	80 Mbps	320 Mbps
Shaping time	Fixed at 22 ns	Programmable to 25 ns, 50 ns, 75 ns or 100 ns
Gain	Fixed	Programmable with dynamic ranges DR10, DR30 or DR60
Internal calibration	Voltage step only	Voltage step and current pulse injection
Temperature monitoring	No	Yes
Calibration/Monitoring ADC	External	Internal / Integrated and External

Table 3.2 – Key differences between VFAT2 and VFAT3

The block diagram for the VFAT3 chip can be seen in Figure 3.23.

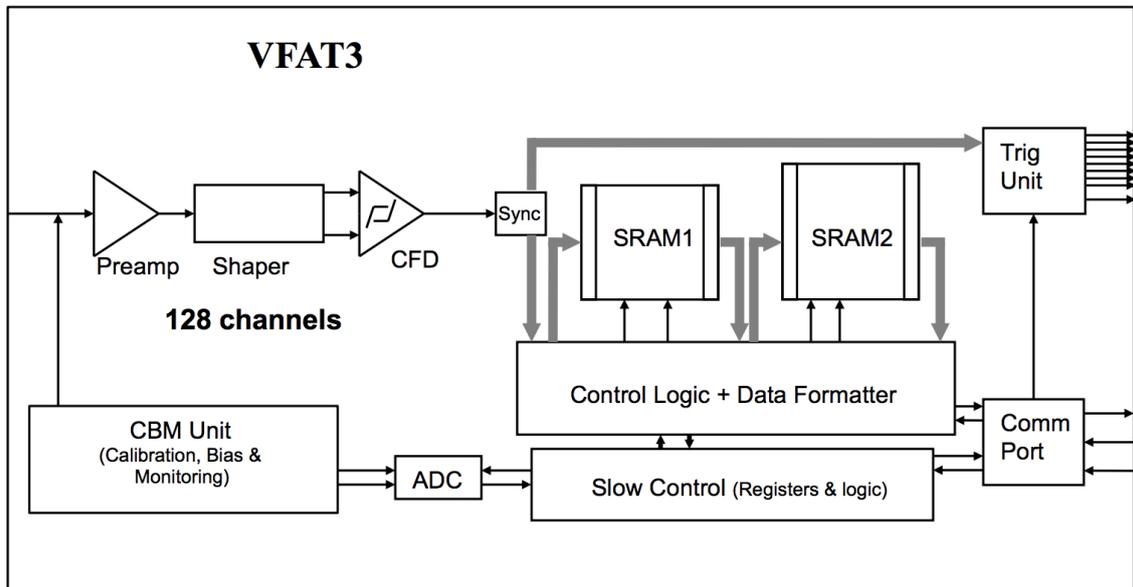


Figure 3.23 – Block diagram of a VFAT3 chip [46]

As with VFAT2, each VFAT3 chip takes the charge which is deposited onto each of the corresponding 128 readout strips of the readout board and converts it to a binary value using a preamplifier and shaper. Here, however, is where things start to change. Rather than being sent to a simple comparator, the signal then moves to a constant fraction discriminator (CFD), the block diagram for which can be seen in Figure 3.24. In order to provide better noise rejection, the implementation of the CFD in VFAT3 is based on a fully-differential architecture.

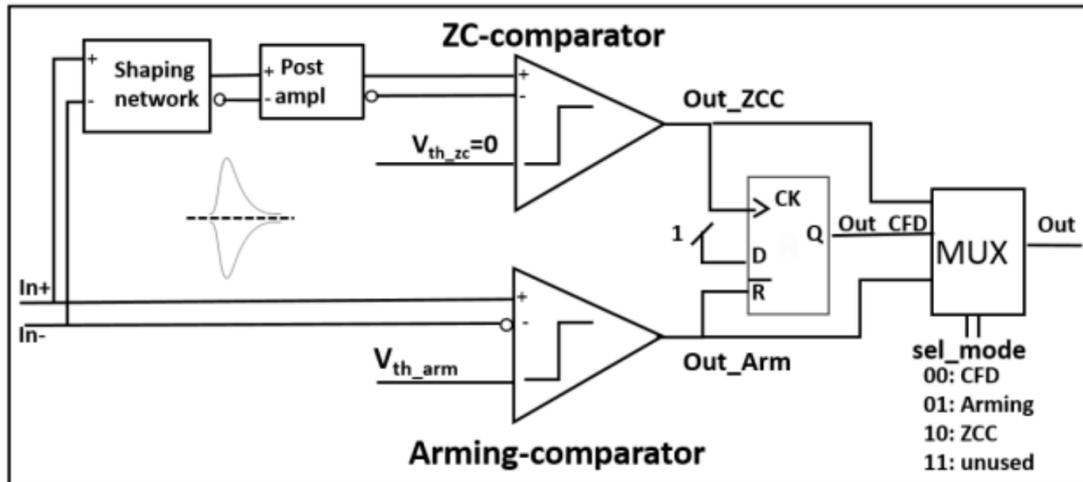


Figure 3.24 – Block diagram of the CFD of a VFAT3 chip [46]

The differential input signals are first sent to a fully-differential cross-coupling passive shaping network, seen on the left of Figure 3.24 and shown in more detail in Figure 3.25. This shaper is based on a four-stage RC filter with programmable time constants of 25, 50, 75, or 100 ns, with 100 ns being the default. This represents a vast improvement from VFAT2, which had a fixed shaping time of 22 ns, leading to a ballistic deficit as a result of the shaping time being shorter in duration than that of the signals themselves. By allowing for programmable time constants up to 100 ns, the VFAT3 is capable of integrating the entire signal charge, thus providing a higher signal-to-noise ratio than VFAT2, as will be seen in Section 3.4.1.

This shaper combines a delayed copy of the signal with its fraction to produce differential bipolar pulses with amplitude-independent zero-crossing time.

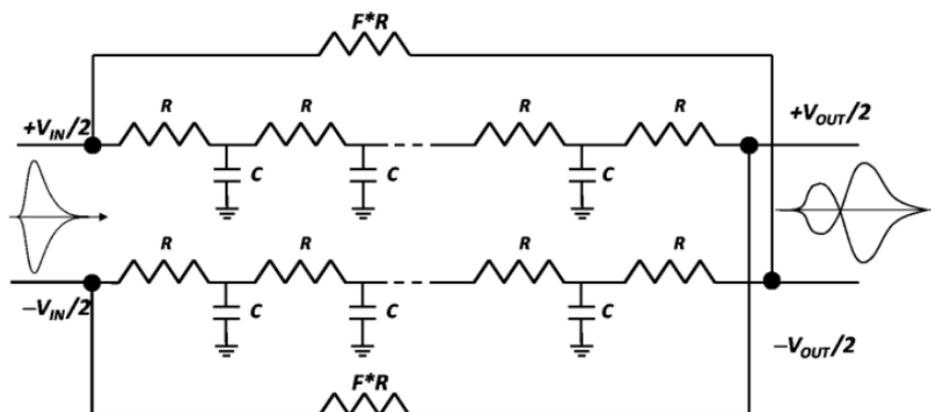


Figure 3.25 – Circuit diagram of the shaping network of a VFAT3 chip [46]

These differential bipolar pulses are then amplified by the post amplifier (“post ampl”) which recovers the signal attenuation which was introduced by the shaping network and also applies a dynamic offset compensation. Finally, the pulses are sent to the zero-crossing comparator (ZCC), which produces a digital pulse whenever its differential input crosses its baseline, much in the same way the simple comparator operated for VFAT2. The threshold for the ZCC is provided by a global 8-bit digital-to-analog converter (DAC) called the ZCC DAC, which has a nominal value of $0.5 \frac{mV}{bit}$ and a range of 0x00 to 0xFF.

In parallel to this, the original input signals are sent to an arming comparator, in order to enable the CFD output only when the input signal is larger than the programmed threshold. This threshold is also provided by a global 8-bit DAC called the ARM DAC, which has a nominal value of $2.0 \frac{mV}{bit}$ and a range of 0x00 to 0xFF. Unlike the simple comparator of the VFAT2, the arming comparator of the VFAT3 has a programmable hysteresis which is provided by an 8-bit DAC with a nominal value of $0.4 \frac{mV}{bit}$.

Note that the ZCC and arming comparator are both based on the same high gain differential amplifier followed by a chain of digital inverters. However, they differ in the way the corresponding threshold is applied. In theory, in order to detect the zero-crossing time, the ZCC should not have any threshold applied. However, as noise is an unavoidable fact of electronic life, it must have some small threshold applied in order to prevent it from firing solely on noise.

The outputs from these two streams are then sent to a multiplexer (MUX), which allows for the selection of the comparator output according to the selection of the CFD output mode. The options for this mode are shown in Table 3.3. For example, if one set the SEL_COMP_MODE register to 0 1, it would allow the user to bypass the full CFD and use the arming comparator output without the time-walk correction.

SEL_COMP_MODE<1:0>	Function
0 0	CFD (default)
0 1	Arming comparator only (leading edge)
1 0	Zero crossing comparator only
1 1	0

Table 3.3 – Programmable options for the comparator modes [46]

Moving on from the CFD, whereas VFAT2 used a monostable block to alter the pulse width of the signal received from the comparator, in VFAT3 the interface between the analog and digital portions of the chip are controlled by a data synchronization and pulse stretching block, which is labeled as “Sync” in Figure 3.23.

Here, for the data synchronization portion, the signal output from the CFD is sampled and synchronized to an internal 40 MHz clock. For the pulse stretching portion, a pulse stretching logic works to stretch the synchronized pulse from one to eight clock periods. Note that this stretching is only valid in the case of the variable latency path; for the fixed latency path, the synchronized pulse will always be the width of one clock cycle.

Note that this Sync block can operate in one of two modes. In “edge” mode, the input signal from the CFD is sampled on the rising edge of the internal clock, and the output is a synchronized digital pulse whose pulse length is equal to the number of clocks set for the pulse stretching (PS) and completely independent of the CFD pulse length.

On the other hand, in “level” mode, the level of the CFD pulse is sampled every clock cycle and the PS pulse length is added to that length, making the pulse length of the output dependent on the CFD pulse length. For example, if the CFD output is N clock cycles long, the output from the Sync block will be N+PS+1 clock cycles.

The output from the MUX then, as with the VFAT2, splits off to the trigger unit in one direction, and the SRAM in the other direction.

In regards to the trigger granularity of 2 strips shown in Table 3.2, each VFAT3 has eight SLVS pairs which run at 320 MHz, in comparison to the 40 MHz of VFAT2. These SLVS pairs can be configured to transmit in either single data rate (SDR) or double data

rate (DDR) mode, where DDR mode uses both edges of 320 MHz clock to transmit data.

In SDR mode, you have “fast OR” information, equal to 8×8 bits per bunch crossing, totaling 64 bits per bunch crossing for all 128 channels. This gives a granularity of two channels, as shown in Table 3.2, and is transmitted directly to the trigger unit in the same way as for VFAT2.

In DDR mode, on the other hand, you have full channel information, equal to 8×16 bits per bunch crossing, totaling 128 bits per bunch crossing for all 128 channels. This gives a granularity of one channel.

The variable latency path, shown in Figure 3.26, is used to transmit the full granularity information via the communication port (comm port), accompanied by a timestamp which identifies the bunch crossing which is associated with that data. This fulfills the project requirement of full-granularity tracking data after L1A and allowing for improved offline muon reconstruction.

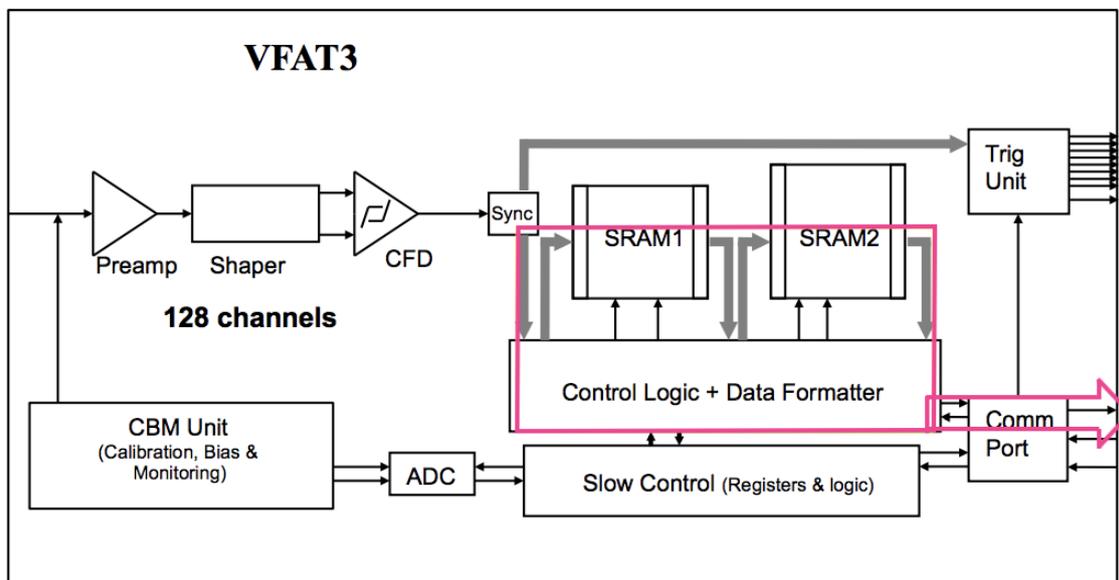


Figure 3.26 – Block diagram of a VFAT3 chip with the variable latency path highlighted in pink

As with VFAT2, VFAT3 has two SRAM memories, SRAM1 and SRAM2. The SRAM1 is a circular buffer which continuously samples all 128 channels every clock cycle, and is 128 bits wide \times 1024 bits deep. This allows it to store data for all 128 channels for up to 1024 continuous bunch crossings before the buffer loops around and overwrites previously-written data.

And as with VFAT2, upon receiving the L1A signal, the corresponding data from SRAM1 is transferred to SRAM2 and has a time stamp added to it to allow for bunch crossing identification later. The SRAM2 then acts as a FIFO, sequentially popping out each block of data and feeding it to the data formatter for the construction of the data packet which is then transferred off-chip. The variable latency path comes from the fact that these data packets are queued within the SRAM2 buffer. This buffer is of dimensionality 176×512 , where the depth of 512 is equal to the maximum number of data time slots which can be simultaneously queued for transmission off-chip. If this maximum value is exceeded, writing to SRAM2 is disabled and any further triggers are lost, making all the physicists very sad. However, in this occurrence, the event counter will continue to incre-

ment with each L1 trigger, so that when the queue is reduced back below the maximum number, the newly-written data will have the correct associated time tags, which makes the physicists happy.

One additional component of the VFAT3 that has yet to be mentioned is at the bottom of the block diagram of Figure 3.23, the slow control. The VFAT3 was designed to be an incredibly versatile and programmable chip, and that comes from the slow control, which allows for the reading and writing of internal registers for monitoring and programming, respectively, at a maximum rate of 40 Mbps.

As is standard for CMS upgrades, the slow control protocol utilizes the IP-bus protocol wrapped within the high-level data link control (HDLC) protocol, which insures the correct chip addressing and slow control packet error checking. As the VFAT3 uses the communication port for all data communication, including the slow control, control characters (CC) are used to distinguish between slow control commands/data and all other commands and data. There are two such slow control CCs, for the binary values of slow control 1 and 0. The reception and transmission of these slow control data and commands must take a lower priority than that of real data transmission, and as such, it is possible (and, indeed, required) to stop the slow control communication mid-flow and resume it again when the communications port is free.

To the left of the slow control block in Figure 3.23 is an analog-to-digital converter (ADC) and the calibration, bias, and monitoring (CBM) unit. This unit provides the analog front-end and CFD biasing in addition to the calibration functions, and is controlled by the slow control.

The required currents and voltages for the analog front-end biasing is programmable via nine 6-bit DACs and eight 8-bit DACs, in addition to 129 local channel DACs for the local threshold adjustment of each channel (128 normal channels plus one test channel). This choice was made as the analog characteristics of the chip are expected to change over time due to both radiation damage (despite the significant radiation tolerance of the chip) and normal process variations, such that re-biasing may be required over the lifetime of the detector. The properties of these DACs can be found in Tables 9, 10, and 11 of Reference [46].

Just as these internal DACs can be programmed, they can also be read back out in order to monitor and precisely measure the calibration pulse charge, current bias, and voltage bias, as indicated by the bi-directional arrows in Figure 3.23. This monitoring system utilizes a multiplexing system to select which DAC, and either current or voltage, to monitor.

The measurement of a voltage can be done directly through the ADC which is integrated into the chip. In this way, the VFAT3 differs from the VFAT2, as the VFAT2 utilized an external ADC for voltage measurements of this kind. This internal ADC, like the CBM unit, is controlled and read out by the slow control system.

The measurement of a current, on the other hand, is slightly more complex, as it cannot be done directly through the ADC, which is only capable of reading voltages. Rather, the DAC currents are first routed to an external precision resistor and are converted into voltages before being read by the ADC.

The other half of the CBM unit is, of course, the calibration circuit. This circuit is capable of delivering internal pulses to each individual channel of the VFAT, which, in addition to the obvious calibration purposes, serves as an incredibly useful feature for the characterization of the chip during production testing, as will be discussed in much greater

detail in Section 3.4.1. This calibration circuit can be seen in Figure 3.27.

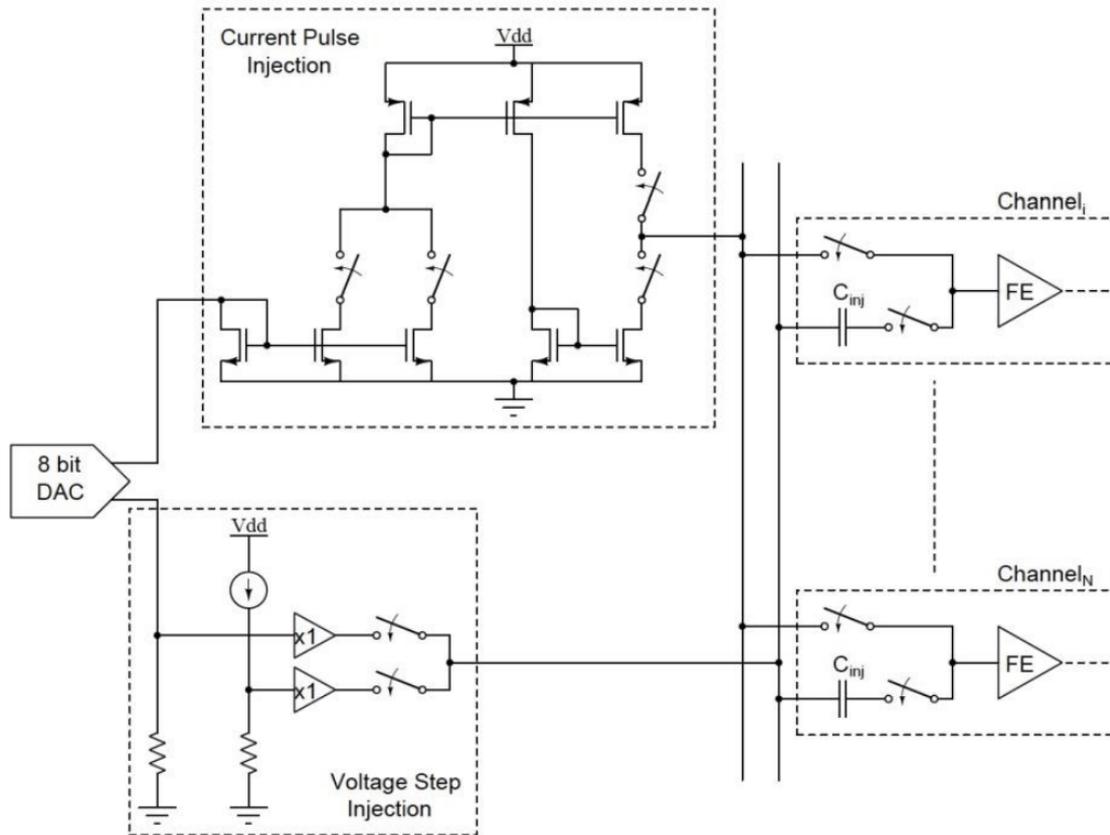


Figure 3.27 – Circuit diagram of the calibration circuit of a VFAT3 chip, including their connections to the channel preamplifiers via switches [46]

Here, we see that the circuit is broken into two main blocks - one for current pulse injection, and another for voltage step injection. These represent the two different types of injection pulse from which the user is able to choose.

The voltage step pulse, from the circuit on the bottom left of Figure 3.27, is applied to a series capacitor, which generates a delta-like pulse similar to that of a silicon detector at the input of the preamplifier. The polarity, amplitude, and phase of this delta-like pulse are all programmable by the user.

In order to produce the voltage step, two DC voltages are used, as shown in Figure 3.28: V_{high} and V_{low} , where V_{high} is programmable via an 8-bit DAC and V_{low} is a constant voltage. Depending on the polarity chosen by the user, one of the two voltages is used to pre-charge the capacitor, and then the other voltage is switched to. The voltage step is then given by the difference between V_{high} and V_{low} , as shown in Figure 3.29.

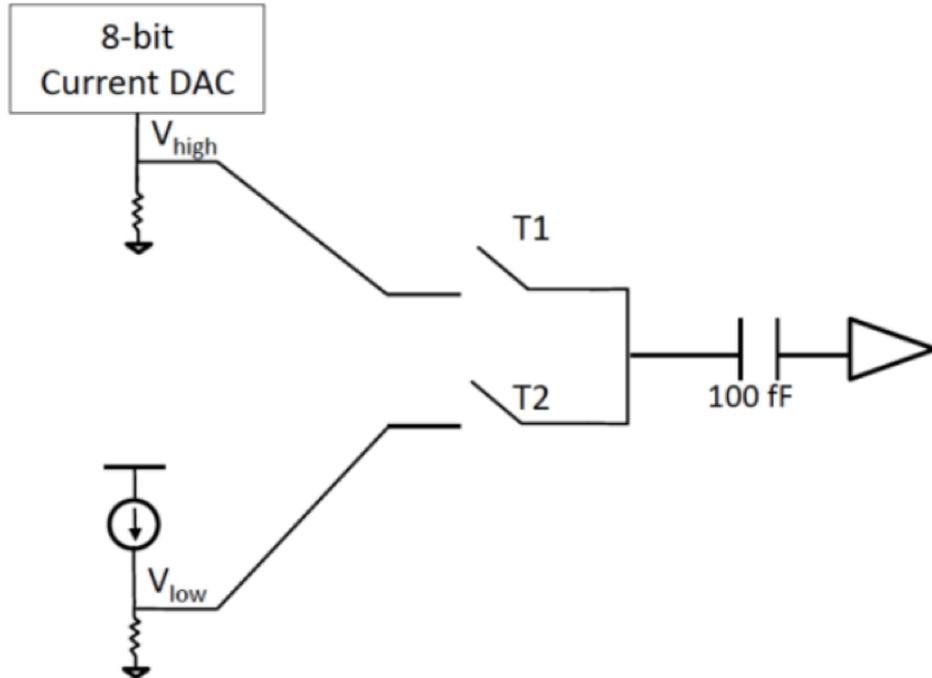


Figure 3.28 – Circuit diagram of the voltage step calibration circuit of a VFAT3 chip, showing more clearly the voltages used in creating the step and the principles of operation [46]

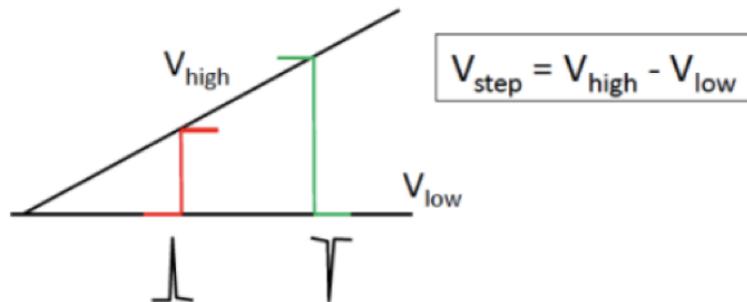


Figure 3.29 – The voltage step as created by V_{high} and V_{low} in the voltage step injection calibration circuit [46]

The monitoring system of the VFAT3 allows the user to scan the CAL DAC to measure V_{high} , V_{low} , and the difference between them as DC voltages. This difference is then multiplied by the 100 fF capacitor in order to get the amount of injected charge, via Equation 3.1, shown again below.

$$Q = CV \quad (3.2)$$

It is important to note that while it is possible to inject charge in this manner onto multiple channels at once, this method of injecting charge relies on charge sharing, and so doing so would result in the possible inaccuracy of the charge value. Thus, when using this calibration circuit, it is highly recommended to proceed channel-by-channel in series.

The current pulse injection, on the other hand, is applied directly to the input of preamplifier, and represents a GEM-like pulse. In this case, only the magnitude and pulse duration of the current pulse are programmable by the user.

To create the pulse, a DC current is switched between a current sink and the preamplifier, as shown in Figure 3.30. Here, a slow buffer, highlighted with the dashed circle, provides a reference voltage V_{in} to the node labeled “A” in order to decrease large transients. The width of the current pulse can be chosen by the user with the CAL DUR register as being 1 - 4 clock periods (25 ns, 50 ns, 75 ns, or 100 ns). These pulse widths then correspond with the resulting full scale current, where 25 ns \rightarrow 2 μ A, 50 ns \rightarrow 1 μ A, 75 ns \rightarrow 750 nA, and 100 ns \rightarrow 500 nA.

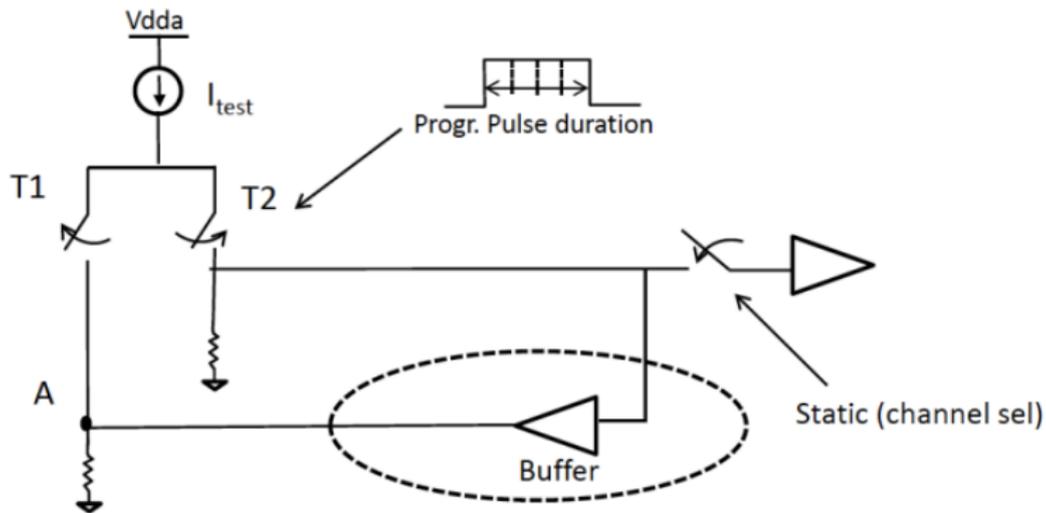


Figure 3.30 – Circuit diagram of the current pulse calibration circuit of a VFAT3 chip detailing the principle of operation [46]

The use of these calibration circuits will be further explained in Section 3.4.

One other important monitoring function that is present in VFAT3 but was not present in VFAT2 is the temperature sensor, whose circuit diagram is shown in Figure 3.31. This sensor has a temperature range of -100°C to $+100^{\circ}\text{C}$, with a sensitivity of $3.78 \frac{\text{mV}}{\text{C}}$, and allows the user to better monitor and care for the health of the electronics by ensuring they are not overheating and that the cooling circuit of the detector is performing optimally.

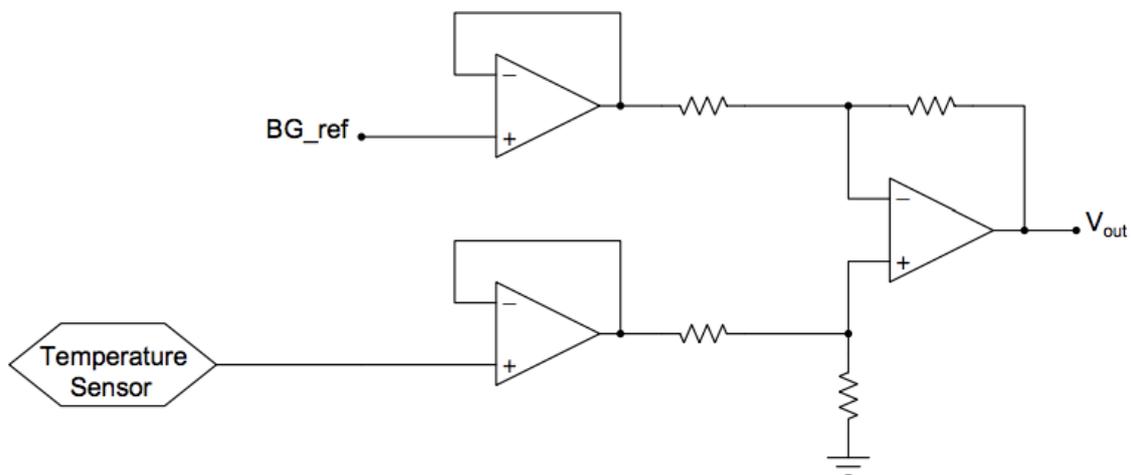


Figure 3.31 – Circuit diagram of the temperature sensor of the VFAT3 chip [46]

More detailed information on the processes described above can be found in Reference

[46].

As before, all of this merely describes the VFAT3 ASIC itself, and is an incomplete picture without the associated v3 hybrid. A photo of an HV3b_v3 hybrid can be seen in Figure 3.32, shown to the right of a v2 hybrid for comparison. Note that here, the VFAT3 chip that is installed on the HV3b_v3 VFAT3 hybrid is covered with a black protective “glob top”, which serves the same purpose as the hard plastic casing shown in Figure 3.22 for v2. Although the glob top offers a slimmer profile, a slightly smaller and more variable footprint, and is easier to install, it has the distinct disadvantage of being completely permanent. Any attempts to remove the glob top will damage both the VFAT3 ASIC and the bonding wires.

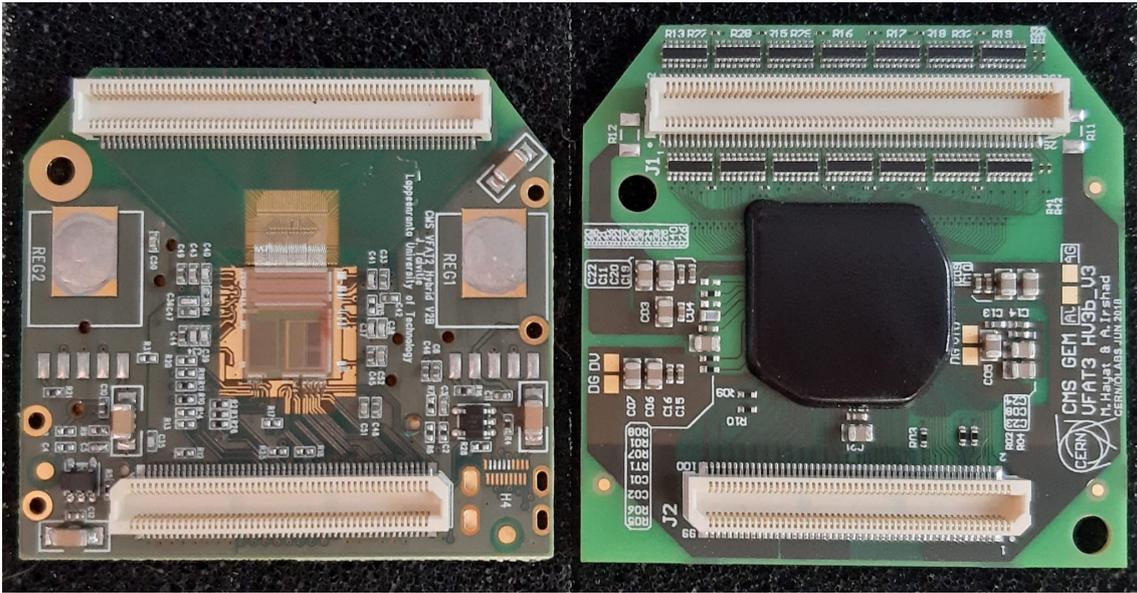


Figure 3.32 – Photo of a v2 VFAT2 hybrid (left) shown in comparison to a HV3b_v3 VFAT3 hybrid (right)

One of the most easily-observable difference between the v2 and HV3b_v3 VFAT3 hybrids is the difference in height between them, with the HV3b_v3 VFAT3 hybrid being 4 mm taller than the v2 hybrid. This difference in height is to accommodate the array of resistors seen to the top and bottom of the 128-pin Panasonic connector. These resistors are perhaps the most crucial addition to the v3 VFAT hybrids, and will be discussed in great detail in Section 7.2.

3.2.4 Optohybrid (OH)

The optohybrid is a large mezzanine board (typically on the order of 10.0 cm × 20.0 cm × 1.1 cm) which is mounted onto the GEB. Its main purpose is to synchronize the data sent by the VFAT chips, zero-suppress the trigger data, encode that data, and send it via optical links to the trigger electronics. Its main components include a field-programmable gate array (FPGA), one or more GBTx chipsets (see: Reference [47]), a slow control adapter (SCA) (see: References [48] and [49]), and versatile optical links (see: Reference [50]).

3.2.4.1 v2

The v2b version of the optohybrid can be see in Figure 3.33.

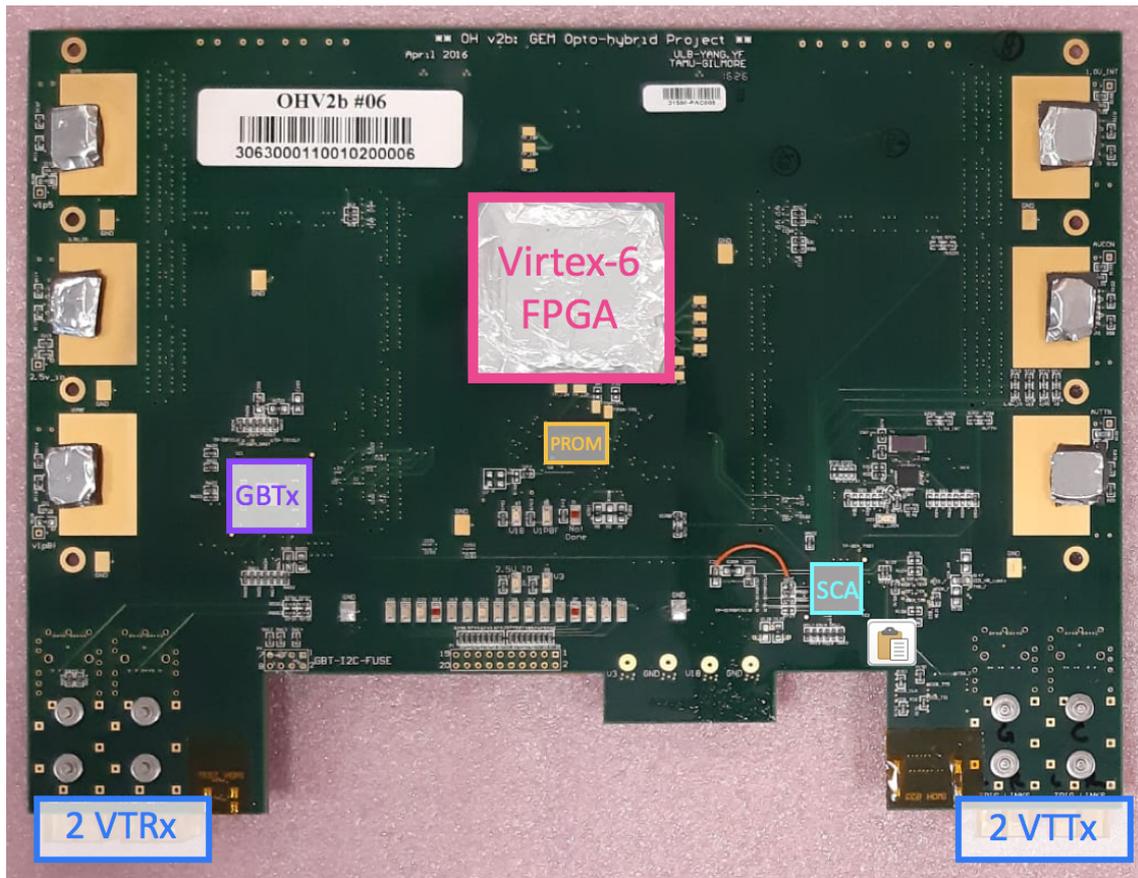


Figure 3.33 – Photo of a v2b optohybrid mezzanine card, with main features labeled

In the center of the optohybrid is a Virtex-6 FPGA, which has been shown to be radiation hard to levels at least two orders of magnitude higher than the expected radiation dosage of the GE1/1 detectors [51]. This FPGA acts as a data concentrator and is responsible for the main operations of the optohybrid board, including the handling of tracking data, trigger data, and slow control commands.

To the left of the FPGA is a GigaBit Transceiver (GBTx), which is a radiation-hard (up to 200 Mrad) optical data link technology which provides bidirectional 4.8 Gbps serial communication. The GBTx acts as a full link transceiver with bidirectional data communication between the front-end electronics and the counting room in USC55. Capable of recovering the frame clock, reducing the jitter from an input clock, and distributing phase-controlled clock signals, the GBTx delivers the global system clock reference from USC55 to all of the VFAT chips through a set of local electrical links (e-links).

This GBTx is connected to a slow control ASIC (SCA), which transmits the slow control commands to the FPGA and is connected to the joint test access group (JTAG) scheme which is integrated into the FPGA.

At the bottom of Figure 3.33, we see two sets of versatile optical links. On the left are two Versatile TransReceiver (VTRx) optical links, which are bidirectional links which run between the μ TCA crates located in USC55 and the optohybrid. In the direction of the detectors from USC55, this data path is used for sending setup and control signals to the front-end chips. In the opposite direction, this data path is used for the VFAT tracking and trigger data packets, as well as for returning slow control data. On the right are two Versatile Twin-Transmitter (VTTx) optical links, which are unidirectional optical links between the GEM optohybrids and the CSC trigger motherboard (TMB). This data

path is used to deliver the VFAT fixed-latency trigger data from the GEM system to the CSC system [52]. A photo of a VTRx module, and the differences between VTRx and VTTx, are shown in Figure 3.34.

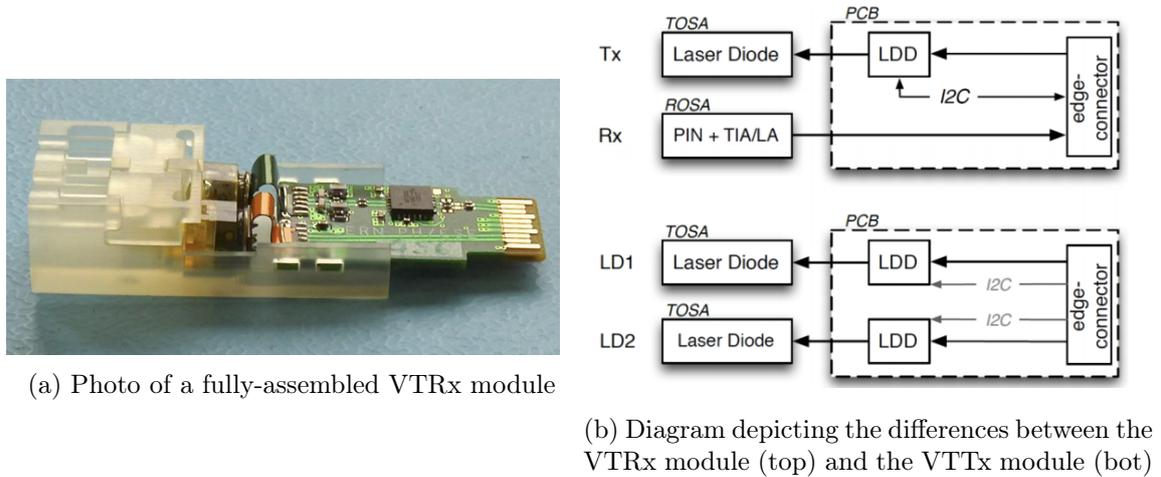


Figure 3.34 – Photo and diagrams of the versatile link optical modules [50]

3.2.4.2 v3

The v3c version of the optohybrid can be seen in Figure 3.35, with its main features labeled.

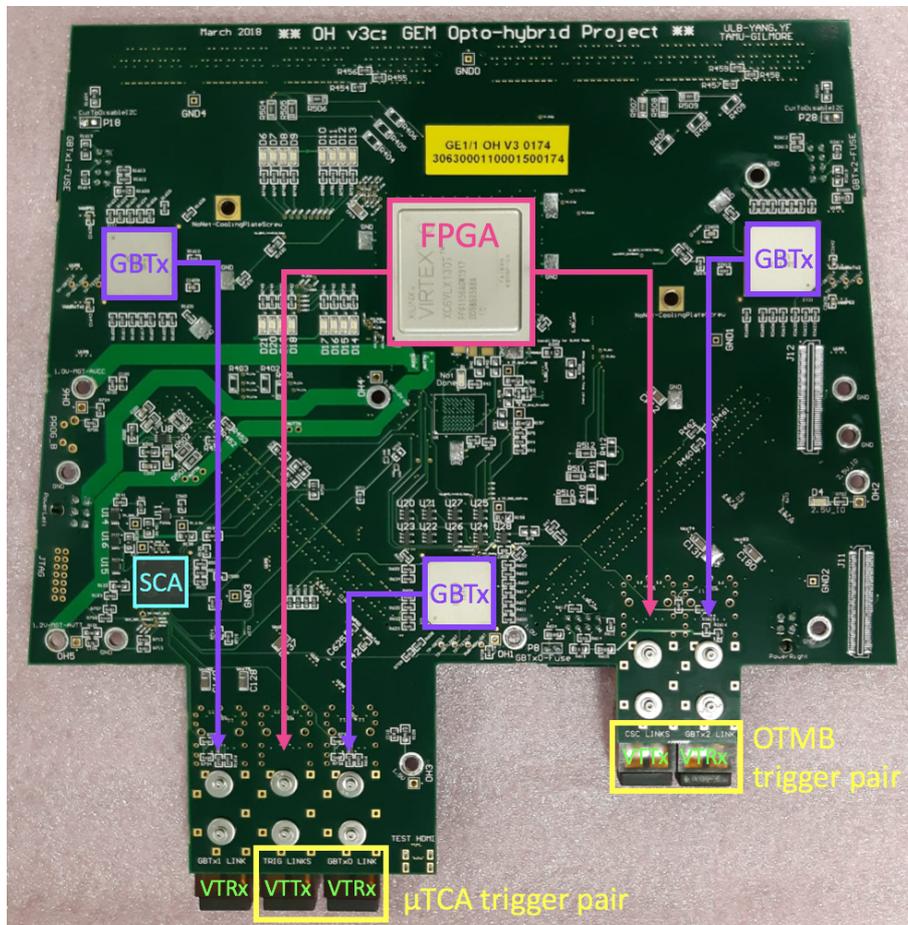


Figure 3.35 – Photo of a v3c optohybrid mezzanine card, with main features labeled

Although the v3 optohybrid features the same Virtex-6 FPGA as the v2 optohybrid, its purpose makes up one of the primary differences between the two versions. Whereas in v2 the FPGA was responsible for tracking, trigger, and slow control data, the v3 optohybrid offloads much of these functions to the GBTx chipsets. As a result, the v3 FPGA is only responsible for the handling of the trigger data. It does so by applying zero suppression algorithms, formatting the trigger data, and transmitting it through the two VTTx links located on the bottom of the optohybrid - one link to the GEM trigger system and one link to the CSC optical trigger motherboard (OTMB). The FPGA and this data transmission path is shown highlighted in pink in Figure 3.35.

In order to handle the increased load, the v3 optohybrid features three GBTx chipsets rather than just the one of the v2 optohybrid. Each GBTx is capable of handling up to 10 VFATs, though in practice each is only responsible for one column of eight VFATs. These are highlighted in purple in Figure 3.35. Each GBTx communicates with its corresponding VFAT chips through a set of local e-links, which are driven by e-ports that are integrated into the VFAT3 chips and run at 320 MHz using SLVS.

The GBTx chipsets are minimally fused, such that they lock to the fiber link, recover the clock, and keep a certain configuration after power loss or other issues. The presence of a GBTx watchdog prevents broken link and communication errors after a clock change. It is reset at power-on, controls the initialization procedure, and continuously monitors the GBTx chipsets, and in the case of a major link failure, it reinitializes the link automatically, without any needed input from the detector operator [53]. Although this watchdog was attempted to be implemented in v2 for the slice test, it was determined to be non-functional in-situ.

The reason for this failure was two-fold: For one, the power-on reset was removed too soon, prior to the GBTx power reaching a stable value. And for another, the decoupling capacitance of the GBTx was too small. Both of these issues were solved with the v3 optohybrid by adding additional capacitance, increasing the time constant of the RC circuit and increasing the decoupling capacitance from 10 μF to 47 μF .

Another change made to the v3 optohybrid is that the programmable read-only memory (PROM) was made redundant, thanks to the “spare” e-links left over on the GBTx chipsets. The PROM, being a commercial, off-the-shelf component, was not designed for high-radiation environments and thus posed a risk for radiation damage during operations, risking the lifetime of the readout electronics. As a result of this change, the FPGA can be programmed directly from the CTP7 card in under 70 ms. This is achieved through the SCA, highlighted in teal in Figure 3.35. This SCA, which is connected to GBTx1, also allows for resets to be sent to the VFAT3 chips and other front-end electronics, and allows for temperatures and voltages to be read out and monitored.

All of the communication with the optohybrid is performed through the versatile optical links seen highlighted in green at the bottom of Figure 3.35. Each of the VTRx links is connected to one of the three GBTx chipsets and allows for the communication of tracking data and slow control between the VFATs and the GBTx chipsets. The rightmost VTRx link, in particular, also transmits the FPGA programming information for the SCA, which is connected to GBTx1. The VTTx links, on the other hand, are connected to the FPGA, and allow for the communication of trigger data. The left VTTx link transmits the trigger data to the GEM back-end electronics, while the right VTTx link transmits this same trigger data to the CSC OTMB.

3.2.5 FEASTs

All of the active front-end electronics are powered by FEASTs³. In particular, the GE1/1 detectors use the FEASTMP_CLP module, a modified FEASTMP which is customized for applications whose module stack height is constrained to < 1 cm. A photo of two FEASTMP_CLP modules can be seen in Figure 3.36.

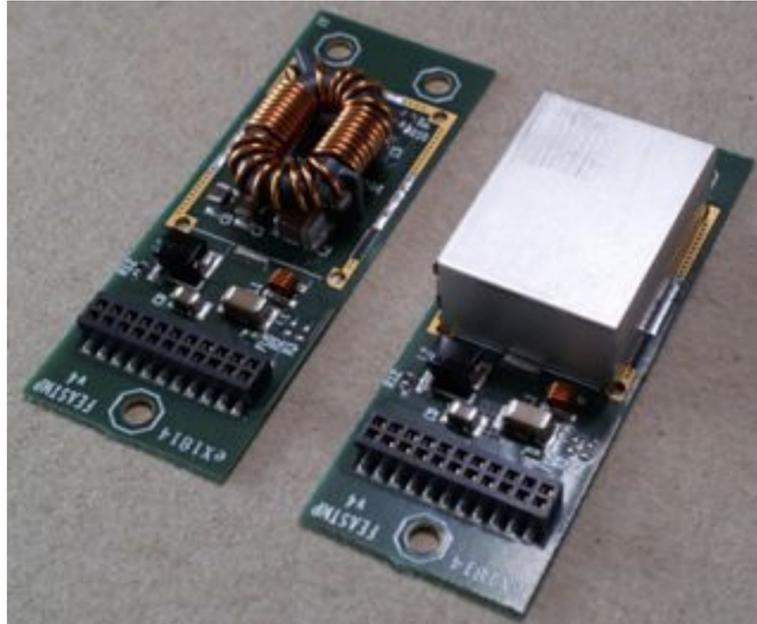
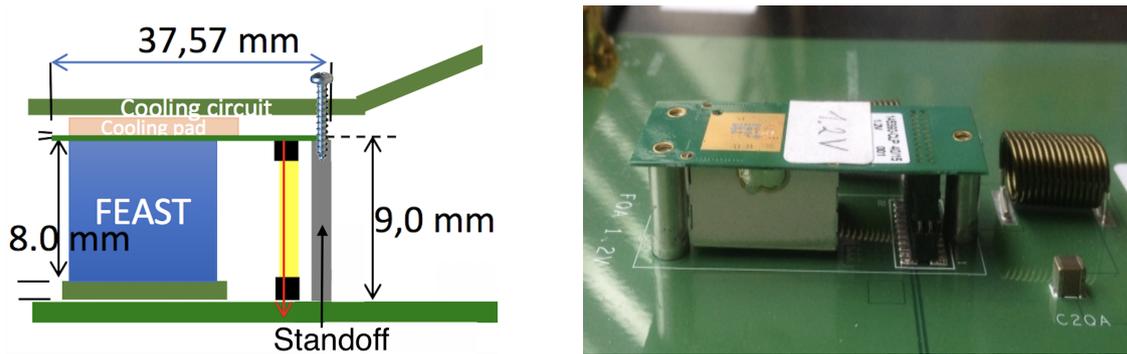


Figure 3.36 – Photo of a fully-assembled FEASTMP_CLP module (right), and a FEASTMP_CLP module with its shield removed to show its main inductor coil (left). The black connector on both is a 12-pin Samtec TFM-106-01-L-D-A male connector, which corresponds to the Samtec SFM-106-01-L-D-A female connector located on the GEB [54]

FEASTs are synchronous step-down buck DC-DC converters which are designed to be both radiation-hard (total ionizing dose of up to > 200 Mrad) and magnetic field-tolerant (tolerant of magnetic fields in excess of 40,000 Gauss), and are produced exclusively by CERN. They take an input voltage of 5 - 12 V, and can provide an output current of 4 A and an output power of 10 W. The output voltage of each module is set at the time of production via the choice of the SMD resistor used in the regulation voltage loop [54]. As they have an overtemperature protection mechanism where they shut down if the temperature exceeds 73°C , they must make good contact with the GEB - as such, they are placed on the GEB upside-down, with a cooling pad placed on top of the FEAST PCB so as to make better contact with the detector cooling circuit. This can be seen in Figure 3.37a.

³Despite FEAST being in all capitals, and thus assumed to be an acronym, the author was unable to find any person or document who knew what it stands for. At the author's best guess, it may stand for "front-end ASIC for science and technology", though this is, admittedly, a complete shot in the dark.



(a) Diagram of a FEAST installed onto a GE1/1 (b) Photo of a 1.2 V FEAST installed on a GEB, showing the connection to both the showing the cooling pad on the bottom of the detector and the cooling circuit of the detector PCB (gold square)

Figure 3.37 – Side-views of installed FEASTs [55]

However, although Figure 3.37 makes the installation appear clean, this is not so in reality. In order to improve the conduction of heat, a layer of thermal paste is applied to the cooling pad. This paste is designed to not dry out and is therefore the perfect recipe for a messy time if not carefully handled - gloves are highly recommended!

3.2.5.1 v2

In total, six commercial DC-DC converters are used to power the front-end electronics. These converters can be seen in Figure 3.33, as the three chips on the left and right sides of the optohybrid. These are used to power the VFATs and the optohybrid itself, including the PROM of the optohybrid, which is a feature that only exists in the v2 electronics.

3.2.5.2 v3

In total, 9 FEASTs are used to power the front-end electronics. The 24 VFATs are powered by four FEASTs in total, each one with an output voltage of 1.2 V. The optohybrid is powered by the remaining six FEASTs: 1.0 V, 2.0 A; 1.0 V, 4.0 A; 1.2 V, 2.0 A; 1.5 V, 3.5 A; and 2.5 V, 3.0 A. There is a tenth slot on the GEB for an additional FEAST of 1.8 V, 1.0 A; however, as a result of the promless nature of the v3 optohybrid, this FEAST ended up not being necessary, and its designated area on the GEB remains simply as an artifact of the design process.

3.3 Back-End Electronics

3.3.1 μ TCA Crate

The back-end electronics are all based on the μ TCA⁴ standard, an open standard embedded computing specification that was first ratified in 2006. In particular, CMS utilizes the MTCA.4 specification for its rear I/O and precision clocking abilities. It utilizes advanced mezzanine cards (AMCs) in a hot-swappable backplane format, and is a highly-reliable standard, providing up to 99.9999% uptime in a smaller form factor / footprint than previous standards such as the 3U and 6U Eurocard systems, with a typical mid-sized

⁴Also called the MTCA standard, outside of high energy particle physics where people aren't cool enough to utilize Greek letters for SI prefixes.

4HP single AMC measuring $73.8 \text{ mm H} \times 18.96 \text{ mm W} \times 181.5 \text{ mm L}$, and even smaller compact 3HP AMCs with a width of just 13.88 mm. A typical μ TCA crate as used in CMS has space for 12 4HP AMCs. To either side of the AMCs are two μ TCA carrier hubs (MCHs), which provide the system with Intelligent Platform Management Interface (IPMI)-based shelf management and switching functionality.

The μ TCA standard uses high-speed serial fabrics, which may include PCIe, Serial RapidIO, and, with the most interest to CMS, gigabit Ethernet (GbE). All of this can be seen in Figure 3.38, which describes the typical backplane port setup for a μ TCA crate being used within CMS.

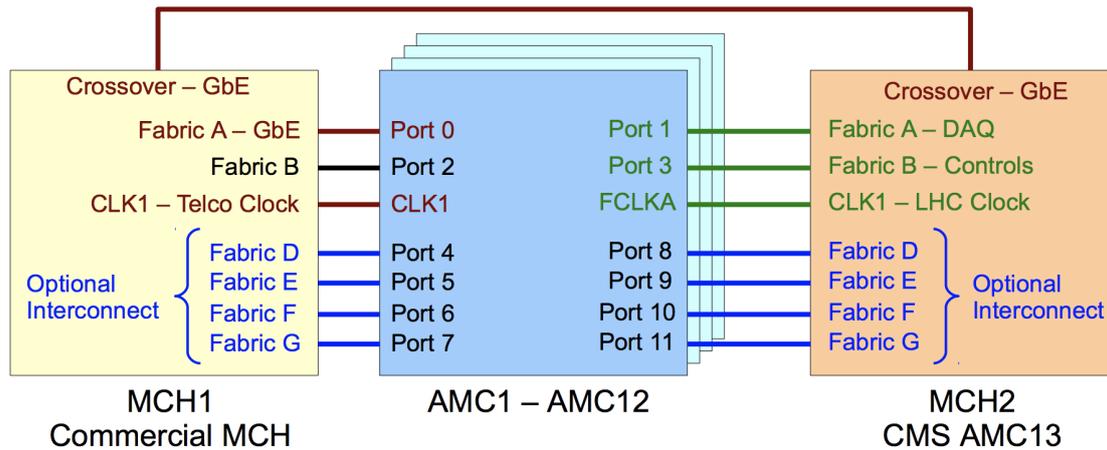


Figure 3.38 – Backplane ports for a typical CMS μ TCA crate setup [56]

3.3.2 Calorimeter Trigger Processing (CTP7) Card

The Calorimeter Trigger Processing (CTP7) card was originally designed for, as the name suggests, the calorimeter trigger, which was upgraded for Run II of the LHC in order to operate under the higher center-of-mass energy and pile-up that would be present then. These upgrades would require a more powerful, flexible, and compact processing card, and so the CTP7 was built with the intent of being implemented for the first layer of the upgrade calorimeter trigger system. However, the needs of the calorimeter trigger would conveniently match well with the needs of the CMS GE1/1 project, and so this card was chosen for use as the GE1/1 AMC card as well. A photo of a CTP7 card is shown in Figure 3.39.

The CTP7 card's primary data processor is a Virtex-7 (XC7VX690T) FPGA, so named as it contains 690 thousand logic cells, and has a speed grade of -2, making it a powerful asset. With 53 Mb of block RAM, it can handle large DAQ buffers, and contains an additional external 36 Mb RAM chip for FIFO use. The Virtex-7 FPGA also features optical links which operate at up to 10 Gbps: on the front side this includes 40 RX and 36 TX optical links, and on the back side this includes 27 RX and 12 TX optical links, for a total of 67 RX and 48 TX links [57].

The CTP7 is the first such card within CMS to utilize a ZYNQ system-on-a-chip (SoC) running embedded Linux. The main purpose of using it is to offload ancillary control functions from the Virtex-7 FPGA to the ZYNQ-7000 (XC7Z045) SoC FPGA, providing transmission control protocol / internet protocol (TCP/IP) communication and board support functions including powerful monitoring and control features. This ZYNQ

FPGA can be seen in the bottom right-hand corner of Figure 3.39.

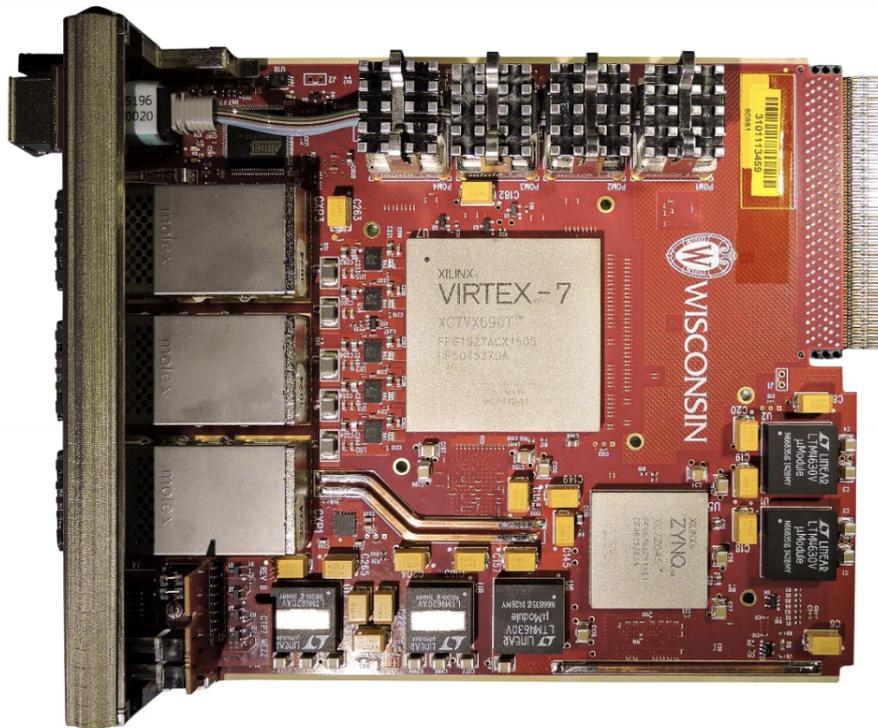


Figure 3.39 – Calorimeter Trigger Processing (CTP7) card [58]

The ZYNQ-7000 SoC provides a dual ARM Cortex-A9 processor which is able to run up to 667 MHz, with 256 kB of on-chip RAM and a Kintex-7 FPGA. As it is implemented on the CTP7, it also contains 512 MB of DDR-1066 memory, a backplane GbE connection, a front panel USB console interface, and a 32 GB file system stored on a microSD card. This microSD card is split into three partitions: The first partition contains the ZYNQ image, which is automatically loaded on boot-up, the second partition contains a root file system, and the third partition provides permanent storage which is available to users.

As previously stated, the ZYNQ SoC runs on an embedded Linux operating system, which includes a full file system, peripheral device drivers, and a functional shell. This streamlines the debugging process and allows for low-latency operations to be performed on the card. A library of custom Linux-based software was developed, including the following desirable features [59], [58]:

- Remote programming and debugging via a virtual Xilinx cable which allows for the remote use of the Xilinx toolset. This virtual cable allows the ZYNQ chip to connect to the Virtex-7 JTAG pins through programmable logic, rather than more commonly-seen crate-based JTAG solutions.
- Embedded applications for monitoring, control, readout, and a web server
- A remote procedure call service which allows for remote card access
- uHAL Hardware Access Library via the TCP/IP interface
- An integrated eye scan engine for link monitoring
- A dedicated memory service for ZYNQ / Virtex-7 memory access

- An IPMI-integrated initialization agent for board initialization

As the back-end AMC card, the purpose of the CTP7 within the GE1/1 system is threefold: to handle slow control requests, interpret transmission, timing, and control (TTC) signals, and read out the trigger and tracking data from the optohybrids.

For the first, the CTP7 acts as the interface between the control and monitoring software and the front-end detector electronics. It utilizes the fabric A lines of MCH1 to port 0 to connect to a GbE network in order to handle TCP requests, and then forwards them to the appropriate subsystem, whether to itself or via the optical link to the optohybrids and VFATs.

For the second, the CTP7 receives the TTC commands via the fabric B lines from MCH2 to port 3, and then forwards them, along with the LHC clock which it receives on the CLK1 line, to the appropriate subsystems.

Lastly, for the third, when the VFATs or optohybrids send data to the CTP7, it formats that data and sends it on the backplane to a dedicated DAQ AMC called the AMC13 [60], which will be discussed in Section 3.3.3.

The optical interface of the CTP7 with the on-detector electronics is accomplished through a total of 36 transmitters and 48 receivers. Of these, 12 receivers are for the trigger data, and 36 transceivers (transmitter/receiver) for the tracking data and slow control [61]. These transceivers are in the form of GBT-FPGA cores, an implementation of the GBT chipset on the FPGA which is explained in much greater detail in Reference [62]. The firmware for this is shown in Figure 3.40. Here, the GBT RX cores appear lime green, the fixed-latency GBT Tx cores appear red, and the rest is largely DAQ (data readout, buffering, processing, error checking, status reporting via TTS, the trigger throttling system), control (fast control, slow control), and trigger (monitoring, simple local trigger, S-bit connectivity scans) [59].

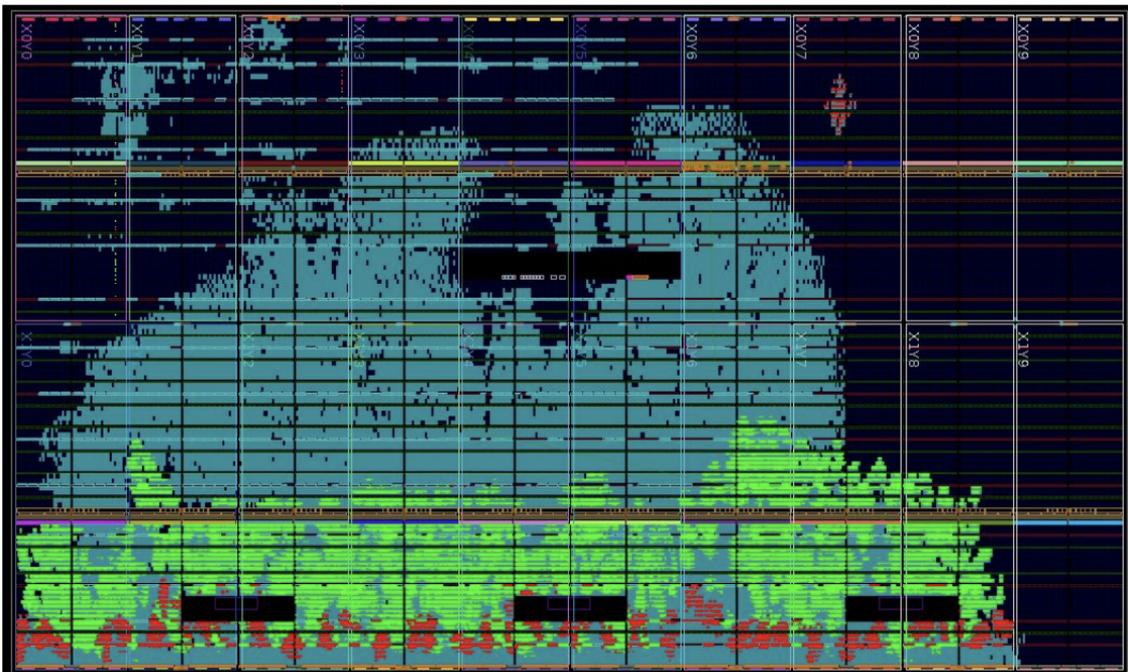


Figure 3.40 – CTP7 firmware, showing the GBT RX cores (green) and fixed latency GBT TX cored (red) [59]

In total, the GE1/1 system will require 12 CTP7s, with only one being required for the

slice test, which will be discussed in Section 4.4. This is due to the fact that a single CTP7 card can run up to 36 GBT cores at 4.8 Gbps, which equates to 12 GE1/1 optohybrids. As the GE1/1 detectors are arranged as two-detector superchambers, a single CTP7 card can therefore cover 60° within the endcap. This coverage is convenient, as the Modular Track Finder (MTF7) of the CMS Endcap Muon Track Finder (EMTF) also covers 60° within the endcap.

However, the ability of a single CTP7 card to run 12 detectors means that the S-bit data must necessarily be concentrated, as 24×3.2 Gbps from the optohybrids and 8×9.8 Gbps to the EMTF [59].

3.3.3 AMC13 Card

For the GE1/1 back-end electronics (and, indeed, all CMS μ TCA-based systems), the MCH2 slot of the μ TCA crate is filled with a custom AMC card, the AMC13XG, commonly simply called the AMC13⁵. The AMC13 is the link between the μ TCA crate and the CMS central DAQ, delivering the TTC signals and LHC clock to the CTP7s, and retrieving the trigger and tracking data from the backplane. The card derives its name from the fact that it is the thirteenth AMC module in the μ TCA crate, and from the 10 Gbps speed of the optical links and backplane connection to the CMS central DAQ, with X being the Roman numeral for 10 and G being short for Gbps. While standard μ TCA crate normally contains 12 AMC modules and two MCHs, the MCH slots are redundant, so the GE1/1 system instead uses the second MCH slot for the AMC13 module, taking advantage of the fact that the MCH slots are the only slots which have the ability to directly communicate with the 12 standard AMC slots.

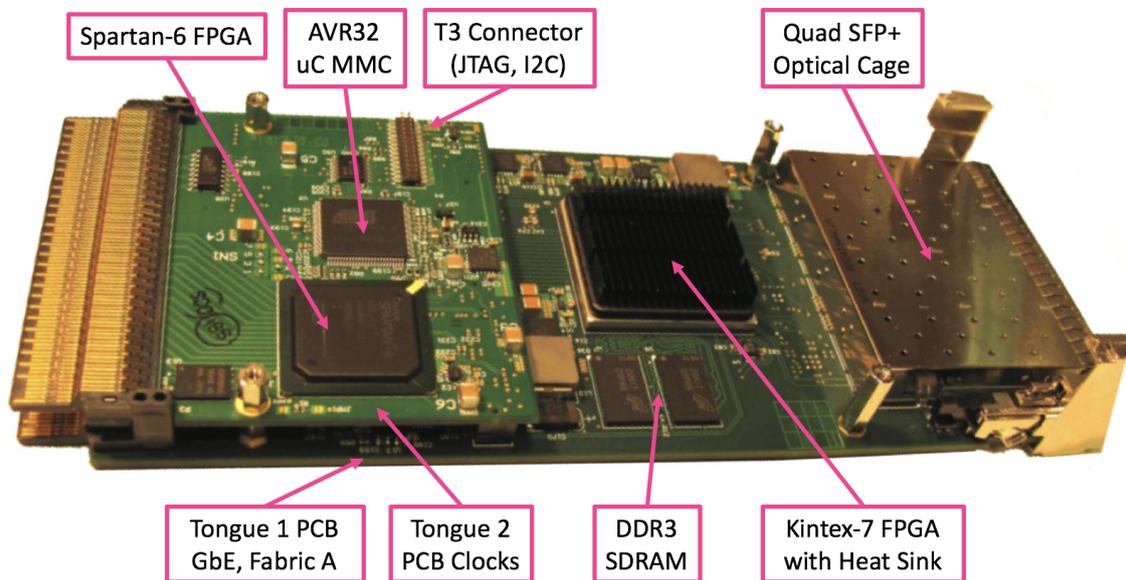


Figure 3.41 – Photo of an AMC13 card with main features labeled. Note that here, the Tongue 3 connector board has been removed in order to better show the internal detail.

A photo of the card can be seen in Figure 3.41, and its block diagram can be seen in Figure 3.42. Here, we see that the AMC13 is equipped with two different FPGAs. The

⁵This common name may lead to confusion with the original, obsolete AMC13, which operated with a Virtex-6 instead of a Kintex-7 as the main data handling FPGA, but it may be assumed that any references to the AMC13 of the CMS GE1/1 system of GE1/1 slice test refer to the newer AMC13XG module, and written off as physicists simply being the confusing creatures we are.

first, on Tongue 1, is a Kintex-7, and is intended for data handling. The second, on Tongue 2, is a smaller Spartan-6, and is intended for control and monitoring.

The Kintex-7 FPGA is connected to the CMS central DAQ front-end readout optical links (FEROLs) through the two enhanced small form-factor pluggable (SFP+) transceiver optical links at up to 10 Gbps, implementing the S-Link Express protocol. If more bandwidth is required than can be provided by these two SFP+s, a spare SFP+ is available, as seen in Figure 3.42.

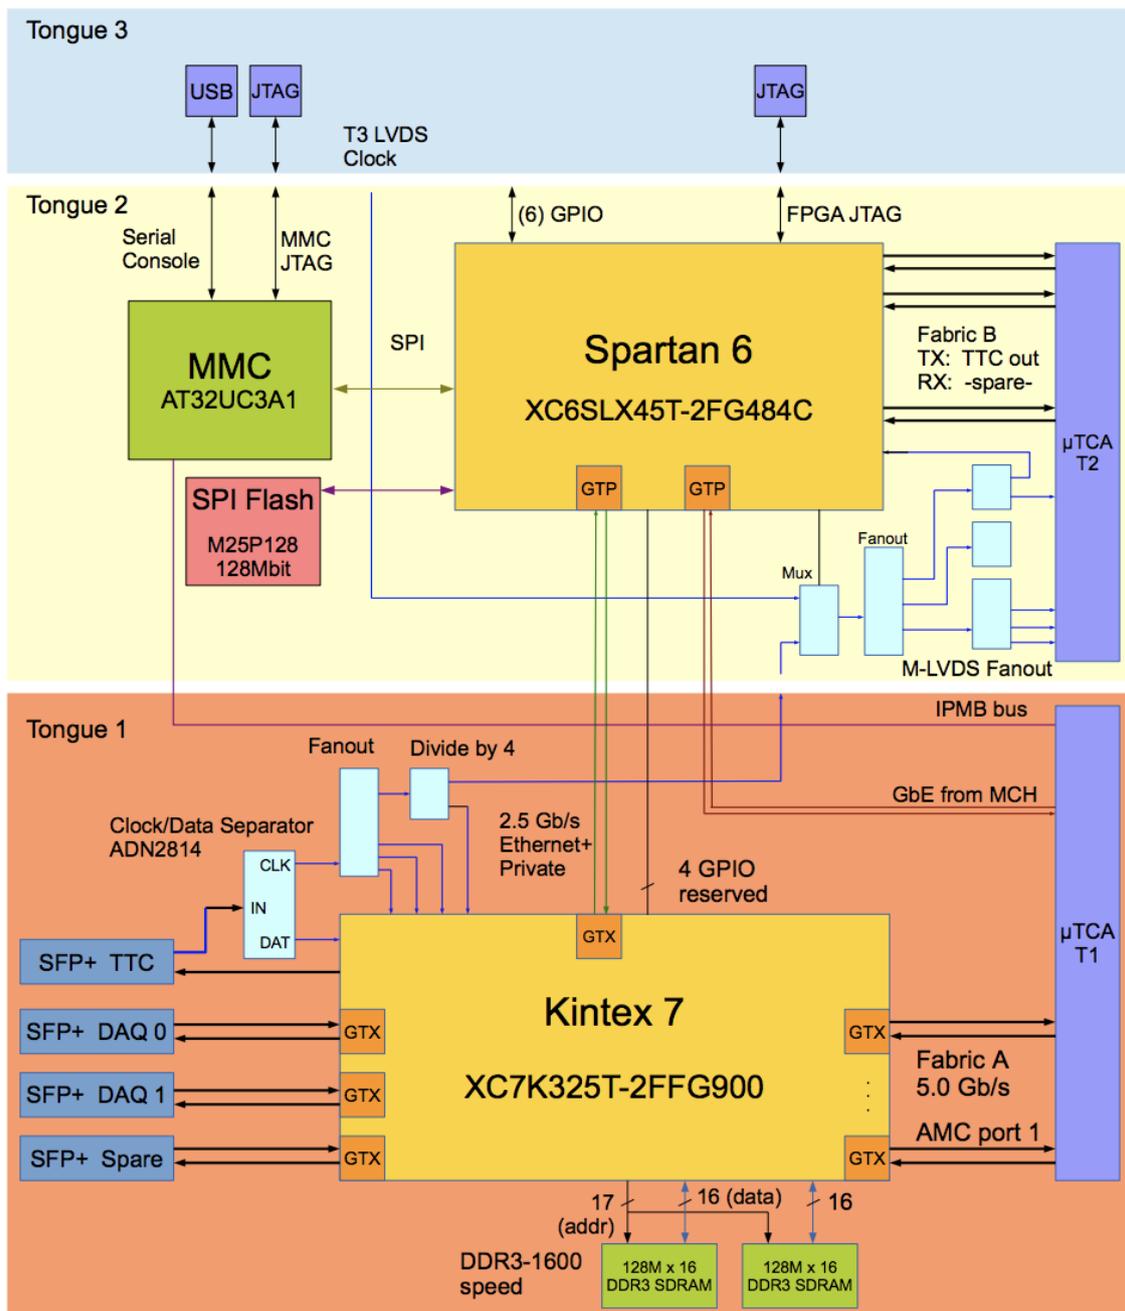


Figure 3.42 – AMC13XG block diagram [56]

As the AMC13 is installed into the MCH2 slot of the μ TCA crate, it has access to all of the AMCs in the crate. It can therefore retrieve all tracking and trigger data over the fabric A ports, which run at 5 Gbps for the purposes of GE1/1 (with the capability to run between 2-6 Gbps if needed). The AMC13 receives the TTC signal over an SFP+ optical

link, labeled in Figure 3.42 as “SFP+ TTC”, from which it extracts the LHC clock. This clock is then forwarded on to the AMCs over the CLK1 line. The TTC commands, which are also extracted from this signal, are passed on to the Spartan-6 FPGA. The Spartan-6 then formats these signals on the fabric B lines at 80 MB/s. The AMC13 also has the ability to produce internal LHC clock generators and TTC commands for developmental purposes where the AMC13 is not installed into a CMS system.

In addition to receiving the TTC signals from the timing and control distribution system (TCDS), the AMC13 sends TTS signals back to the TCDS to inform it about the readout status of the GE1/1 system, including ready status and error information such as hardware failure, imminent data overflow, loss of synchronization, and busy status. If an error status is received through these TTS signals, the TCDS can regulate the L1A rate to allow the GE1/1 system to be able to cope [60].

3.4 Scan Types

In order to assess and monitor the health of the front-end electronics, a number of scan routines have been developed. Although not an exhaustive list, the following sections describe the most commonly-used scan routines within the context of this dissertation.

3.4.1 S-Curves

One method of monitoring the health and behavior of the chambers is through the taking of a calibration pulse scan known as an S-Curve, the process for which is described in Figure 3.43 for a single VFAT.

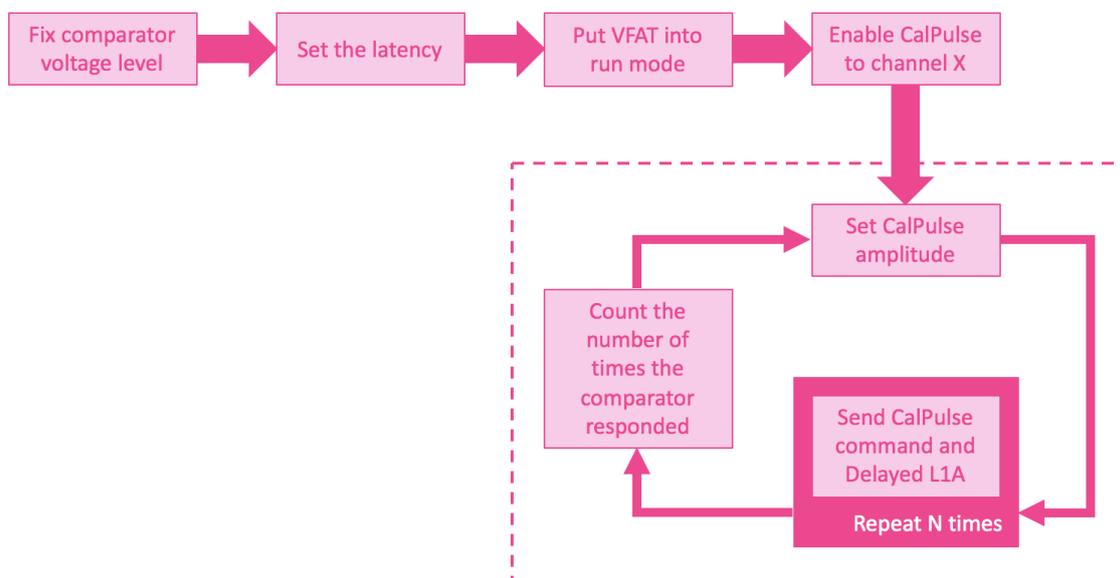


Figure 3.43 – Flow diagram illustrating the process for taking an S-Curve scan

The first four steps outside of the dashed box are performed only once, in order to fix the settings of the VFAT. First the voltage level of the comparator (VFAT2) or CFD (VFAT3) and the latency of the VFAT are set, and the VFAT is put into run mode. Then, for a given channel, the calibration pulse (CalPulse) is enabled - for VFAT2, this uses the calibration circuit described in Section 3.2.3.1, and for VFAT3, this uses the current pulse calibration circuit described in Section 3.2.3.2 by default, although the scan routine code

ultraScurve.py provides the command line option “-voltageStepPulse” if the user wishes to use the voltage step calibration circuit instead.

From here, an iterative process begins. The amplitude of the CalPulse is set to the minimum value the user desires, and a configurable number N of these CalPulses are injected along with a delayed L1A signal. The number of times the comparator responds in reaction to these CalPulses is then recorded. The cycle then begins again, with the CalPulse amplitude incremented by some value, until the scan reaches the maximum desired CalPulse value. By default, the scan range is the equivalent of $[0,255]$ DAC units, incrementing each time by 8 DAC units. These DAC units are converted into charge units during the analysis process using the conversion factor determined during the initial calibration of the circuitry (using an external ADC for VFAT2 and the integrated ADC for VFAT3).

The result for a given channel should theoretically take the form of a step function, with the comparator having a zero response until the CalPulse’s amplitude exceeds the comparator’s threshold, at which point the comparator should respond 100% of the time. However, unavoidable electronic noise inevitably degrades this step function, resulting in the characteristic S shape, seen in Figure 3.44, from which the S-curve derives its name. Because of this degradation, we take the 50% point of the curve to be the channel’s threshold, μ , which is also referred to as the “turn-on point”. Here, the comparator responds to the injection of a CalPulse 50% of the time.

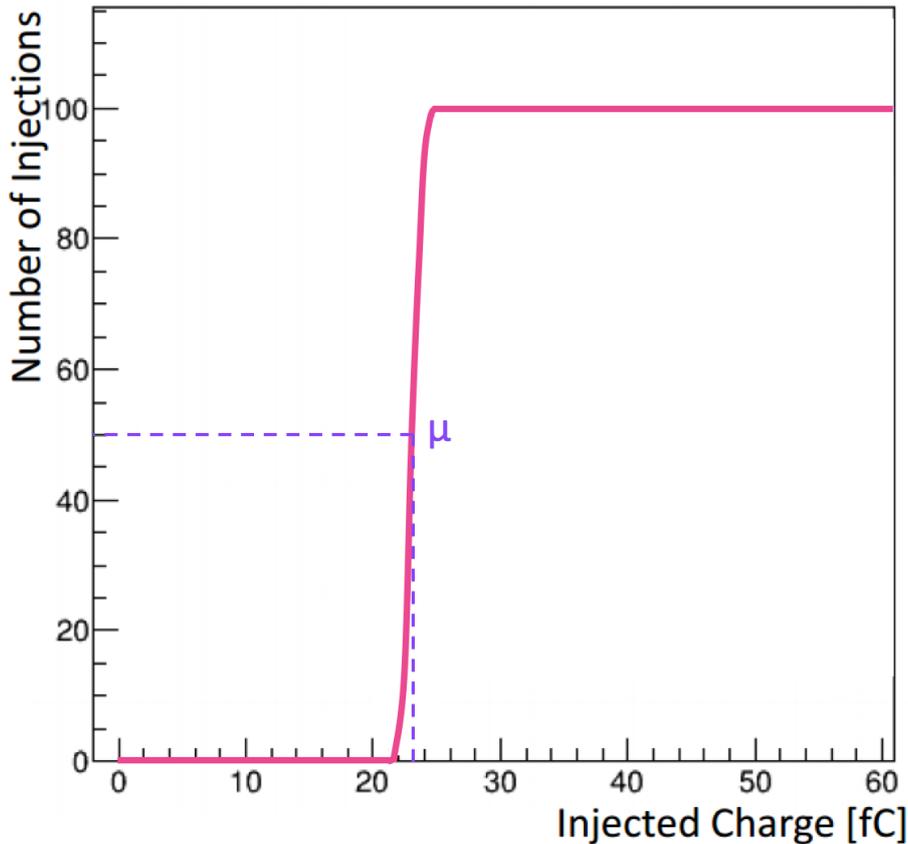


Figure 3.44 – Example of an S-Curve of a single channel on one VFAT showing the location of the turn-on point

Each S-curve is then fit with a modified error function, given by Equation 3.3, where $A = (\frac{N}{2})$ for N charge injections, P is the number of hits at zero injected charge (the

“pedestal”), Q is the charge injected, μ is the channel’s threshold (also known as the 50% response point, turn-on point, and the mean of the error function), and σ is the equivalent noise charge (ENC) or simply the noise.

$$f(q) = A \cdot \text{erf}\left(\frac{\max(P, Q) - \mu}{\sigma - \sqrt{2}}\right) + A \quad (3.3)$$

For Q , a calculation is performed which is dependent on the type of VFAT and, in the case of VFAT3, the type of calibration circuit used. For VFAT2, the calculation uses a simple linear fit, where the slope and intercept are given for all VFAT2 as the average values seen in Equation 3.4.

$$Q = mx + b = \left[0.05\left(\frac{fC}{DAC}\right)\right](CAL_DAC) + (-0.8fC) \quad (3.4)$$

The VFAT3 voltage step also uses a simple linear fit, though the slope m and intercept b are given as distinct values for each VFAT rather than an average as for VFAT2. These calibration coefficients are determined per-VFAT as a part of production testing.

For the VFAT3 current pulse, Q is instead given by Equation 3.5, where CAL_DUR is the duration of the current pulse, CAL_DAC is the amplitude of the current pulse, 10 nA is the hard-coded current least significant bit (LSB), and CAL_FS is the register value in the range of [0,3] (See Reference [46]).

$$Q = (CAL_DUR)(CAL_DAC)(10nA)(CAL_FS) \quad (3.5)$$

This fit is performed using an iterative algorithm which moves in steps of eight DAC units through the full range of [0,255], starting with initial values for A , P , μ , and σ . If the initial fit fails, it moves onto the next step, exiting either when a fit is found with $\chi^2 < 50$, or when the full DAC range has been stepped through, in which case it returns the fit with the minimum value for χ^2 .

This scanning-and-fitting process is performed for each of the 128 channels of a given VFAT, and, in the case of a fully-outfitted detector, each of the VFATs of the detector in parallel. The results are then plotted as in Figure 3.45, which is also colloquially called an “S-curve” plot and, indeed, what is most often referred to by that name.

This particular S-curve was taken using the short GE1/1 detector GE11-X-S-CERN-0011, operated in Ar:CO₂ 70:30, with voltages $V_{drift} = 756$ V, $G1 = 378$ V, $V_{transfer1} = 294$ V, $G2 = 370$ V, $V_{transfer2} = 588$ V, $G3 = 352$ V, and $V_{ind} = 420$ V, corresponding to a gain of 10^4 . The VFAT3 chip, integrated into a HV3b_v2 hybrid, was installed onto the detector’s readout sector 10, set to a medium gain and CFD mode, and set to a threshold of $ARM_DAC_THR = 100$. This S-curve therefore represents “normal operation”.

This S-curve is presented as a 2D plot, with the x-axis representing each of the 128 channels of the VFAT, the y-axis representing the amount of injected charge, and the z-axis representing the number of times the comparator fired in response to that injected charge as a percentage of the total number of injections.

Although not normally included in a standard S-Curve plot, Figure 3.45 includes, for illustrative purposes, a marker for the S-Curve fit, in red, on each channel. Here, the red marker represents the mean value from the error function fit, i.e. the the turn-on point for that channel, and the vertical error bars represent the ENC.

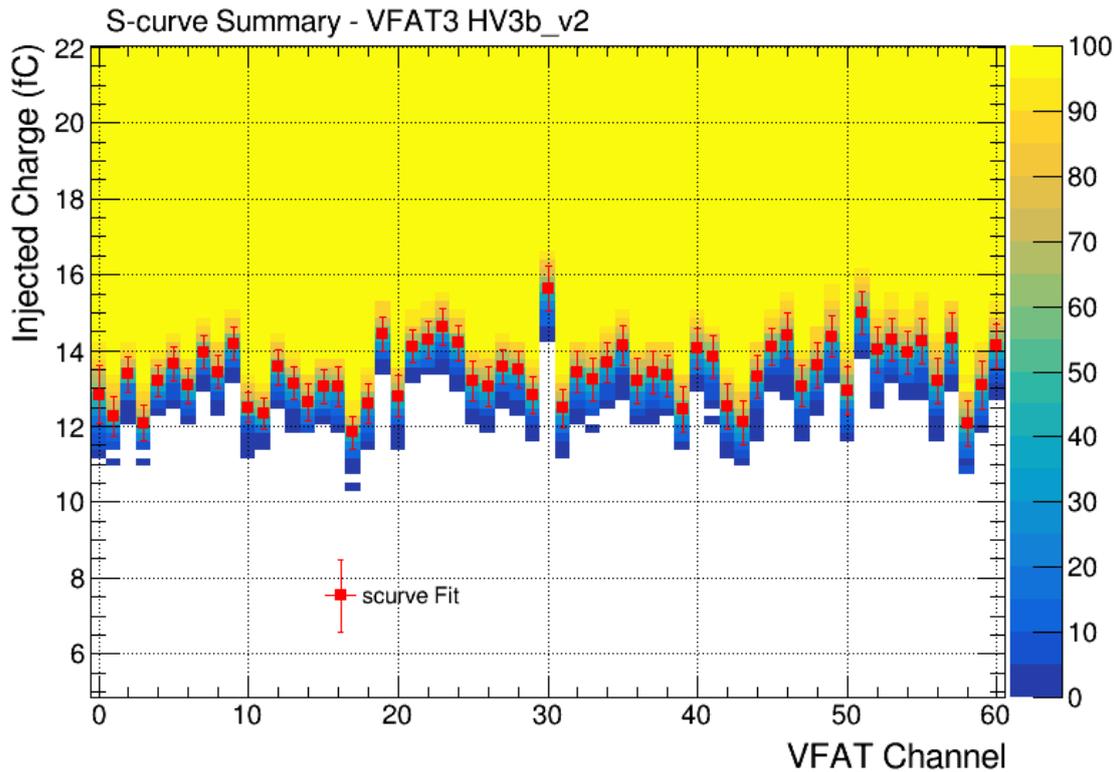


Figure 3.45 – S-curve scan for all channels on a single VFAT3 chip on a HV3b_v2 PCB installed on GE11-X-S-CERN-0011

In general, a “good” S-curve would have a very thin, flat line of green bins across the channels, representing a uniform threshold across each of the channels and a low amount of noise for the chip. An example of a good set of S-curves for a detector is seen in Figure 3.46. This set of S-curves was taken using HV3b_v3 VFAT3s on a production-grade GE1/1 detector during qualification testing in the “QC8” cosmic stand (to be discussed in a later chapter).

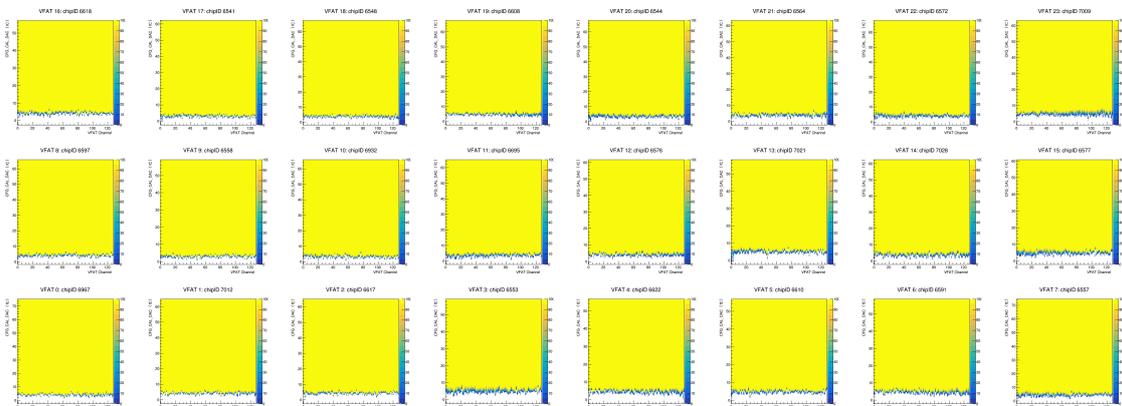


Figure 3.46 – Set of S-curve scans for all VFAT3 chips, integrated with HV3b_v3 PCBs, installed on GE11-X-S-PAK-0001, as taken in Run 206 of the QC8 stand as part of GE1/1 detector qualification

This is considered to be a good set of S-curves, as each VFAT shows a low, uniform set of thresholds across all channels, and very little noise. As the size of these plots makes this

difficult to see on a standard sheet of paper, “candle” or “box” plots can be used to better visualize the data. Figure 3.47 and 3.48 are two such box plots. Such plots are not an automatic part of the S-curve analysis routine, but nevertheless are a useful tool.

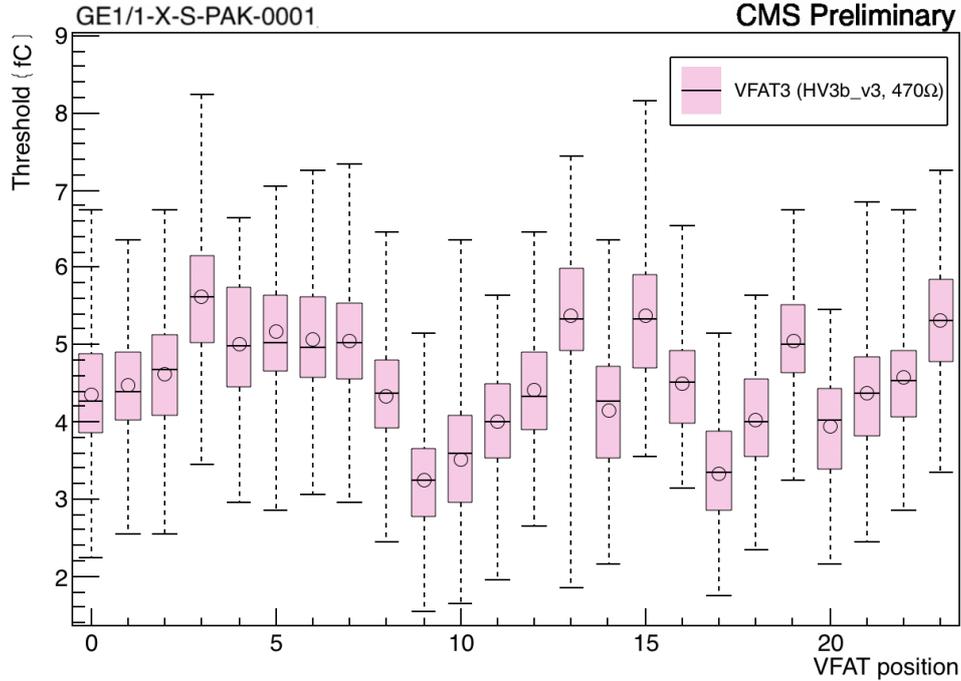


Figure 3.47 – Box plot showing the thresholds of the VFATs for GE11-X-S-PAK-0001, taken from the S-Curves of Run 206 of the QC8 stand [63]

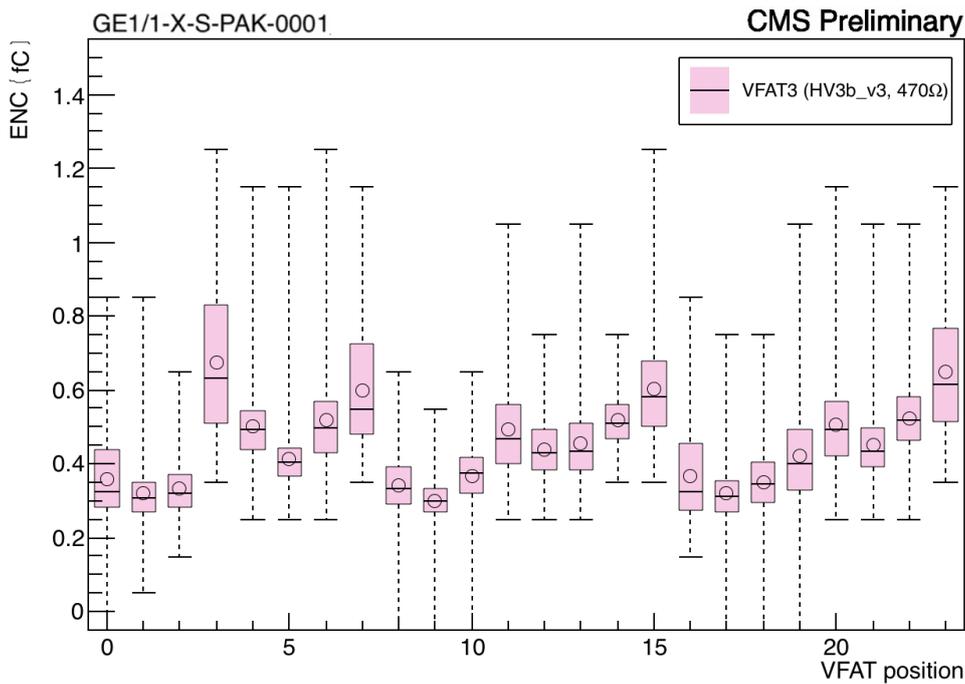


Figure 3.48 – Box plot showing the ENC of the VFATs for GE11-X-S-PAK-0001, taken from the S-Curves of Run 206 of the QC8 stand [63]

These plots visualize each VFAT on the detector as one “candle”. The body of the candle (colored box) encompasses 50% of the channels for that given VFAT, and the “wick” of that candle (dotted lines) encompass 100% of the channels for that VFAT. The mean value is represented by the circle within the body of the candle, and the median is represented by the horizontal line within the body of the candle.

Figure 3.48 is a particularly useful visualization, as it makes a certain trend much more visible than it is in Figure 3.46: that the noise of the VFATs increases from the narrowest eta sector to the widest eta sector, with additional noise for the VFATs closest to the optohybrid in the center of the detector. This trend is expected and, to a degree, present in all such plots for v3 GEM detectors, as the VFAT has an intrinsic noise when coupled to the detector. This noise increases proportionally to the capacitance of the strips which are coupled to each VFAT channel. As the detector width widens, the strips widen as well, increasing the capacitance, and thus increasing the noise, as visualized in Figure 3.48. The additional noise on the VFATs closest to the optohybrid is also expected and generally seen, as a result of the additional electronics and power lines. These same two trends were also seen in the v2 electronics, to a much greater degree; so much greater, in fact, that it was easily visible on a simple set of S-curve plots, such as Figure 3.49, which shows a set of S-curve plots for one of the v2 slice test detectors.

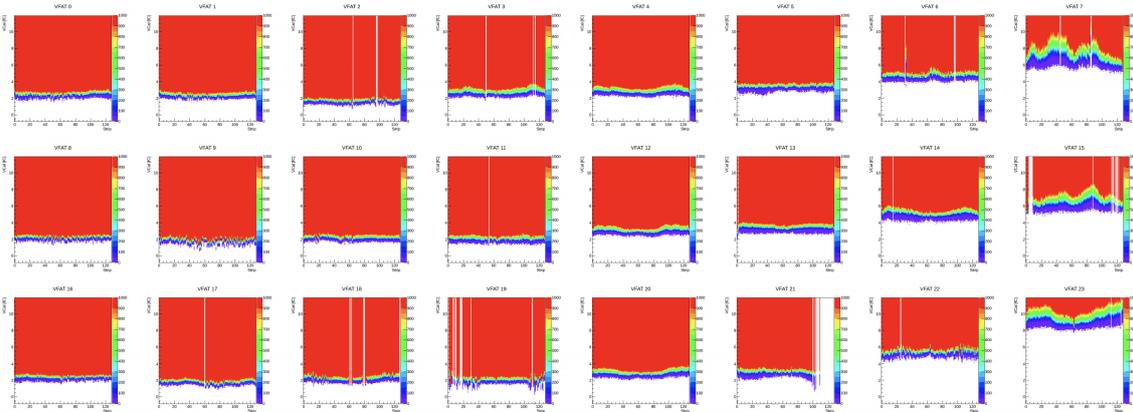


Figure 3.49 – Set of S-curve scans for all VFAT2 chips installed on a v2 slice test detector within the endcap of CMS [64]

Note that these plots were produced prior to a change in the analysis code which modified the color scale, such that red represents the same value (100%) here as yellow does for the more recent v3 plots. Here, the increase of noise in VFATs from the narrow end of the detector (left of the plot) to the wide end of the detector (right of the plot) where the optohybrid is located and the LV power lines converge is abundantly clear.

In case, however, “abundantly” is not clear enough, and because the author of this dissertation is shamelessly vain, Figure 3.50 shows said author at gesturing in the direction of the increased noise, as it was presented at the Topical Workshop on Electronics for Particle Physics in 2018.



Figure 3.50 – Elizabeth Rose Starling, presenting Figure 3.49 at Topical Workshop on Electronics for Particle Physics in Antwerp, Belgium in 2018

One other important thing to note in Figure 3.49 is the presence in some VFATs of white vertical lines. These indicate a problem with either the S-curve fit or the channel itself, and will be discussed in excruciating detail in Chapter 6.

3.4.1.1 Trimming

One method of obtaining more uniform S-curves is through the so-called “trimming” of the VFATs, in which a given value in the S-curves is aligned for all channels.

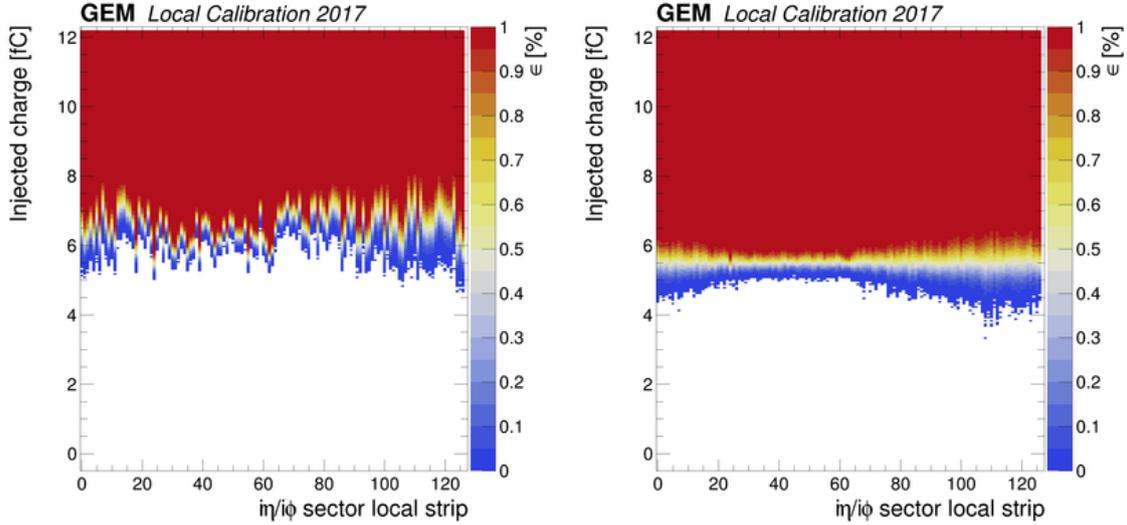
For VFAT2, this was accomplished using the 5-bit TrimDAC register discussed in Section 3.2.3.1, and traditionally aligned the 99% or 50% response point of each channel. This method of trimming resulted in the improvements seen, for example, in Figure 3.51.

For VFAT3, the process was greatly improved upon, first by implementing an iterative trimming procedure which followed the following steps [65]:

- Take a set of reference S-curves. From these reference S-curves, calculate the first set of trim values.
- Use these initial trim values to take another set of S-curves. From these, calculate a set of adjusted trim values
- Take a second set of trimmed S-curves. From these, calculate a set of re-adjusted trim values.
- Take an S-bit threshold scan (as described in Section 3.4.3), obtain new values for THR_ARM_DAC at 100 Hz noise rate⁶ per VFAT.

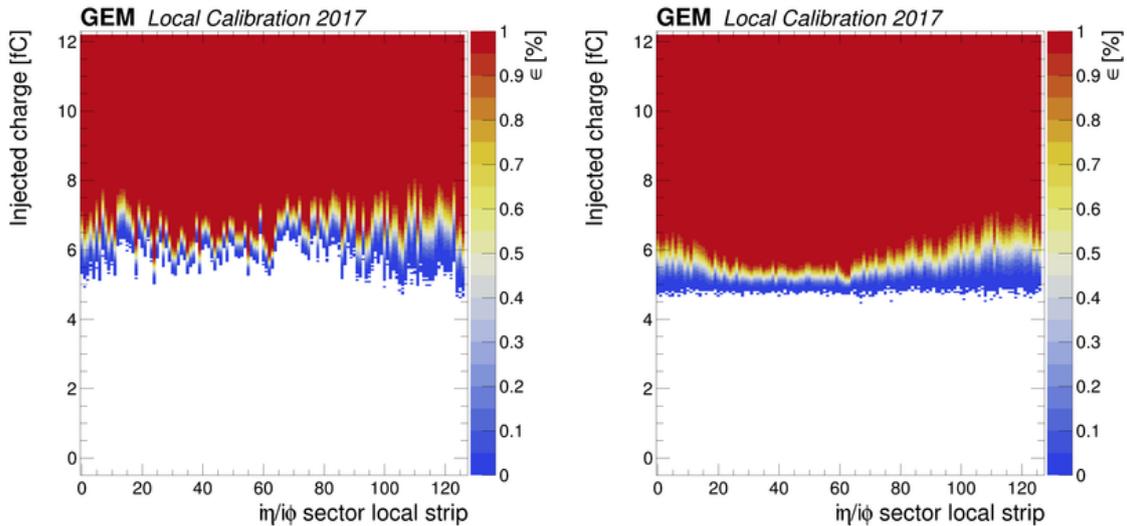
⁶This value of 100 Hz was chosen as it reduces the thresholds well without negatively impacting the efficiency of the detector to a significant degree.

- Take a set of S-curves with the new threshold values and the same trim values as were used for the S-bit threshold scan. From these, calculate a final set of trim values.
- Take a final set of S-curves to verify that the procedure worked as intended.



(a) S-curve scan of slice test detector GEMI-NIm01L2 VFAT #0, pre-trimming

(b) S-curve scan of slice test detector GEMI-NIm01L2 VFAT #0, post-trimming to 50%



(c) S-curve scan of slice test detector GEMI-NIm01L2 VFAT #0, pre-trimming

(d) S-curve scan of slice test detector GEMI-NIm01L2 VFAT #0, post-trimming to 99%

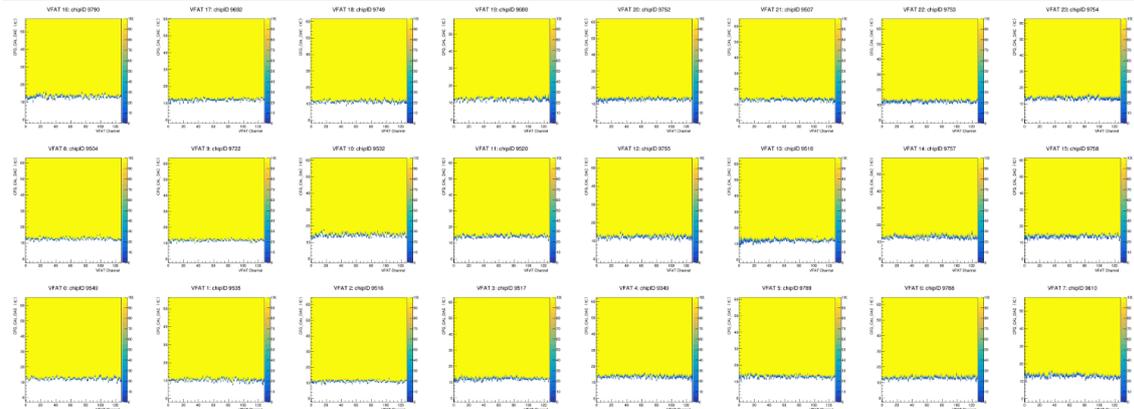
Figure 3.51 – Example VFAT2 S-curve plots before and after a trimming algorithm was applied to align the 50% (top) and 99% (bottom) response points across the VFAT

Then, rather than applying this iterative procedure to the VFATs in a way that aligns their 50% or 99% response point, the procedure is applied to align the S-curves based on their noise level. That is, the value of $X = (\text{mean} - n\sigma)$ for the S-curves of each of the channels is aligned, where “mean” is the 50% response point of the S-curve, σ is the ENC of that channel, and n is some number ≤ 4 that can be chosen by the user for best results, with $n = 0$ being the equivalent of the previous trimming method.

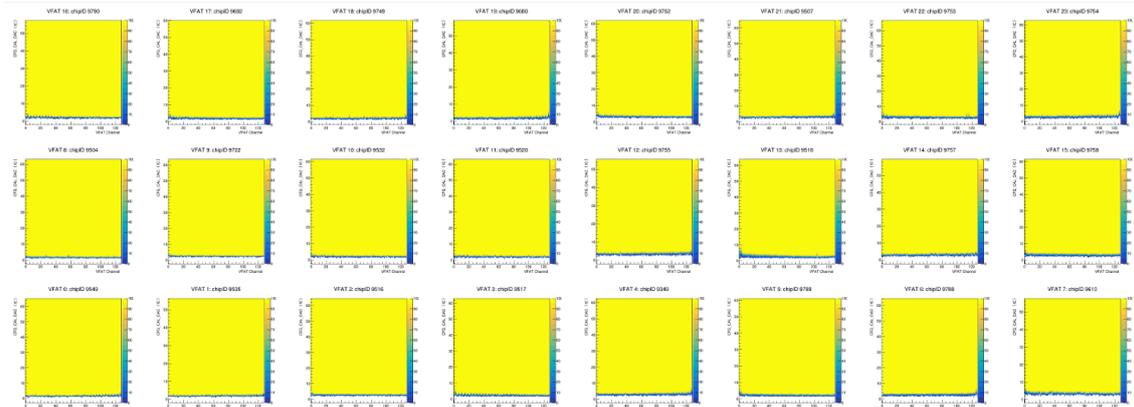
For each iteration, the trim values are calculated using Equation 3.6, where trimDAC is rounded to the nearest integer value.

$$\text{trimDAC} = |(X_{\text{avg}} - X_{\text{ch}}) \cdot 15| \quad (3.6)$$

The result of applying this new, improved algorithm can be seen in Figure 3.52



(a) S-curve scan of GE11-X-S-BARI-0014, pre-trimming



(b) S-curve scan of GE11-X-S-BARI-0014, post-trimming with the iterative σ -aligning method

Figure 3.52 – Example S-curve plots of GE11-X-S-BARI-0014 before and after an iterative trimming algorithm was applied to align the value of $(\text{mean} - n\sigma)$ across the VFATs

3.4.2 Latency Scans

The “latency” of the VFAT chip is defined as the time, in bunch crossings / clock cycles, between an external trigger input and the desired point in memory where the tracking data is stored. In the case of the CMS GEMs, the L1 trigger determines which events are interesting enough to proceed onto the next level of data analysis, and which events are thrown out as being uninteresting or of poor quality. This trigger distributes L1A signals to the GEMs and other subdetectors, and there is a certain period of time between when a muon arrives in the GEM detectors and when the L1A associated with that muon event arrives. In order to ensure that we are reading out data from the correct VFAT memory block, this latency value must be fixed accordingly.

Each VFAT2 has one 8-bit register which is used for programming this latency into the chip, where each bit represents a single clock cycle. The latency can therefore be programmed as between 1 and 256 clock periods (corresponding to $0000\ 0000 = 0$ and $1111\ 1111 = 256$), or up to $6.4\ \mu\text{s}$. The default setting is $1000\ 0000 = 128$ clock cycles = $3.2\ \mu\text{s}$.

On the other hand, each VFAT3 uses a 16-bit register rather than an 8-bit register. However, unlike for VFAT2, the full register is not used. Rather, here the maximum latency is determined by the SRAM memories of the VFAT3 chip. As stated in Section 3.2.3.2, the SRAM1 is a circular buffer which continuously samples all 128 channels every clock cycle, and is 128 bits wide \times 1024 bits deep. This allows it to store data for all 128 channels for up to 1024 continuous bunch crossings before the buffer loops around and overwrites previously-written data. Therefore the maximum programmable latency is 1024 bunch crossings, or $25.6 \mu\text{s}$.

In either case, a latency scan allows the user to determine the best latency value for a given VFAT to be set to, by scanning over a set of latency values. For each latency value, the total number of events is compared to the number of events with verified hits. Figure 3.53 shows an ideal latency scan for a VFAT2 chip. Here, the set of latency values that were scanned over was $[80, 120]$ bunch crossings. This scan is idealized because latency values 95, 96, 97, and 98 show 100% of hits vs. the total number of events (Ratio = 1), while all other latency values are 0%. Thus, the user would choose one of these non-zero latency values, likely 96 or 97 as the middle values to account for margin of error. Note that this window of four 100% latency values would be due to the choice of $\text{MSPL} = 4$, and therefore could be changed with a different MSPL value.

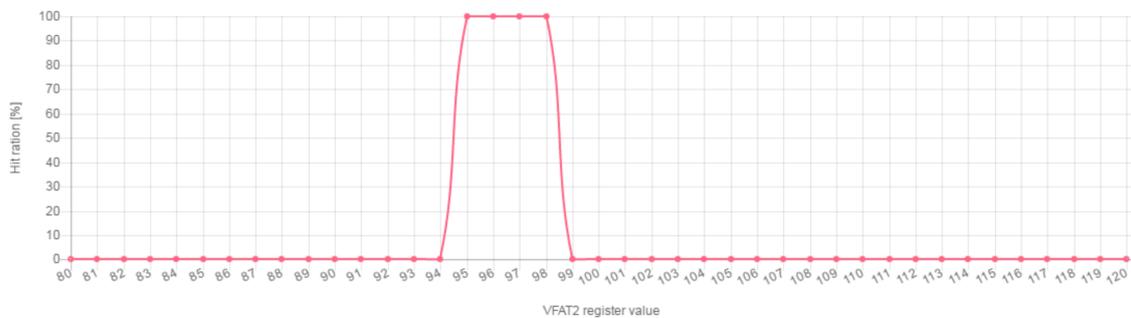
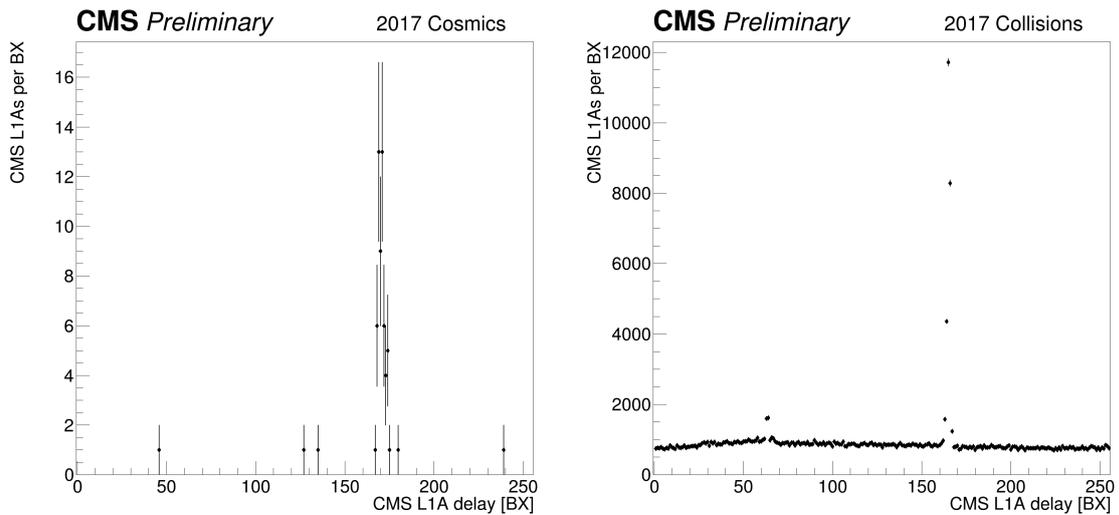


Figure 3.53 – Ideal latency scan for a VFAT2 chip [60]

Of course, the idealized case is rarely, if ever, the real case. Figure 3.54 shows two real latency scans of a VFAT2 chip installed on a slice test detector. In both plots, the x-axis is the L1A delay, or latency value, in number of bunch crossings, and the y-axis is given as an absolute number of hits rather than a percentage or ratio as in the idealized plot. For Figure 3.54a, the scale of this axis is several orders of magnitude lower than for Figure 3.54b, as the former used cosmic ray muon data and the latter used LHC collision data. In both cases, however, we see that, unlike the idealized case, there is not a perfect binary choice of 0% or 100%. As such, a certain degree of user consideration must be used in order to choose the most efficient latency value to set.



(a) Latency scan taken with a slice test detector in 2017 using cosmic ray muons
 (b) Latency scan taken with a slice test detector in 2017 using LHC collisions

Figure 3.54 – Latency scans of a VFAT2 chip installed on a slice test detector

Another thing that must always be kept in consideration is that the optimal latency value will be different for each VFAT in each detector. Figure 3.55 shows the optimal latency values for VFATs in different positions on a single detector, over several detectors. Here, we see that there is a general trend around latency = 175, but that this is variable among VFATs both within a small range around that value, and over the full scan range. Thus, when commissioning a detector, it is crucial to perform this latency scan for each VFAT on that detector and not simply assume a given uniform value.

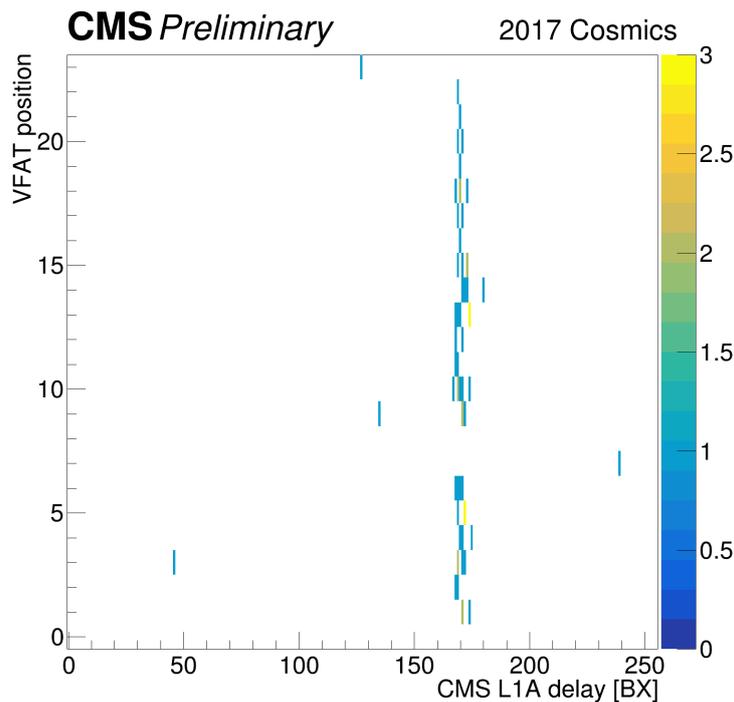


Figure 3.55 – Latency values over multiple VFATs per detector position

3.4.3 S-bit Threshold Scans

Prior to describing the S-bit threshold scans, we must first define what an S-bit is. In Sections 3.2.3.1 and 3.2.3.2, we discussed the building of a “fast OR” for the trigger, which has a granularity of 16 strips for VFAT2 and two strips for VFAT3. This raw trigger data, which is sent over trigger lines from each VFAT and packaged by the optohybrid to indicate where a muon hit occurred within the GEM detector, is called a “select bit”, or S-bit [66].

An S-bit threshold scan aims to determine if any of the trigger lines in either the VFATs or the GEB are broken or damaged in any way, while simultaneously providing a noise rate for any given threshold. To do so, it scans over the range of thresholds of [0,255] DAC units, and for each threshold value, it records the number of trigger signals given by a logical OR of all 128 channels within a given VFAT. That is, if even a single channel records a hit, the S-bit OR value is 1, whether it is only that one channel or all 128 channels. If no channel records a hit, the S-bit OR value is 0.

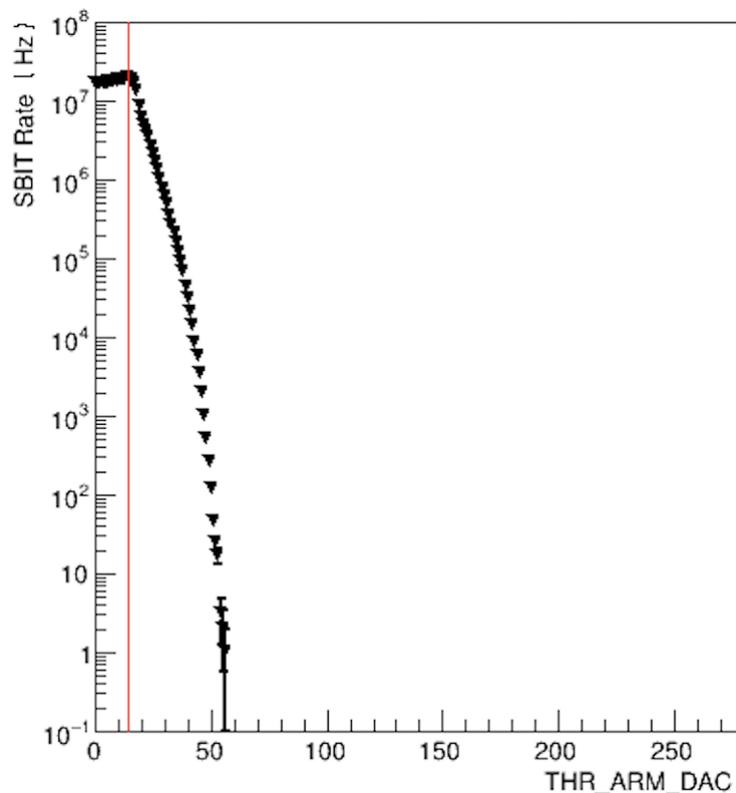


Figure 3.56 – Example of an S-bit threshold scan for a single VFAT, shown for VFAT #11 of GE11-X-L-PAK-0004, as taken in the QC8 stand on April 23rd, 2020⁷

Figure 3.56 is an example of such an S-bit threshold scan. This particular scan was taken in the QC8 stand on April 23rd, 2020 with VFAT #11 of GE11-X-L-PAK-0004.

Here, the x-axis represents the range of THR_ARM_DAC thresholds scanned over, from 0 to 255 DAC units. The y-axis is the measured S-bit rate, in Hz, at each threshold value. In the absence of true muon signal, this rate is representative of the rate of noise hits within the detector for that particular VFAT. The threshold value with the maximum S-bit rate is marked with a red vertical line. From this plot, the threshold value for the VFAT can be determined based on the maximum desired noise rate.

⁷<http://cmsonline.cern.ch/cms-eelog/1104031>

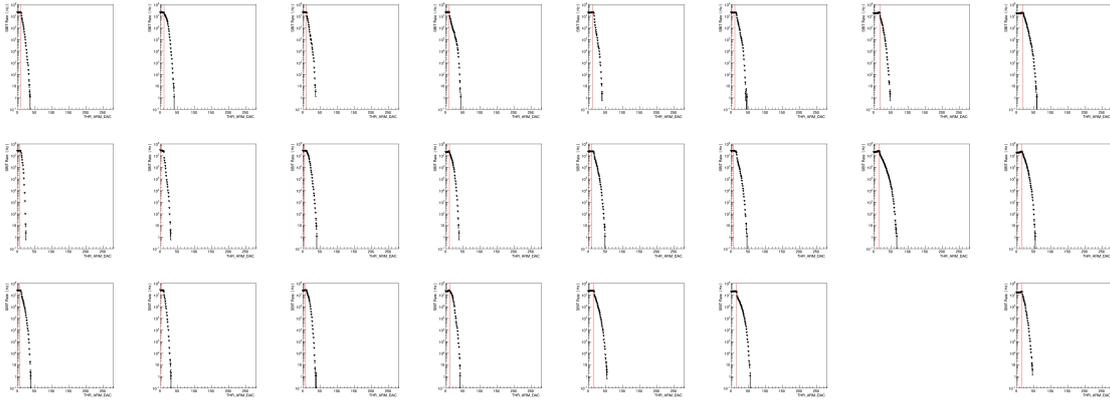


Figure 3.57 – Example of a set of S-bit threshold scans for a full detector, shown for GE11-X-L-CERN-0025, as taken in the QC8 stand on April 23rd, 2020

Figure 3.57 shows a set of S-bit threshold scans for all 24 VFATs on a GE1/1 detector, GE11-X-L-CERN-0025, as taken in the QC8 stand. Here, the plot for VFAT #6 is missing. This is one way in which the S-bit threshold scan can be used to diagnose problems with the trigger lines. Since this particular scan was taken in the QC8 stand, after the initial electronics installation and verification of QC7, this missing plot could indicate one of two issues: For one, it might indicate that the trigger line for this VFAT was damaged in the course of installing the chamber into the QC8 stand, or following any interventions on the detector within the stand. Alternatively, it might indicate that there was a non-physical communication error during the scan, which could be fixed following another scan. If, on the other hand, such a problem was seen in QC7, it might indicate that the trigger line of that VFAT or GEB was broken from the point of production and would need to be replaced.

Chapter 4

GE1/1 and the Slice Test

Having covered the basics of what a GEM detector is, and how it is read out electronically, we now turn to the GE1/1 system, the first GEM-based subsystem of CMS and the first Phase II upgrade system to be installed in preparation for the HL-LHC, and its demonstrator system, the “slice test”.

4.1 The Ancestors of GE1/1

Before we delve in, however, it is important to provide some context to this dissertation in the form of the history of GEM detectors within high energy physics. The GE1/1 system was not the first time that GEMs have been installed within a high energy physics experiment, least of all one located at CERN, and yet many of the problems that we will discuss in the coming chapters were unique to the slice test. Thus it is important to understand precisely what makes the CMS GEMs so different from their predecessors.

4.1.1 COMPASS

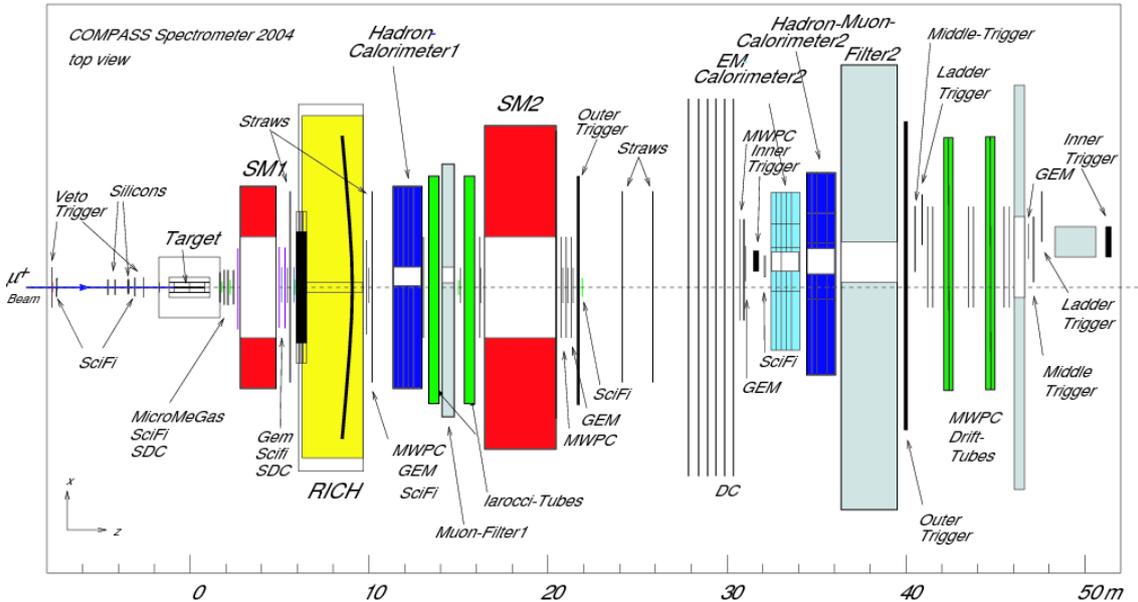


Figure 4.1 – Schematic of the two-staged COMPASS spectrometer as it existed in 2004, with the main components labeled [67]

One of the earliest applications of GEMs to high energy physics was the Common Muon Proton Apparatus for Structure and Spectroscopy (COMPASS) experiment, located at the M2 beam line of the SPS. Prior to the commissioning of the LHC experiments, COMPASS was the largest data-taking experiment located at CERN. It consisted of three main parts: the beam telescope, the fixed target area, and the two-staged spectrometer. It is this

spectrometer where we find our first real application of GEMs to a high energy physics experiment, as shown in Figure 4.1.

The COMPASS GEMs were considered large for their time, at $31 \times 31 \text{ cm}^2$. Their structure, shown in an exploded view in Figure 4.2, consisted of three GEM foils at a gap configuration of 3 mm. These gaps were maintained using grids of thin fiberglass strips between the electrodes, covering the entire active area of the detector. The gain of each amplification stage was decreased by 20% from the first to the third GEM foil, in order to decrease the probability of a discharge within the detector by more than two orders of magnitude with respect to double-GEM detectors of equal gain sharing. Furthermore, the segmentation of each GEM foil into individually-powered sectors decreased the amount of energy stored within the foil, minimizing the amount of charge released during a discharge and minimizing the probability of a discharge propagating to the anode. The detectors, like the CMS GEM detectors, were operated using a gas mixture of Ar:CO₂ 70:30.

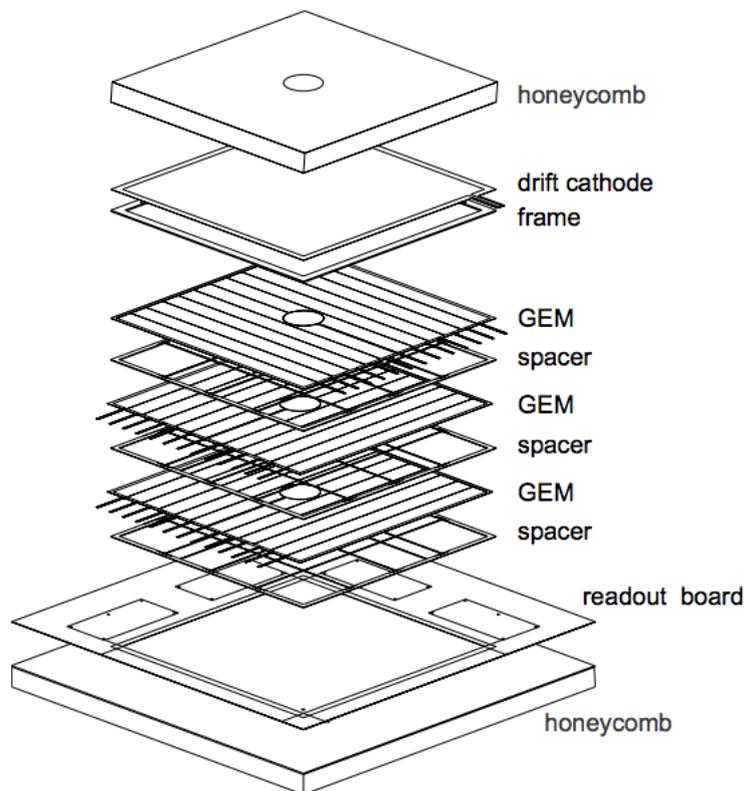


Figure 4.2 – Exploded view of a COMPASS GEM detector [68]

The readout board consisted of two layers of orthogonal readout strips, insulated by thin ridges of 50 μm Kapton. These strips were varied in width such that each layer collected approximately half of the total charge of a given avalanche. This charge was then read out by APV25 chips with an analogue sampling rate of 40 MHz [68].

The main support frames of the detector were two sandwiches of Stesalit - Honeycomb Nomex - Stesalit, where the Stesalit was 125 μm thick and the honeycomb Nomex was 3 mm thick. One honeycomb of dimension $33 \times 33 \text{ cm}^2$ contained the drift electrode and another honeycomb of dimension $50 \times 50 \text{ cm}^2$ contained the readout board, the front-end readout electronics, and the HV distribution scheme. Each of these contained a hole (33 mm diameter in the drift honeycomb to minimize deformation of the foil during vacuum-assisted assembly, and 50 mm in the readout honeycomb), matching the central region of

the detector where a unique feature of the third GEM foil was located - an independently-powered sector of the foil which could be deactivated in high intensity runs, and re-activated in low intensity runs for calibration and alignment studies [69]. Colloquially called the “beam killer”, the application of a voltage lowered by 200 V to this sector, (with respect to the rest of the foil, would entirely prevent the detection of the main beam [70]. This region and its corresponding hole in the drift honeycomb can be clearly seen in Figure 4.3. More details on these detectors can be found in Reference [71].

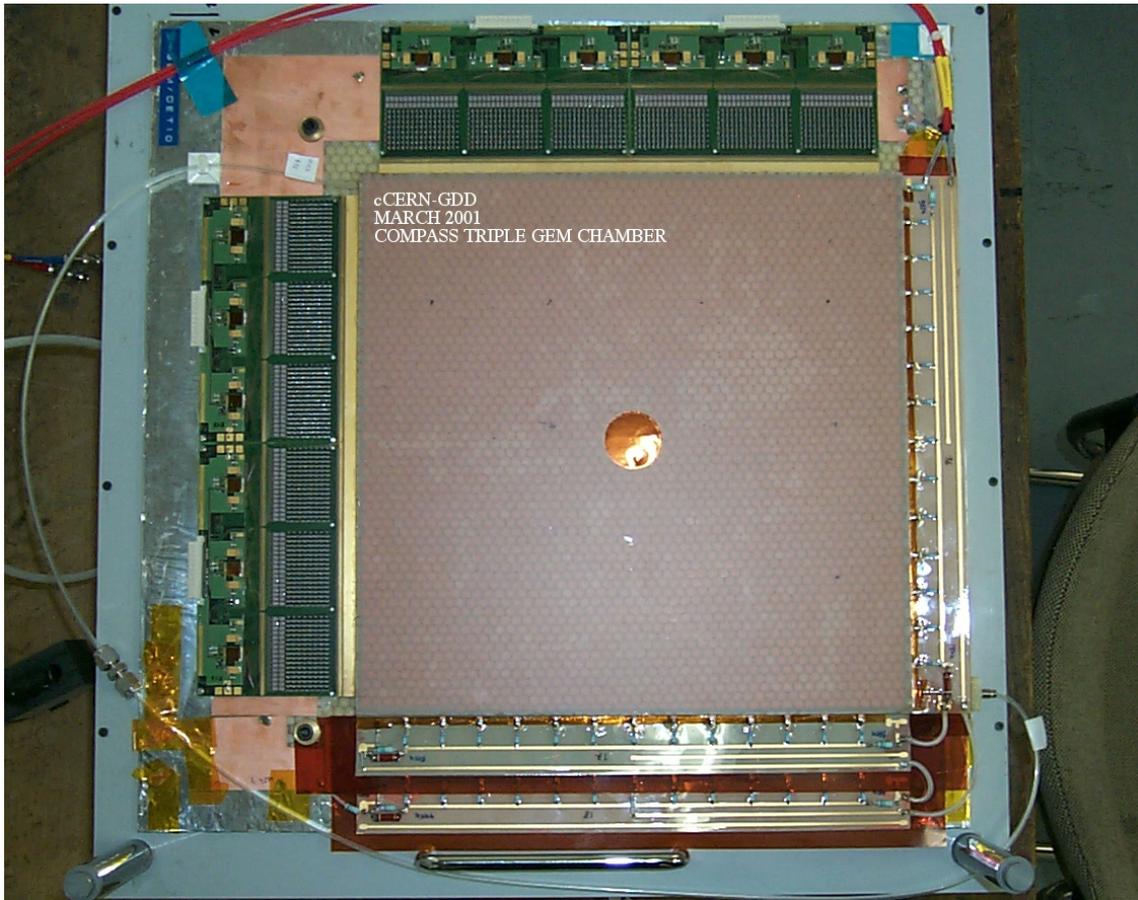


Figure 4.3 – Photo of a COMPASS GEM detector [67]

While there were very few direct similarities between the CMS GEMs and the COMPASS GEMs, this experiment nevertheless paved the way for future GEM-based experimental systems, by proving the operational capabilities of GEMs with a high energy physics experiment.

4.1.2 TOTEM

CERN’s longest-running experiment, the TOTAl Elastic and diffractive cross section Measurement (TOTEM) experiment, used triple-GEM detectors in its T2 telescope until 2019, when the telescope was removed during Long Shutdown II (LS2) as a result of both extensive radiation damage and its incompatibility with the new beam pipe redesign for LHC Run 3. Although similar in size to the square detectors of COMPASS with an active area 28 cm in width, TOTEM utilized a unique semi-circular geometry for its GEM detectors, shown in Figure 4.4. Eight such GEM detectors were arranged in two opposing planes, as shown in Figure 4.5. One plane of four detectors was dedicated to providing a precise posi-

tion resolution, while the other plane of four detectors was dedicated to providing improved trigger and pattern recognition. The difference between the detectors of the two planes lies solely in their readout boards - for the tracking detectors, the orthogonal readout strips were designed such that the charge of a given electron avalanche is always shared on at least two electrodes in order to provide the best resolution, whereas the readout board of the trigger detectors utilized larger pads whose size varied in order to ensure that they would be traversed by the same particle flux regardless of their location [72].

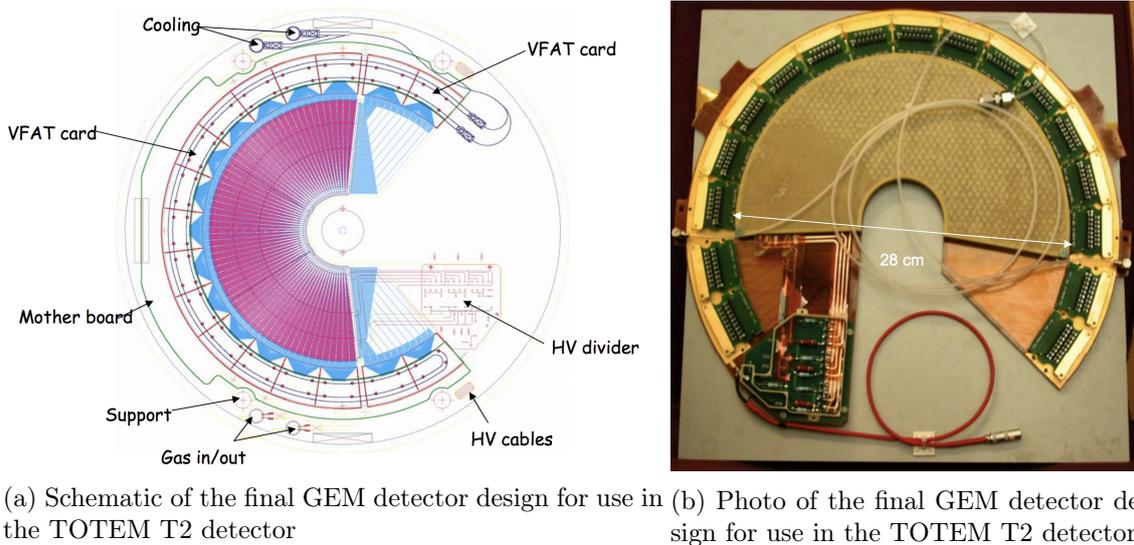


Figure 4.4 – Schematic and photo of a TOTEM GEM detector [70]

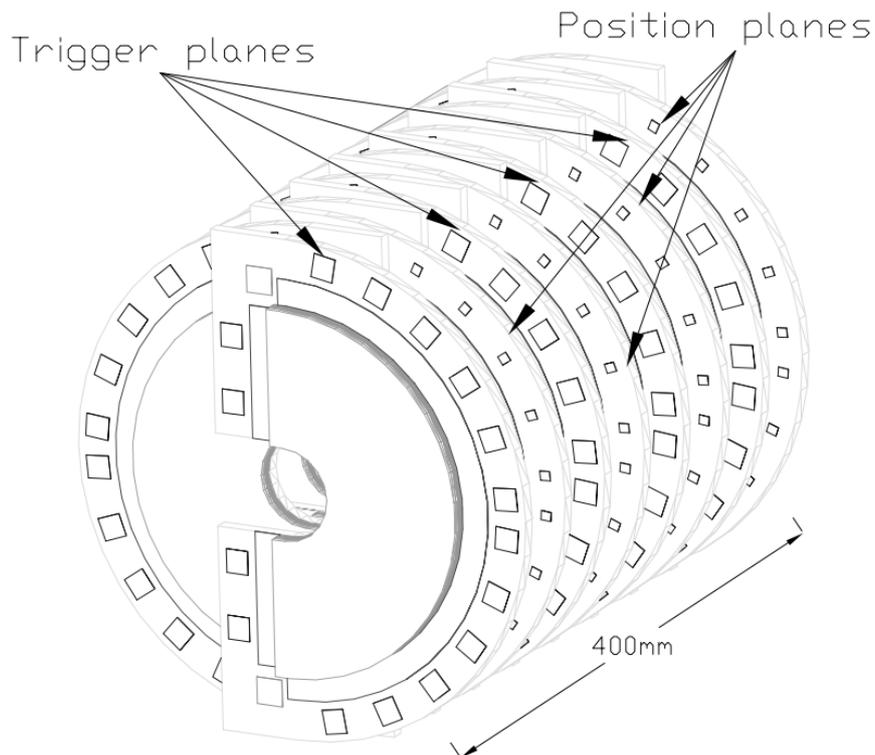


Figure 4.5 – Schematic of the alignment of the GEM detectors of the TOTEM T2 experiment [72]

As with the COMPASS GEMs, the TOTEM GEMs utilized a triple-GEM structure, shown in Figure 4.6, which was enforced on the anode and drift sides by a honeycomb structure. And, in the same way that the COMPASS GEMs segmented the foils into multiple individually-powered HV sectors, the TOTEM GEMs were segmented into four HV sectors, as shown in Figure 4.7.

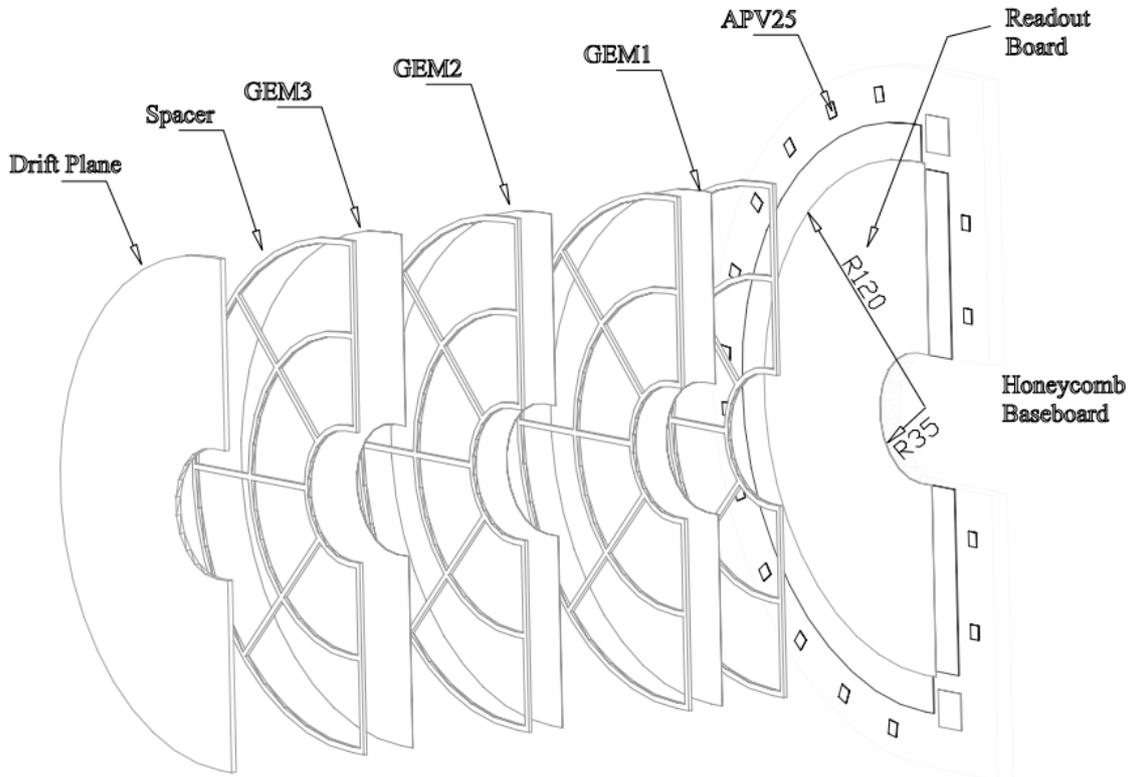


Figure 4.6 – Schematic of the GEM detectors of the TOTEM T2 experiment. Note that this is an earlier schematic depicting the presence of APV25 readout chips rather than the newer VFAT2 chips [72]

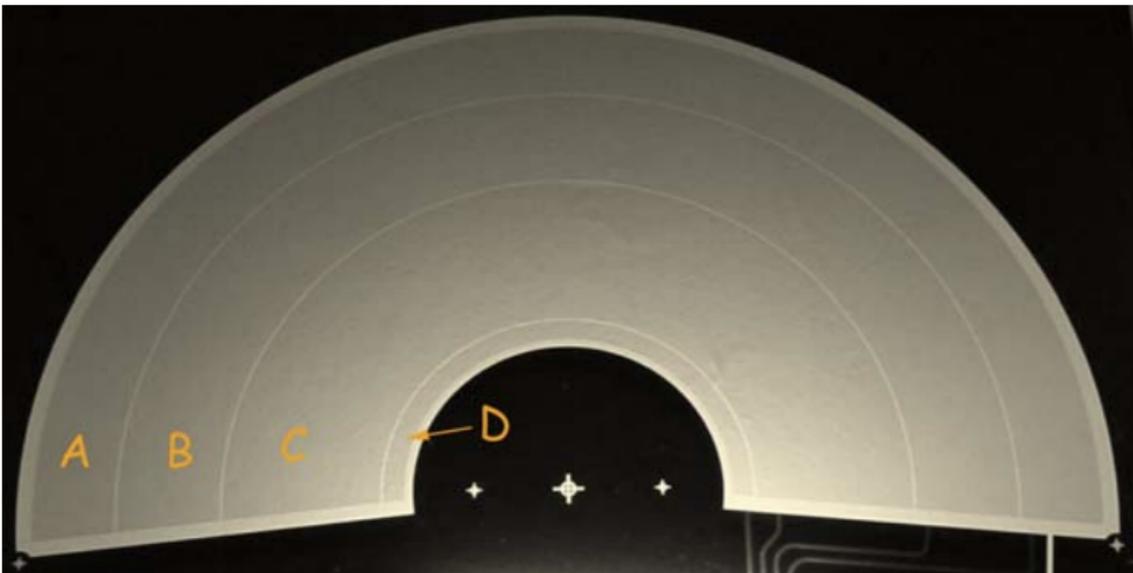


Figure 4.7 – Photo of a TOTEM GEM foil, with the individual HV sectors labeled A through D [70]

However, more important than the similarities to the COMPASS GEMs are the similarities to the CMS GEMs. Crucially, the readout of all TOTEM subsystems was based on the VFAT2 chip, the very same chip used in the v2 CMS GEMs. But where the slice test GEMs saw significant damage to the VFAT2 ASIC (as will be discussed in-depth in Chapter 6), the VFAT2s of the TOTEM GEMs did not. Despite the two detector systems' vastly different designs and locations relative to the interaction point, this perhaps led to a certain false sense of security in which a significant amount of damage-causing discharges was not expected in the CMS GEMs.

4.1.3 LHCb

At the same time as TOTEM's construction, LHCb was equipped with triple-GEM detectors in the innermost region of the first muon station (M1R1), shown in its location within the LHCb detector in Figure 4.8. Here, the GEMs were expected to handle a particle rate of $\mathcal{O}(500 \frac{kHz}{cm^2})$, far greater than the $\mathcal{O}(10 \frac{kHz}{cm^2})$ of GE1/1 [73]. With approximately 20% of trigger muons coming from M1R1, this was the only one of the four muon stations of LHCb to utilize GEMs - the others used the older multiwire proportional chamber (MWPC) technology.

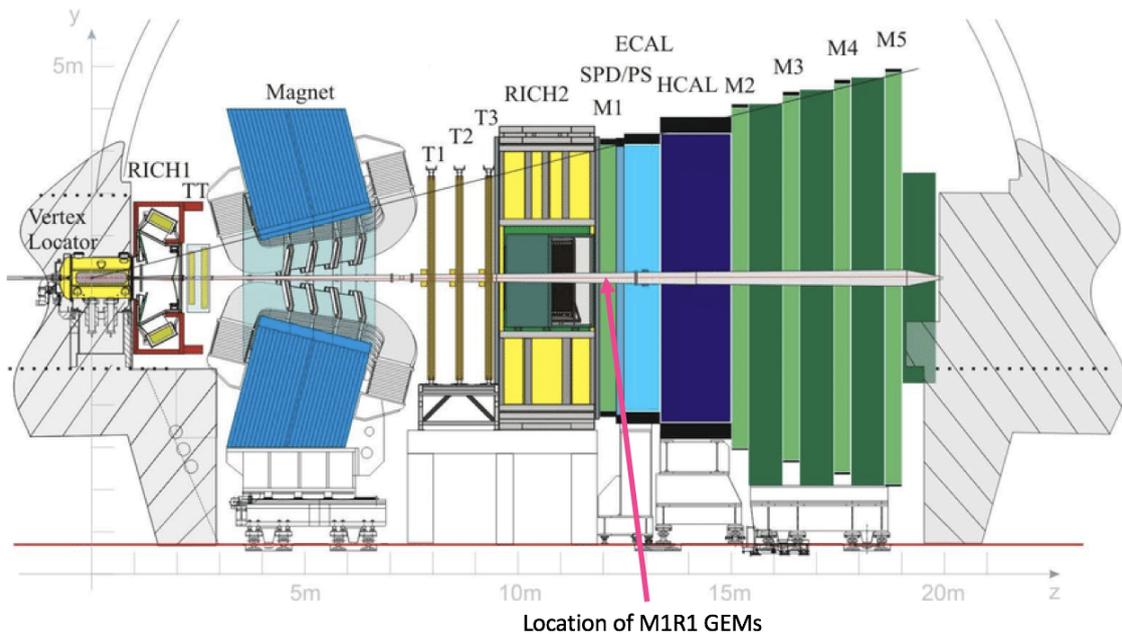
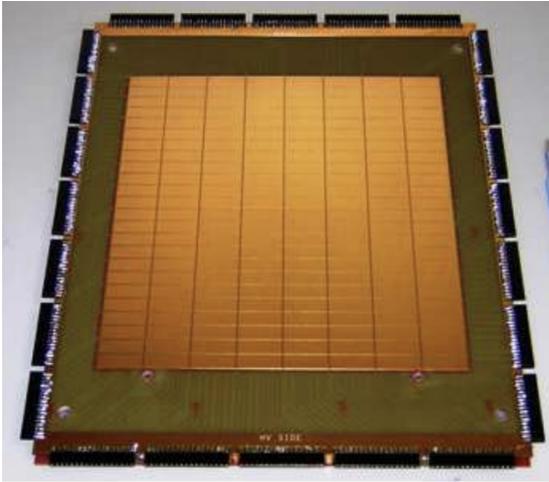


Figure 4.8 – Cross-sectional schematic of the LHCb detector with the various subsystems labeled

Here, rectangular detectors with dimensions 20×24 cm were used, with the bottom of each foil broken down into six HV sectors of approximately 66×240 mm² [73]. Unlike the COMPASS and TOTEM GEMs, the LHCb GEMs used the same 3/1/2/1 mm gap configuration that is used in the CMS GEM detectors, though with a different gas mixture - Ar:CO₂:CF₄ 45:15:40 instead of the more eco-friendly Ar:CO₂ 70:30. And, as with GE1/1, two GEM detectors were coupled together to make one “superchamber”.

The gap configuration and coupling is essentially where the similarities end, however. The LHCb GEMs used rectangular pad-shaped readout electrodes, shown in Figure 4.9a. These pads were read out using the CERN and Rio Current-mode Amplifier (CARIOCA) chip - a chip which was designed specifically for the MWPCs of LHCb's muon detector. As

a result, they were not fully optimized for use with the GEM detectors of M1R1, including ion tail cancellation and baseline restoration circuits which were not necessary for the purely electronics GEM signal and added unnecessary noise into the system [74].



(a) Photo of the pad electrodes of an LHCb M1R1 GEM detector



(b) Photo of two LHCb M1R1 GEM detectors enclosed within a protective box, as they would be installed into the LHCb detector

Figure 4.9 – Photos of an LHCb M1R1 GEM detector [74]

And, as with the COMPASS and TOTEM GEMs, the LHCb GEMs once again utilized a glued honeycomb structure to reinforce the detector on both ends and act as support plates. The LHCb GEMs were also maintained within a strong metal box structure, shown in Figure 4.9b, which gave them a certain additional degree of physical protection, particularly during transport, installation, and LHCb maintenance.

Thus, while these GEMs might give some small insight into the operation of a 3/1/2/1-configured GEM detector, and in particular showed their operations in an extremely high-rate environment, the stark differences between the two systems make direct comparison quite difficult.

4.1.4 Summarizing History

While the above sections represent just a handful of the many members of the GEM family tree, the knowledge of these previous systems allows us to get a better idea for what made the CMS GEMs so unique. Prior to the slice test, a “large” GEM detector was on the order of 30 cm wide - significantly smaller than the meter-long CMS GEMs. These smaller detectors included a honeycomb structure to reinforce both sides of the detector, which was not included in the CMS GEM, and used glue to keep the stretched GEM foils in place rather than the self-stretching technique utilized in the CMS GEMs. Each detector system utilized a sub-sectioning of the foils into smaller HV sectors in order to reduce discharge damage, though the location (top vs. bottom) of the sectioning on the foil and the size of the sectors varied with each experiment, none directly equal to the chosen sectioning of the GE1/1 foils. The 3/1/2/1 gap configuration used in the CMS GEMs made only one prior appearance in the GEMs of the LHCb experiment, though they cannot be directly compared as the LHCb GEMs used a gas mixture of Ar:CO₂:CF₄ 45:15:40 rather than the more eco-friendly mixture of Ar:CO₂ 70:30. And finally, the VFAT2 chip saw only one prior use, with the majority of other experiments preferring the older APV25 readout chip

or the specialized CARIOCA chip.

Thus, while the CMS GEMs could draw on experience piecemeal from previous experiments, there was no truly comparable experiment from which to gather the full picture. Only the slice test, which will be discussed in greater detail in Section 4.4, could truly provide the full understanding of the operations of such a detector within the CMS experiment.

4.2 Motivation for GE1/1

Having now provided the historical context, we move now to the present motivation that brought GE1/1 into being: despite forward RPCs being in the original conception of the CMS muon system, RPCs were not implemented beyond $|\eta| > 1.6$ due to concerns about their ability to handle the high background particle rates [52], which are on the order of $\mathcal{O}(10 \frac{kHz}{cm^2})$. This leaves only the CSCs to cover that high η region and creating a concerning lack of redundancy. In order to rectify this, a call was put out for a muon subsystem which was sufficiently compact enough to fit within the limits of the space originally planned for the forward RPCs, while being robust enough to handle the high background particle rates of the region, ideally operating well at particle fluxes far above those expected in the region after the implementation of the HL-LHC. This call was satisfied by GEM detectors in what is now known as the GE1/1 project, whose location within CMS can be seen in red in Figure 4.10. These detectors are intended to serve an important role in improving both the tracking and trigger performance of the CMS muon system and CMS detector as a whole.

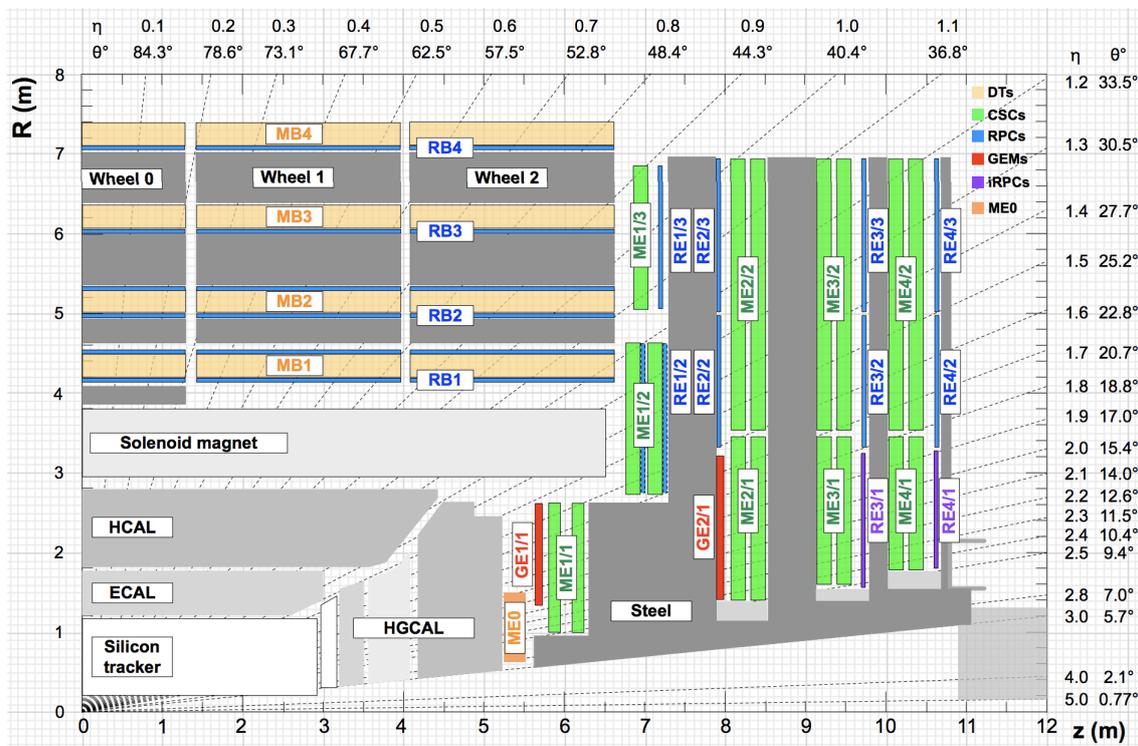


Figure 4.10 – Cross-section of CMS, detailing the subdetector locations, including the new and future GEM systems [7]

The bending of muon tracks within CMS is greatest at the location of the first muon

station, lessening at subsequent stations due to the bending of the magnetic field lines in the endcap flux return. Because of this magnetic field reduction and the previously-mentioned high background rates at higher η values, the contribution to the trigger rate within the GE1/1 coverage area of $1.61 < |\eta| < 2.18$ is particularly large and difficult to control. The GE1/1 chambers, in conjunction with the existing ME1/1 CSC chambers, effectively multiply the muon's path length within the first muon station by 2.4 - 3.5 compared to that of the ME1/1 chambers alone. They also allow for the exploitation of the measurement of the bending angle at trigger level between the two stations in an area where muons emerge at an angle of around 10° relative to the beam axis. A simulation of this measurement is shown for odd-numbered GEM chambers for muons of two different momenta - $p_T = 20 \frac{\text{GeV}}{c}$ (so-called "hard muons") and $p_T = 5 \frac{\text{GeV}}{c}$ (so-called "soft muons") - in Figure 4.11.

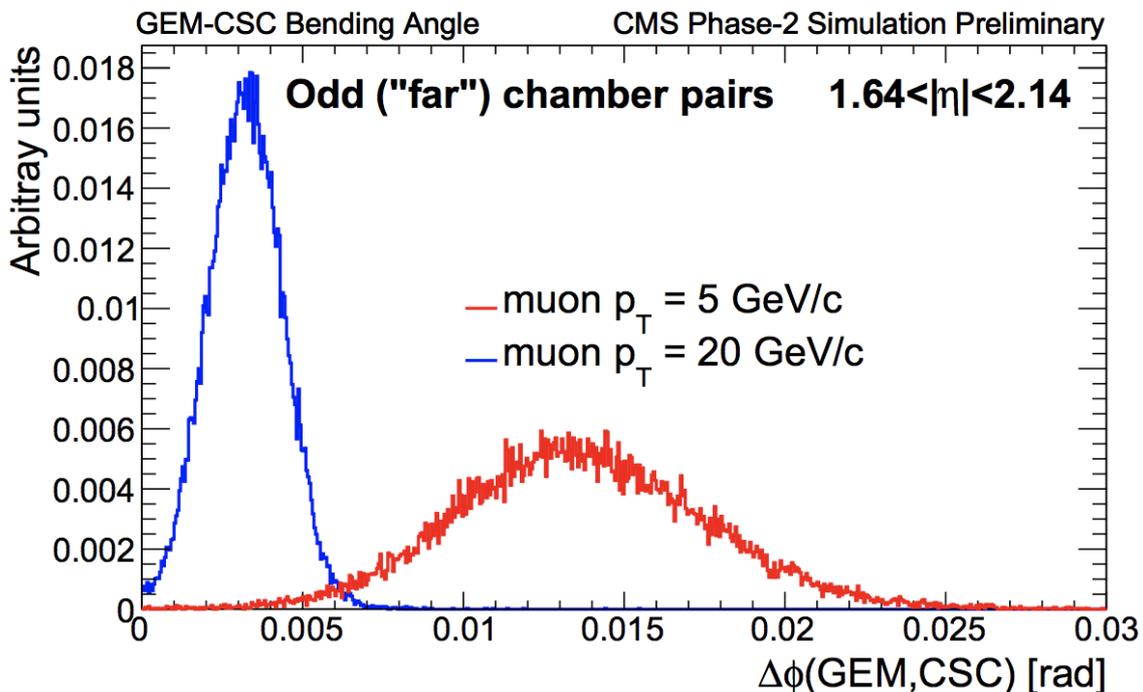


Figure 4.11 – Simulation of the GEM-CSC bending angles for odd-numbered GEM chambers for soft and hard muons [75]

This significantly improves the momentum resolution of the L1 standalone muon trigger and drastically reduces its disproportionately-large contribution to the overall L1 muon trigger rate. This improvement can be seen in Figure 4.12.

By allowing the L1 muon trigger thresholds to be maintained at low p_T values, the GE1/1 project will increase the efficiency of capturing interesting physics processes featuring soft leptons. Taking, for example, a single muon trigger, the GE1/1 upgrade will allow the preservation of the L1 threshold at 12 - 14 GeV, which in turn will provide nearly full efficiency for offline muons with $p_T > 18 - 20$ GeV. The maintaining of the low p_T muon trigger thresholds is crucial for a broad array of physics studies, including Higgs measurements and searches for new physics. For the latter, split supersymmetry and anomaly-mediated supersymmetry searches are often dependent on the ability to trigger on soft leptons, particularly where "compressed" mass spectra are involved. For the former, one interesting example is the study of Higgs coupling to third generation leptons via $H \rightarrow \tau\tau$. Although there are many decay channels, the semileptonic channel $\tau\tau \rightarrow \mu\tau_h + X$

holds particular importance due to its clean signal and large branching fraction. However, given the low average lepton p_T , the measurement of this channel relies on the ability to trigger on the events efficiently. Simulation studies have shown that if the muon p_T threshold is lowered by just 5 GeV (from 25 to 20 GeV), the kinematic acceptance for this channel will increase by as much as 35%. This would be impossible without the addition of the GE1/1 detector system, as muon trigger studies which were performed for the CMS Phase I upgrade showed that achieving an acceptable L1 trigger rate for muons with $p_T < 25$ GeV after LS2 would not be possible without substantial efficiency losses in the endcaps [52].

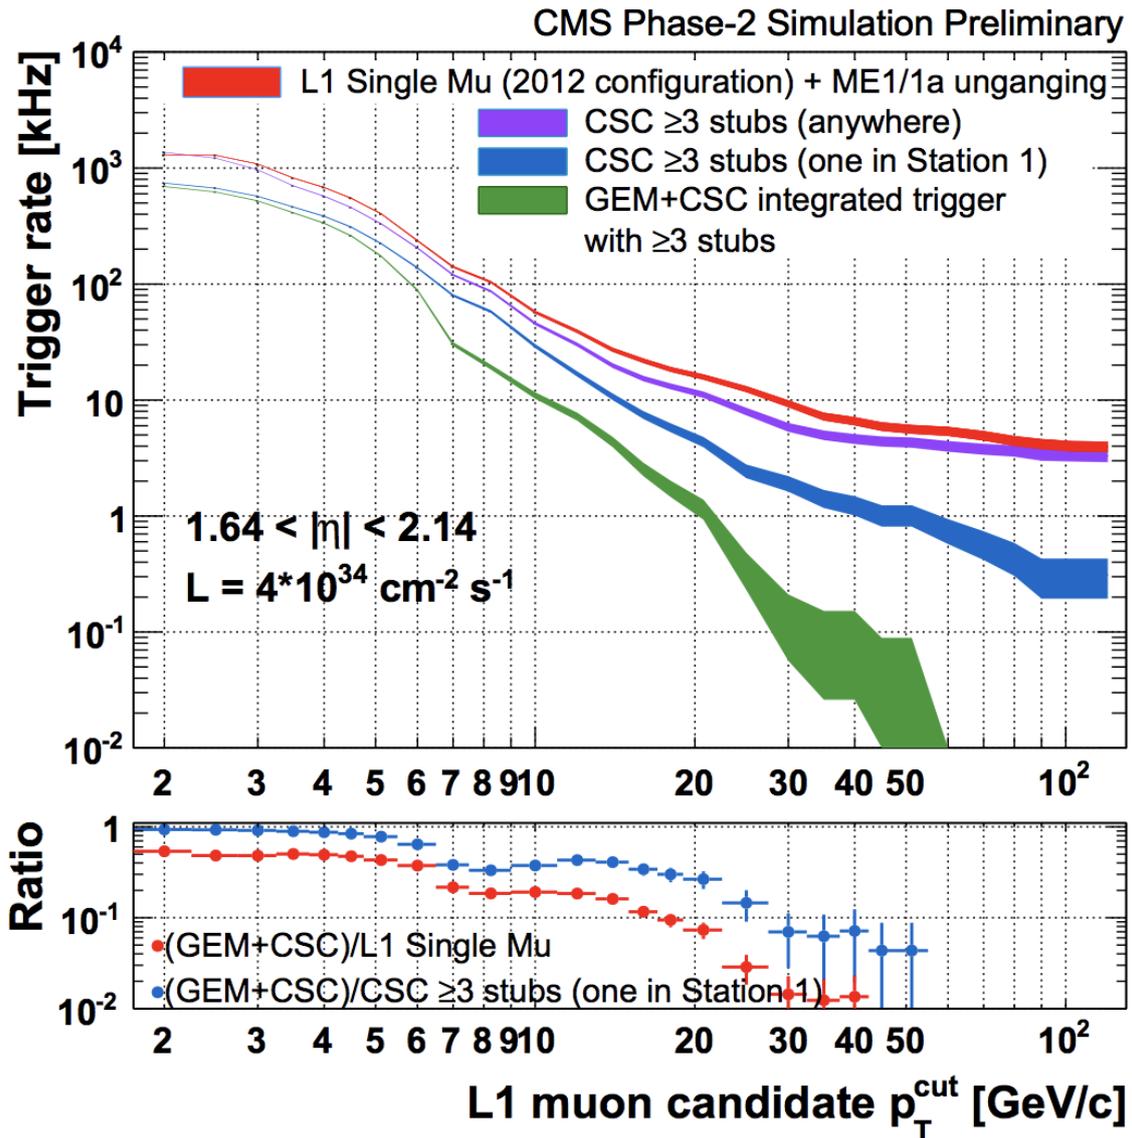


Figure 4.12 – Simulation of the L1 muon trigger rates before and after the GE1/1 upgrade at a luminosity of $4 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ [75]

Although the GE1/1 detectors will be most directly interfaced with the CSCs, they also play a key role in other detector systems. For example, during Long Shutdown III (LS3), the new silicon tracker and track trigger will be commissioned. They will then be used in coincidence with the L1 muon trigger to form a so-called “combined muon trigger”, in which the momentum resolution for most muons from the primary vertex will be determined by the resolution achieved by the track trigger. GE1/1 will be used for

position matching with the track trigger, so that the standalone muon trigger will run in parallel with the combined muon trigger, albeit at a much higher p_T threshold. This will allow the standalone muon trigger to provide a high efficiency for displaced muons and other exotic particles, in addition to serving as a backup to the combined muon trigger.

4.3 GE1/1 Detector System

The GE1/1 system is composed of 144 GEM detectors, arranged as 72 two-detector superchambers called “GEMINI”¹. Each endcap of CMS (referred to as the positive and negative, or +1 and -1, endcaps) contains 36 of these superchambers, as shown in Figure 4.13.

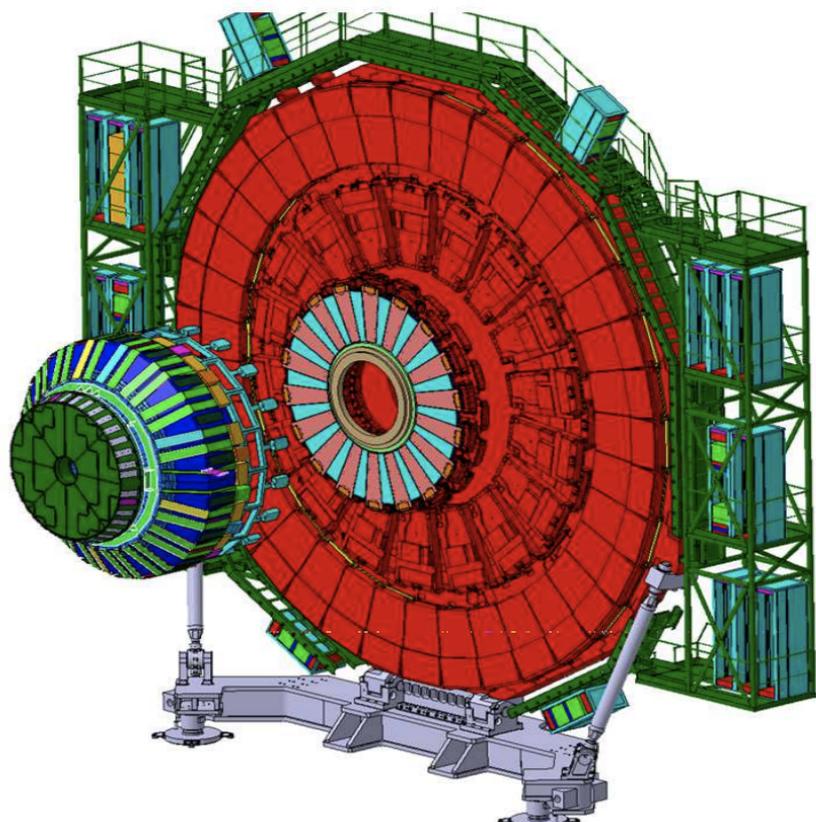


Figure 4.13 – Positions of the GE1/1 superchambers within one of the CMS muon endcaps [52]

As the GE1/1 detectors are the first completely new subdetector system to be introduced into CMS where there previously was only empty space since its construction, they must rise to meet unique challenges. One of these challenges is the constraints imposed by the existing mechanical structure of CMS. As mentioned in Section 4.2, the GEMs were chosen for inclusion into CMS specifically due to their compact nature, as the total empty space left in front of the CSCs was limited and could not be increased. This space can be seen in Figure 4.14, denoted by the purple arrow. Here, the HE back-flange, on which the GE1/1 superchambers are installed, is located 5,674 mm away from the interaction point,

¹This naming scheme was actually first suggested by the author in July 2014, as her birthdate of June 12th makes her a Gemini astrologically. Although the post-doc she suggested this to never propagated the suggestion onwards, and thus she has never been credited for it, nevertheless the naming scheme was chosen, and this is a particular point of pride for her.

and the empty space extends as far as 5,541 mm away from the interaction point, for a total of 133 mm, or 13.3 cm. Thus, the full width of the detector had to be less than 13.3 cm, including any margins required for a smooth installation.

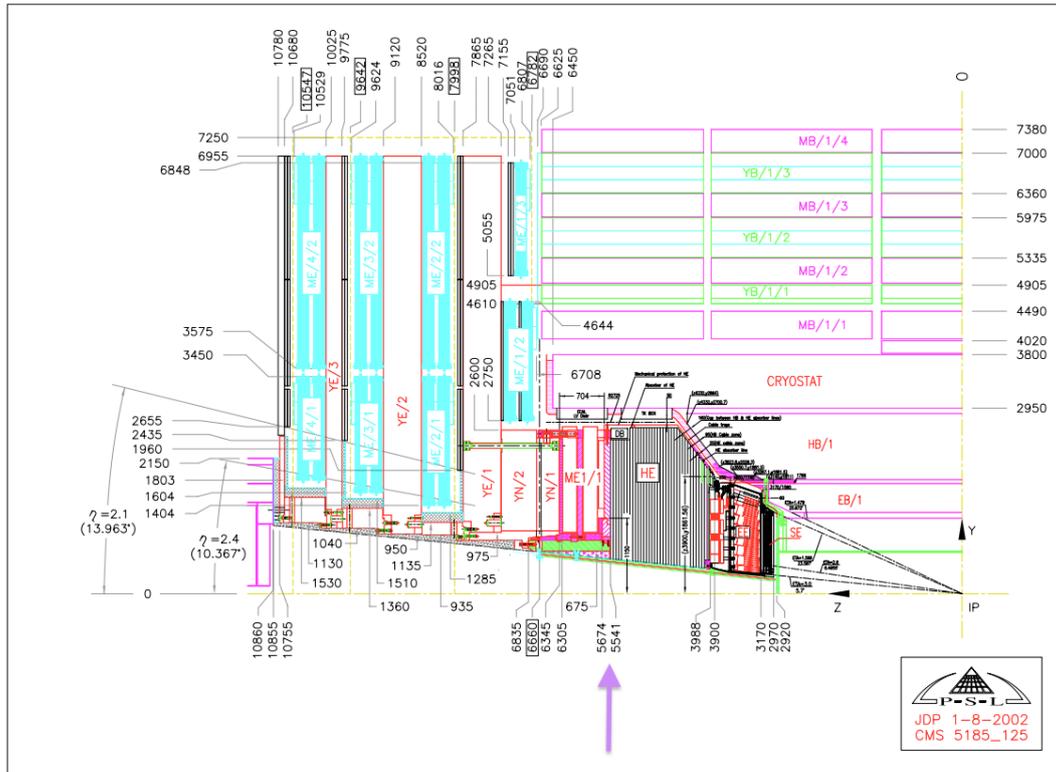


Figure 4.14 – Blueprint of the CMS endcap as the configuration was set in 2002, detailing the space available to the CMS GE1/1 system (as denoted by the purple arrow) [76]

This width limitation was not the only physical constraint, however. In Figure 4.13, the superchambers are colored in an alternating pattern of red and blue, to indicate a difference in length. The blue superchambers, called “short” superchambers, have a geometrical acceptance of $\eta = 1.61 - 2.18$, whereas the red superchambers, called “long” superchambers, have a geometrical acceptance of $\eta = 1.55 - 2.18$. This equates to a physical difference down the center line of the detector of 148.16 mm. The reason for this is better seen in Figure 4.15. Here, we see that the slots available for installation of the GE1/1 detectors are discretized, with steel supports in-between each slot. These supports extend slightly inwards towards the center of the ring, thus reducing some of the available space for the detectors length-wise. The short detectors compensate for this through their length reduction, whereas the long detectors make use of all of the available space that is afforded them due to being in the open area of the slots. This structure does make it slightly harder to access the short chambers once installed, though their patch panels are still available to members of the team with smaller hands². Note that Figure 4.15 shows empty space in each slot where the CSC ME1/1 chambers are installed - this space is not available to the GEM detectors.

²Women of STEM, rejoice, for your delicate ladyhands may find their true purpose in accessing tight detector spaces that are unavailable to beefy menfolk.

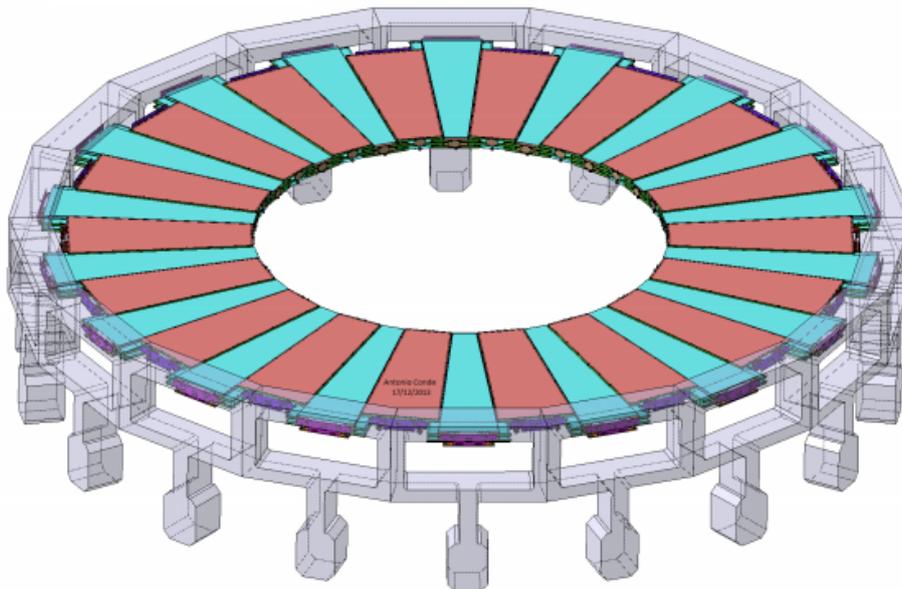


Figure 4.15 – Figure detailing the locations of the short and long GE1/1 superchambers and the mechanical structure of CMS which necessitates the two detector lengths [52]

In addition to meeting these physical construction constraints, the GE1/1 detectors had to meet and/or exceed a strict set of fundamental criteria in order to fulfill the needs of CMS at their installation location. These include [28]:

- A particle rate capability of up to $4.5 \frac{kHz}{cm^2}$
- A per-chamber detection efficiency higher than 97% for detecting minimally-ionizing particles
- An angular resolution of $300 \mu rad$ for $\Delta\Phi = \Phi_{GE1/1} - \Phi_{ME1/1}$
- A timing resolution of 8 - 10 ns per detector
- A gas gain uniformity of 30% across a single detector and between detectors
- The ability to tolerate an integrated charge of $\sim 18 \frac{mC}{cm^2}$ without suffering damage or performance loss
- A discharge rate which does not impede performance or operations

This latter criterion will be the focus of the discussion of the remaining chapters of this dissertation, as we will see that the detectors of the GE1/1 demonstrator system very much did not fulfill this requirement.

The basic outline for the design specifications of the GE1/1 detectors can be found in Table 4.1, as taken from the CMS Technical Design Report for the Muon Endcap GEM Upgrade.

Specification / Parameter	GE1/1
Detector technology	Micropattern gas detector (MPGD)
Charge amplification element	GEM foil (triple, cascaded, tensioned at $\sim 5 \frac{N}{cm}$)
Number of chambers in overall system	144 [72 (2×36) in each endcap]
Chamber shape (active readout area)	Trapezoidal; opening angle 10.15°
Active area overlap in adjacent chambers	2.6 mrad (5.7 readout strip pitches)
Short chamber dimensions (active volume)	L: 106.1 cm (center line), W: (23.1 - 42.0) cm, D: 0.7 cm
Long chamber dimensions (active volume)	L: 120.9 cm (center line), W: (23.1 - 44.6) cm, D: 0.7 cm
Total chamber thickness	H: 3.5 cm
Active readout area	0.345 m^2 (short); 0.409 m^2 (long)
Active chamber volume	2.6 liters (short); 3 liters (long)
Radial distance from beam line	130.2 cm (at inner edge of active readout area)
Geometric acceptance in $ \eta $	1.61 - 2.18 (short); 1.55 - 2.18 (long)
Signal readout structure	Truly radial readout strips
Readout strip dimensions	$230 \mu\text{rad}$ angular strip width; $463 \mu\text{rad}$ angular pitch
Number of η -segments in readout	8
Number of readout strips per η -segment	384
Number of readout strips per chamber	3,072
Counting gas mixture	Ar:CO ₂ 70:30
Nominal operational gas flow	1 chamber volume per hour
Number of gas inlets	1
Number of gas outlets	1
Nominal HV applied to drift electrode	3200 V (Ar:CO ₂)
Nominal operational gas gain	$1 - 2 \times 10^4$
Demonstrated rate capability	$100 \frac{MHz}{cm^2}$

Table 4.1 – Main specifications and parameters for the design and operation of GE1/1 [52]

The first thing of note in this table is that although the chambers alternate between short and long, the opening angle for both types of chambers is the same - 10.15° - which is achieved by the same narrowest width for both chamber types (23.1 cm in terms of active volume, 27.9 cm in total). Similarly, the height of both chamber types (denoted as D, for depth) is also the same, with an active volume thickness of 0.7 cm and a total thickness of 3.5 cm. Not listed is the distance between the two chambers of a given superchamber, 3.7 mm, which is maintained with mounting brackets attached to both the narrow and wide ends of both chambers.

The second thing to note is the “active area overlap in adjacent chambers” of 2.6 mrad. This is achieved by alternating the orientation of the chambers, with the top chimney of the short chambers facing towards the interaction point, and the top chimney of long chambers facing away from the interaction point. The aluminum chimneys which cover

both types of chambers can then interface with each other, sliding into the grooves of the adjacent chambers. This chimney interfacing can be seen in Figure 4.16, which shows two real-sized dummy short chambers and one real-sized dummy long chamber, illustrating how the detectors overlap for the CMS Technical Design Report for the Muon Endcap GEM Upgrade (Reference [52]).

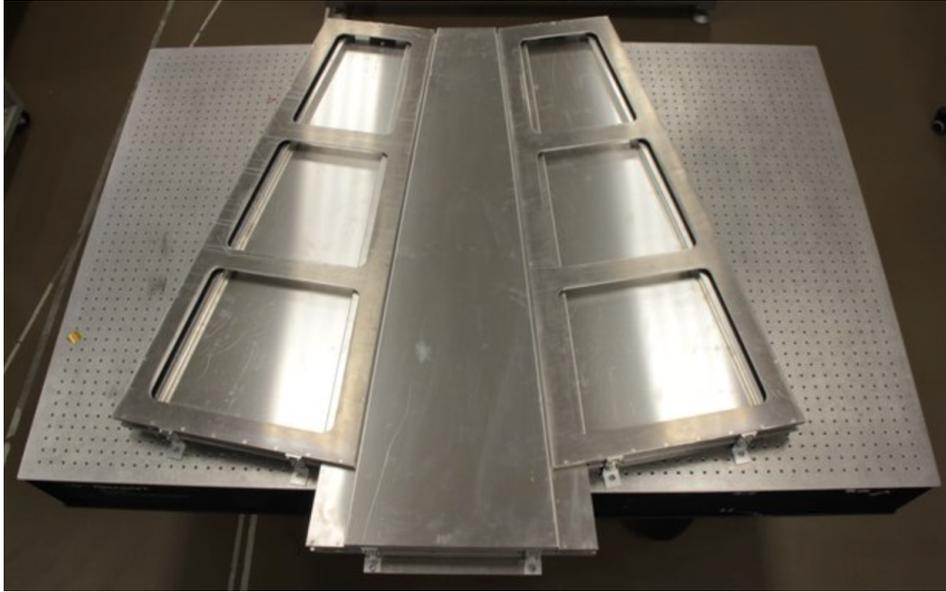


Figure 4.16 – Photo taken of two real-sized dummy short chambers and one real-sized dummy long chamber, illustrating how the detectors overlap within the CMS endcap [52]

This permits some overlap of the chambers, equal to 5.7 readout strips, so as to maximize coverage in the detector and not permit any detection gaps in the ring of chambers. In order to keep labeling consistent, when referring to a given chamber within a superchamber, layer 1 (L1 when in a chamber name, not to be confused with the L1 trigger) is the chamber which is closer to the interaction point of CMS and layer 2 (L2) is the chamber which is further from the interaction point, regardless of the orientation of the superchamber.

Towards the bottom of Table 4.1, we see the gas mixture of choice for the GE1/1 system is Ar:CO₂ 70:30. Although this was the mixture which was ultimately decided on, it is not the only mixture which was considered. Triple-GEMs were in use at the LHCb experiment during the 2010-2012 data taking period of Run I at the LHC, these GEMs utilized a gas mixture of Ar:CO₂:CF₄ 45:15:40. As such, there was precedence for its use, and a significant amount of R&D was done using this mixture. The 40% CF₄ allows for a high drift velocity while maintaining a small Lorentz angle, making it an ideal gas in this respect. A deeper look into the properties of this gas mixture within gaseous detectors and the studies that were performed with it for the GE1/1 project can be found in Reference [52]. However, despite its useful properties, ultimately the gas mixture was rejected for one crucial reason: CF₄ is a highly potent greenhouse gas. With an atmospheric lifetime of 50,000 years and a greenhouse warming potential of 6,500 (compared to CO₂, with a greenhouse warming potential of 1), any leaks in the GEM detectors, while inconvenient for the project, would have disastrous ramifications for the planet. As the leaking of potent greenhouse gases from both the CMS CSCs and RPCs is already a significant concern, the decision was made for Ar:CO₂ to be used in this and all future CMS GEM systems, a

decision for which the project should be commended.

4.4 The Slice Test

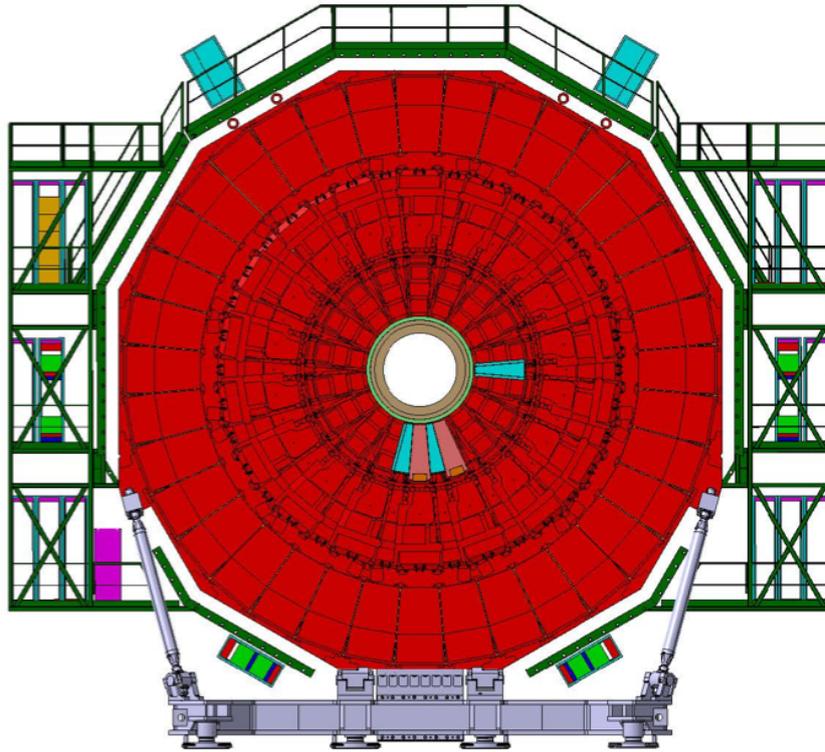


Figure 4.17 – Positions of the slice test superchambers

In preparation for the full installation of the GE1/1 detector system, in January 2017, a demonstrator system called the “slice test” was installed into CMS. This demonstrator system consisted of ten triple-GEM detectors, arranged as five two-detector superchambers. These superchambers followed a naming scheme of “GEMINI” to indicate they were a two-detector superchamber, “m” for module, and a number which was assigned based on its location within the endcap, with “01” starting at the 3 o’clock position and incrementing by 1 per detector position, moving counter-clockwise around the ring when facing the nose of the endcap until a maximum value of 36. Two of these superchambers, GEMINIm27 and GEMINIm29, were short superchambers, and two, GEMINIm28 and GEMINIm30 were long superchambers, and were located at the 6 o’clock position, as seen in Figures 4.17 and 4.19. An additional short superchamber, GEMINIm01, was added in the 3 o’clock position.



Figure 4.18 – Photo of me in a cherry picker during the installation and commissioning of the slice test detectors in January 2017. As the only student certified to drive mobile elevated work platforms, I assisted in the installation of the optical communication fibers used to transmit data from the slice test detectors to the back-end electronics.

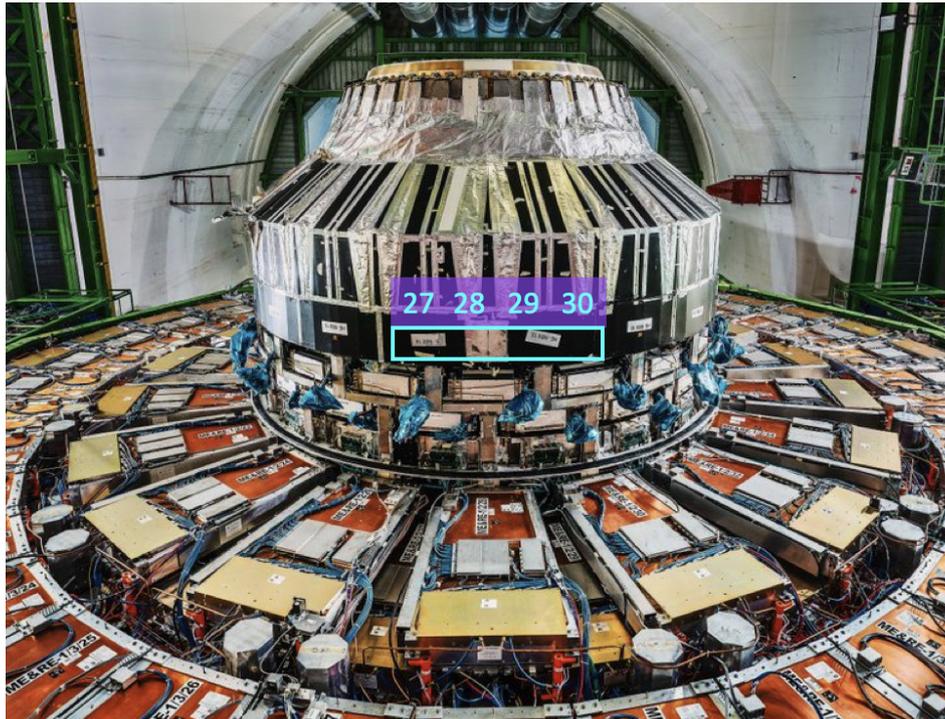


Figure 4.19 – Photo of the positions of GEMINIm27-30 within the CMS forward muon endcap, as seen from below

Due to delays in production, the final v3 electronics had not been available at the time of installation. As such, these ten original slice test detectors consisted of v2 electronics, as described in Section 3.

The eight detectors at the 6 o'clock position operated with the standard single-channel HV power supply (CAEN A1526N) each, the output from which was distributed to the various foils through an HV divider located at the wide end of the chambers. By operating these eight chambers with a single-channel power supply, no additional cabling was necessary, as a sufficient number of excess HV cables had already been routed into the CMS endcap in view of a potential upgrade of the RPC system in the area [77]. The CAEN A1526N power supply has been used successfully in a wide variety of high energy physics experiments and its behavior is well-known. However, it has its drawbacks - namely, that by using a ceramic HV divider, the voltage ratios between foil layers are fixed, disallowing the user from individually adjusting the voltage applied to a single foil in case of problems, or even turning off a single foil layer. In addition, the current which can be observed through the power supply is largely determined by the equivalent resistance of the HV divider, on the order of $700 \mu\text{A}$. This prevents the observation of current fluctuations which may be occurring in the detector, one of the main interests of the slice test, as described later in this section.

GEMINIm01, on the other hand, operated with a multi-channel power supply (CAEN A1515TG) in an attempt to more closely replicate the final electronics setup for the GE1/1 detectors. Unlike the CAEN A1526N module, the CAEN A1515TG module was specifically designed for use with the CMS triple-GEM detectors, as it provides 14 HV channels in order to independently power the seven electrodes of two triple-GEM detectors, or one superchamber. These 14 channels can be seen in the A1515TG's block diagram in Figure 4.20. These two groups of seven channels are internally stacked, allowing the user to set the desired voltage of a given channel with respect to the previous channel, rather than an

absolute value, which is automatically calculated by the CAEN module.

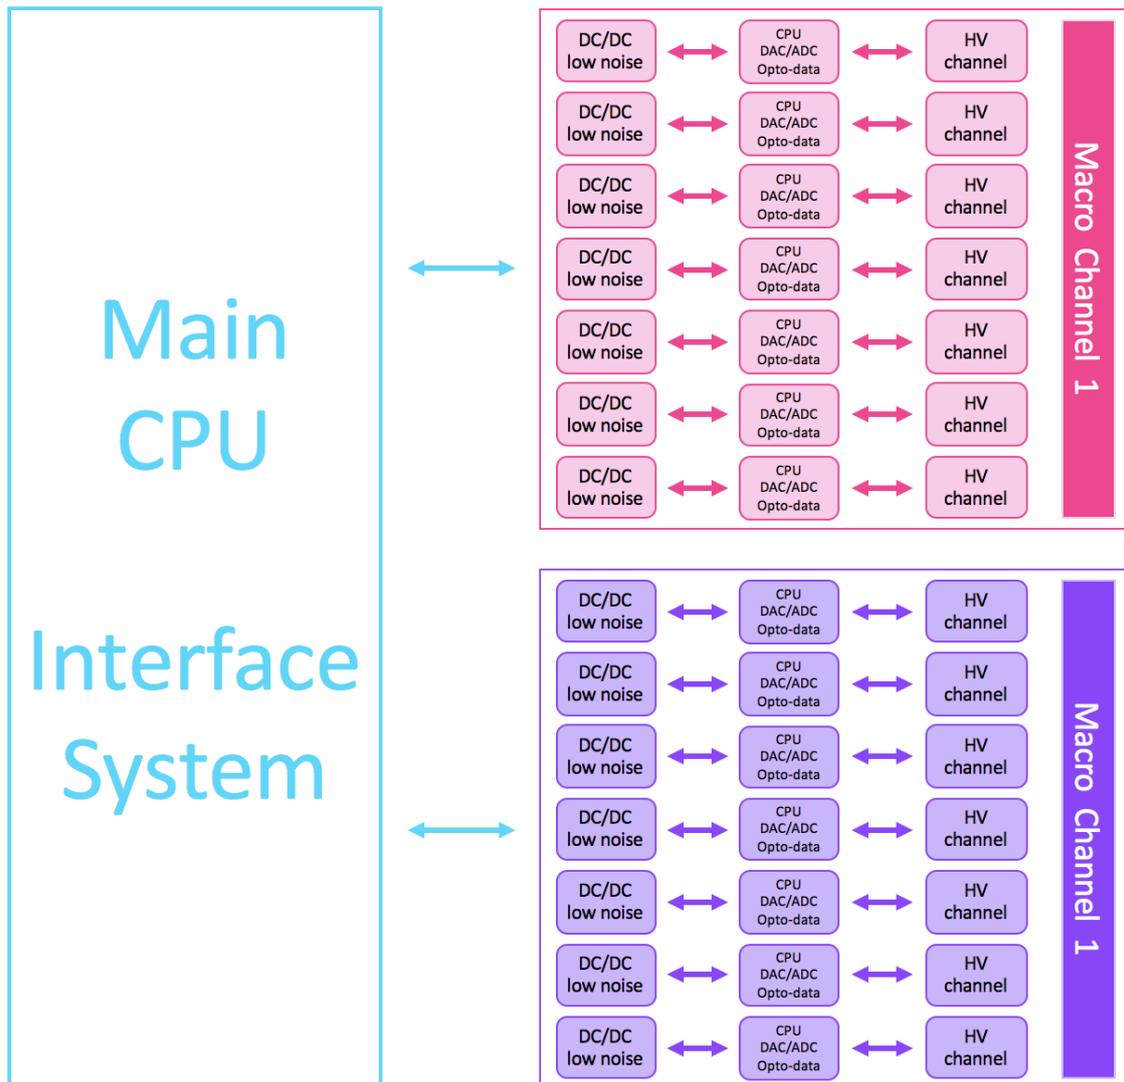


Figure 4.20 – Simplified block diagram of the CAEN A1515TG power supply, as provided by CAEN engineer Alessandro Iovene

For example, say a user wants to apply 350 V to the $G3_{bottom}$ foil and 400 V to the drift field between $G3_{bottom}$ and the readout board. The user can set those values to their respective channels, and the module will automatically apply 750 V to $G3_{top}$. If the user later wishes to turn off the drift field, they can set the drift channel to 0 V and the module will automatically apply $G3_{top} = 0$ V and $G3_{bottom} = 350$ V. Thus, the user can operate individual channels without needing to manually take the interdependence of the voltages of the different channels into account, minimizing human error that could potentially damage the detector or otherwise not provide the desired result [77].

The CAEN A1526N module also has the benefit compared to the single-channel CAEN A1526N module in that the current observed through it is nearly zero, which allows for the observation of discharge events within the detector and on each individual foil. This ability will be discussed in much greater detail in Chapter 6.

In March 2018, GEMINIm01 was replaced with a v3 chamber, also called GEMINIm01. To differentiate between the v2 and v3 GEMINIm01 superchambers, we can refer to each

slice test chamber by its official production designation, which can be found in Table 4.2, or else simply call them GEMINIm01_{v2} and GEMINIm01_{v3}. This replacement offered two distinct benefits to the project. For one, the inclusion of a superchamber with the “final” electronics allowed us to confirm the operational properties of those electronics within CMS, and ensure that everything would work to standard. For another, it allowed for the removal of GEMINIm01_{v2} midway through the slice test, providing the opportunity to examine a detector after a year’s presence within CMS and look for any damage caused by that presence. This would, in fact, indirectly lead to the study which the majority of this dissertation is based upon and provide insight into the conclusions of that study, as will be seen in Chapter 6.

Common Name	Official Name
GEMINIm27L1	GE1/1-VII-S-CERN-0003
GEMINIm27L2	GE1/1-VII-S-CERN-0004
GEMINIm28L1	GE1/1-VII-L-CERN-0004
GEMINIm28L2	GE1/1-VII-L-CERN-0002
GEMINIm29L1	GE1/1-VII-S-CERN-0005
GEMINIm29L2	GE1/1-VII-S-CERN-0006
GEMINIm30L1	GE1/1-VII-L-CERN-0001
GEMINIm30L2	GE1/1-VII-L-CERN-0003
GEMINIm01L1 _{v2}	GE1/1-VII-S-CERN-0002
GEMINIm01L2 _{v2}	GE1/1-VII-S-CERN-0001
GEMINIm01L1 _{v3}	GE1/1-X-S-CERN-0005
GEMINIm01L2 _{v3}	GE1/1-X-S-CERN-0006

Table 4.2 – Slice test chamber names and official designations³

Although the initial 10 detectors did not represent the final electronics foreseen for the GE1/1 detectors, the slice test provided a valuable opportunity to test the GEM detectors within the working environment of CMS, operating them under the harsh radiation conditions that are present in the forward CMS endcap when beam is present. Such harsh conditions are not readily recreated in a laboratory environment outside of occasional test beams, and even so not to nearly the same degree. Thus, despite intensive R&D campaigns on the detectors in the years prior, only the slice test could truly provide proof of concept for the GE1/1 project. The slice test, therefore, was undergone with the following distinct goals in mind:

- Prove that the GEM detectors fulfilled the requirements set out in Section 4.3, particularly the ability to tolerate an integrated charge of $\sim 18 \frac{mC}{cm^2}$ without suffering damage of performance loss, and the presence of a discharge rate which does not impede performance or operations.
- Demonstrate an operational stability which is compatible with CMS global running
- Test the combined GE1/1-ME1/1 GEM-CSC integrated local trigger

³Note that for both layers of GEMINIm01_{v3}, the numbers may be reversed, i.e. GEMINIm01L1_{v3} may actually be GE1/1-X-S-CERN-0006 and GEMINIm01L2_{v3} may actually be GE1/1-X-S-CERN-0005. This is due to the fact that the superchamber was removed from CMS and the chambers were decoupled and removed of their aluminum chimneys prior to their identities being noted.

4.4.1 Slice Test Operations

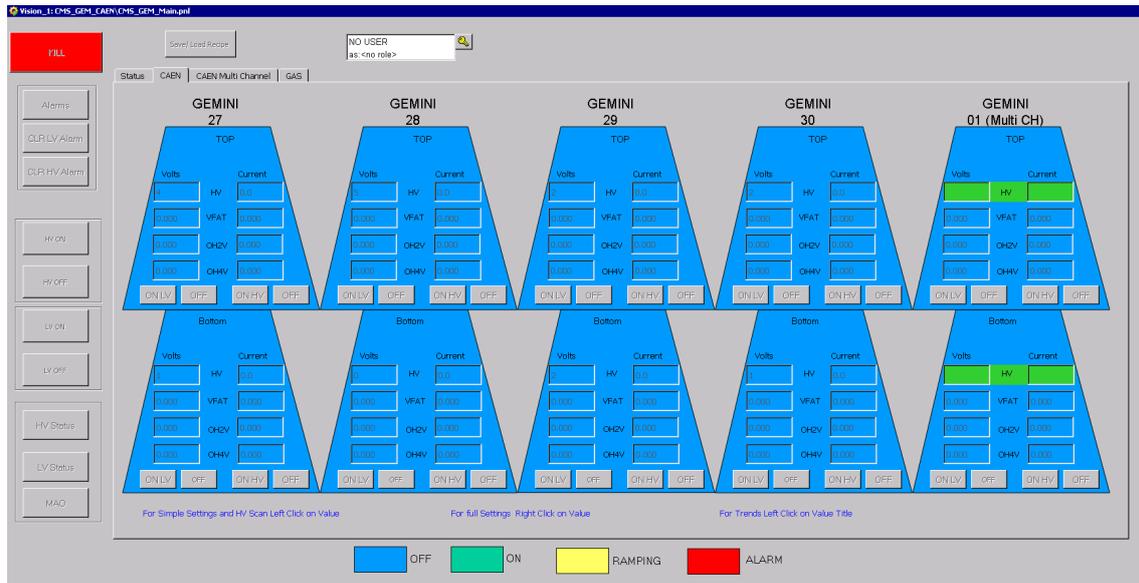


Figure 4.21 – Screenshot of the main page of the 2018 slice test DCS, indicating the LV and HV is off for all slice test detectors except for potentially the HV of GEMINI_{m01}, which, being a multi-channel power supply, requires a nested control panel within the main page

In order to accomplish the goals of the slice test, a team of dedicated shifters, known as detector experts on-call (DOCs), were trained in the daily operations of the slice test detectors. These shifters, of which I was one, were the primary operator and point of contact for all operational matters involving the GE1/1 slice test. The on-call portion of the name stems from the fact that during their week-long shift, a GEM DOC was responsible for any operational problems with the detectors 24 hours a day, and would be expected to take control of debugging the system and making critical decisions at any hour of the day or night.

The GEM DOC's most important tool was the slice test detector control system (DCS). This set of control panels, hosted on the main GEM controlling computer at P5, consisted of all of the necessary control and monitoring mechanisms for the GEM DOC to operate the detector system. The main panel of the slice test DCS, as it was in 2018 after the installation of GEMINI_{m01,v3}, can be seen in Figure 4.21. From here, the DOC could control the LV and HV settings of each detector individually, turning them on and off as required.

For the v2 detectors, HV was controlled with a simple application of a single voltage and current value, and this could be performed from the panel shown. For the v3 detectors, this process was slightly more complicated, as the voltages and currents were delivered via a multi-channel power supply to each foil individually. In order to change these values, the DOC would click on the HV box of either GEMINI_{m01,v3} detector (shown in Figure 4.21 as being highlighted in green) to bring up a more detailed HV panel, not shown.

In the case of the multi-channel detectors, since the voltages supplied to individual layers could be controlled separately, a power on/off sequence was created, such that they do not ramp up or down concurrently. When one of the GEMINI_{m01} layers is powered on or the total voltage is increased, the voltages ramp up in the following order: $G3_{bottom}$, $G2_{bottom}$, $G1_{bottom}$, drift, $G3_{top}$, $G2_{top}$, $G1_{top}$. By powering the GEM foils ($G_{\#bottom}$) first

and then adjusting the drift fields ($G_{\#top}$) allows the detector to be ready for physics faster than it otherwise would. On the contrary, when one of the GEMINIm01 layers is powered off or the total voltage is decreased, the voltages are ramped down in the exact opposite order. This brings the detector to a state without charge amplification faster than it otherwise would [77]. During ramping, it is recommended to change the I0 limit to a greater value than it is set to during normal operations, in order to prevent the detector being tripped by any large current spikes which may occur during the process.

For either the v2 or v3 detectors, the LV was controlled via three sets of voltage and current boxes on this main panel. These three sets are labeled VFAT, OH2V and OH4V. The former, VFAT, was the LV required to power the 24 VFAT2 front-end readout hybrids, while the latter two, OH2v and OH4v, were the LV required to fully power the v2 optohybrids. Typical settings/readings for these detectors can be found in Tables 4.3, 4.4, 4.5, and 4.6.

Detector	HV (V)	VFAT (V)	OH2V (V)	OH4V (V)
GEMINIm27L1	3755.0	2.50	2.00	3.41
GEMINIm27L2	3675.0	2.51	2.01	3.41
GEMINIm28L1	3577.0	2.51	2.03	3.41
GEMINIm28L2	3626.0	2.51	2.02	2.09
GEMINIm29L1	3776.0	2.51	2.01	3.39
GEMINIm29L2	3701.0	4.50	1.99	3.38
GEMINIm30L1	3650.0	2.49	2.19	3.29
GEMINIm30L2	3627.0	2.49	2.00	3.04

Table 4.3 – Typical readings of the HV and LV voltage settings of the single-channel slice test detectors, as taken from the DOC daily elog on July 1st, 2018⁴

Detector	HV (μ A)	VFAT (A)	OH2V (A)	OH4V (A)
GEMINIm27L1	701.2	6.45	1.96	3.12
GEMINIm27L2	699.8	6.20	1.99	3.13
GEMINIm28L1	672.9	6.29	1.86	3.15
GEMINIm28L2	677.6	6.47	2.09	3.19
GEMINIm29L1	712.5	6.78	1.86	3.13
GEMINIm29L2	692.4	6.51	1.91	3.15
GEMINIm30L1	688.0	6.31	2.19	3.29
GEMINIm30L2	682.7	6.47	2.03	3.18

Table 4.4 – Typical readings of the HV and LV current settings of the single-channel slice test detectors, as taken from the DOC daily elog on July 1st, 2018

⁴<http://cmsonline.cern.ch/cms-elog/1050445>

Detector	GEMINIm01L1 _{v2} (V)	GEMINIm01L2 _{v2} (V)	GEMINIm01L2 _{v3} (V)
Drift	528.84	528.76	751.72
G1Top	263.26	263.02	373.84
G1Bottom	205.84	205.92	292.46
G2Top	258.46	258.46	367.18
G2Bottom	411.16	410.68	584.44
G3Top	246.9	246.58	350.20
G3Bottom	293.94	293.56	417.26
VFAT	2.51	2.51	5.99
OH2V	2.02	2.00	0.0
OH4V	3.40	3.39	0.0

Table 4.5 – Typical readings of the HV and LV voltage settings of the multi-channel slice test detectors, as taken from the DOC daily elogs on June 30th, 2017⁵ and July 1st, 2018

Detector	GEMINIm01L1 _{v2}	GEMINIm01L2 _{v2}	GEMINIm01L2 _{v3}
Drift (μA)	0.03	0.02	0.09
G1Top (μA)	0.01	0.02	0.03
G1Bottom (μA)	0.00	0.11	0.01
G2Top (μA)	0.01	0.01	0.02
G2Bottom (μA)	0.01	0.03	0.04
G3Top (μA)	0.02	0.01	0.06
G3Bottom (μA)	0.02	0.05	0.01
VFAT (A)	5.98	6.15	2.05
OH2V (A)	1.99	2.28	0.0
OH4V (A)	3.11	3.25	0.0

Table 4.6 – Typical readings of the HV and LV current settings of the multi-channel slice test detectors, as taken from the DOC daily elogs on June 30th, 2017 and July 1st, 2018

Note that in Tables 4.5 and 4.6, GEMINIm01L1_{v3} is not included. Unfortunately, the decision to replace GEMINIm01_{v2} with GEMINIm01_{v3} was an extremely last-minute one, requiring 'round-the-clock work from the team in Building 904 for a period of approximately a week. A photo of one of the late night sessions preparing the VFAT3 hybrids can be seen in Figure 4.22.

Despite the best efforts of the team to produce a quality detector, upon installation, GEMINIm01L1_{v3} was found to not be operationally stable, with an unsupportable amount of trips and other issues, and as such was kept off for the remainder of the slice test, save for occasional test attempts by v3 electronics experts.

⁵<http://cmsonline.cern.ch/cms-elog/994364>



Figure 4.22 – Photo of one of the late-night production sessions for GEMINIm01_{v3}. Foreground: Michele Bianco, Jeremie Merlin, Muhammad Ali, and Mohsin Hayat prepare the VFAT3 hybrids. Background: Elizabeth Rose Starling prepares other necessary material for the chamber

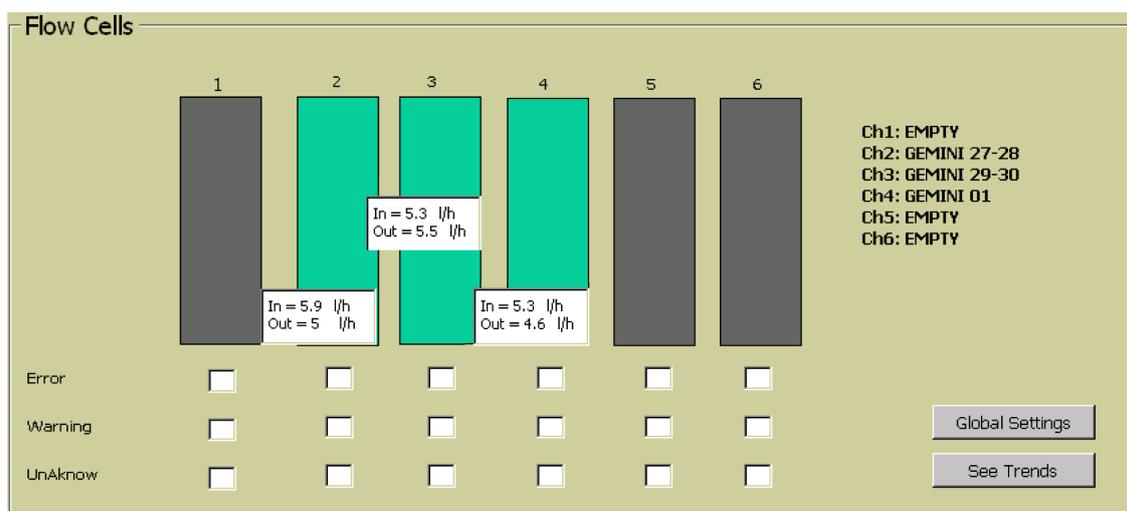


Figure 4.23 – Screenshot of the key portion of the gas panel page of the 2018 slice test DCS, with mouseovers showing the in/out gas flow rate for each set of chambers in liters per hour

In addition to controlling and monitoring the HV and LV for each of the slice test detectors, a DOC would also monitor the gas flow in and out of the detectors, to ensure no leaks had occurred and to ensure proper gas flow was being supplied to the chambers to allow their optimal operation. The main portion of the gas flow panel of the DCS can be seen in Figure 4.23. Here, the DOC could mouse over each flow cell to see the gas flowing into and out of the chambers. It is good to note that these measurements are not made at the gas inlets and outlets of the chambers themselves, but at the level of the closest gas rack of CMS. As a result, the numbers may be very slightly different from reality, and thus a discrepancy the size of the ones shown in the figure are not necessarily a cause for alarm. It is also important to note that each chamber, or even superchamber, is not supplied with its own gas. Rather, the gas cells are grouped into two superchambers each in cells 2 and 3, and one superchamber in cell 4.

Sh	Device DP element	Description	Alarm text	Dir	Value	Ack	Time	Co
W	cms gem dcs 1:CAEN/GEM CAEN L		Channel Error	WENT	FALSE	!!!	2017.03.10 18:48:47.062	
W	cms gem dcs 1:CAEN/GEM CAEN L		Channel Error	WENT	FALSE	!!!	2017.03.10 18:48:47.068	
W	cms gem dcs 1:CAEN/GEM CAEN L		Channel Error	WENT	FALSE	!!!	2017.03.10 18:48:47.073	
W	cms gem dcs 1:CAEN/GEM CAEN L		Channel Error	WENT	FALSE	!!!	2017.03.10 18:48:47.079	
W	cms gem dcs 1:CAEN/GEM CAEN L		Channel Error	WENT	FALSE	!!!	2017.03.10 18:48:47.085	
W	cms gem dcs 1:CAEN/GEM CAEN L		Channel Error	WENT	FALSE	!!!	2017.03.10 18:48:47.091	
W	cms gem dcs 1:CAEN/GEM CAEN L		Channel Error	WENT	FALSE	!!!	2017.03.10 18:48:47.097	
E	cms gem dcs 1:CAEN/GEM CAEN L		Temp High	CAME	36	!!!	2017.03.10 18:54:06.189	
W	cms cen dcs 3:CMSfwClass/CMSfwS	[CMSfwClass/CMSfwScheduler/RDBCh	error	CAME	TRUE		2017.03.10 21:24:51.847	
W	cms cen dcs 3:CMSAlertSystem/Sum	CMSfwScheduler Notification for RDBCh	TRUE	CAME	TRUE		2017.03.10 21:24:51.847	
W	cms gem dcs 1:CMSGEM Mx PT100		Warning: low pressure	CAME	0.0039072041	xxx	2017.03.11 00:51:06.306	
E	cms gem dcs 1:CMSGEM Mx PT100		Alert: too low pressure	CAME	0.0039072041	xxx	2017.03.11 00:51:06.306	
W	cms gem dcs 1:CMSGEM Di PT6124		Warning: high pressure	CAME	12.844933509	xxx	2017.03.11 00:51:08.475	
E	cms gem dcs 1:CMSGEM Di PT6124		Alert: too high pressure	CAME	12.844933509	xxx	2017.03.11 00:51:08.475	
W	cms gem dcs 1:CMSGEM Di PT6122		Warning: high pressure	CAME	13.196581840	xxx	2017.03.11 00:51:08.507	
E	cms gem dcs 1:CMSGEM Di PT6122		Alert: too high pressure	CAME	13.196581840	xxx	2017.03.11 00:51:08.507	
W	cms gem dcs 1:CMSGEM Di PT6128		Warning: low pressure	CAME	1.0622711181	xxx	2017.03.11 00:51:08.536	
W	cms gem dcs 1:CMSGEM Di FE6102		Warning: low input flow	CAME	0	xxx	2017.03.11 00:51:11.313	
E	cms gem dcs 1:CMSGEM Di FE6102		Alert: too low input flow	CAME	0	xxx	2017.03.11 00:51:11.313	
W	cms gem dcs 1:CMS GEM Di Ch4.V		Warning: check the inpu	CAME	1.6999999092	xxx	2017.03.11 00:51:11.402	
E	cms gem dcs 1:CMS GEM Di Ch4.V		Alert: check the input a	CAME	1.6999999092	xxx	2017.03.11 00:51:11.402	
W	cms gem dcs 1:CMSGEM Di FE6102		Warning: low input flow	CAME	0.5	xxx	2017.03.11 00:51:11.431	
E	cms gem dcs 1:CMSGEM Di FE6102		Alert: too low input flow	CAME	0.5	xxx	2017.03.11 00:51:11.431	
W	cms gem dcs 1:CMSGEM Di FE6102		Warning: low input flow	CAME	0	xxx	2017.03.11 00:51:11.460	
E	cms gem dcs 1:CMSGEM Di FE6102		Alert: too low input flow	CAME	0	xxx	2017.03.11 00:51:11.460	
W	cms cen dcs 3:CMSfwClass/CMSfwS	[CMSfwClass/CMSfwScheduler/copyPr	error	CAME	TRUE		2017.03.12 00:41:17.221	
W	cms cen dcs 3:CMSAlertSystem/Sum	CMSfwScheduler Notification for copyPr	TRUE	CAME	TRUE		2017.03.12 00:41:17.221	
W	cms gem dcs 1:CMSGEM Di PT6104		Warning: high pressure	CAME	3.1477413177	!!!	2017.03.12 01:03:09.280	
E	cms gem dcs 1:CMSGEM Di PT6104		Alert: too high pressure	CAME	4.1294269561	!!!	2017.03.12 03:22:37.040	
W	cms cen dcs 3:CMSfwClass/CMSfwS	[CMSfwClass/CMSfwScheduler/checkN	error	CAME	TRUE		2017.03.12 08:20:12.473	
W	cms cen dcs 3:CMSAlertSystem/Sum	CMSfwScheduler Notification for checkN	TRUE	CAME	TRUE		2017.03.12 08:20:12.473	
W	cms gem dcs 1:CMSGEM Mx TotalF		Warning: high total inpu	CAME	15.734155654	!!!	2017.03.12 11:16:03.915	
W	cms cen dcs 3:CMSfwClass/CMSfwS	[CMSfwClass/CMSfwScheduler/checkP	error	CAME	TRUE		2017.03.12 11:32:07.756	
W	cms cen dcs 3:CMSAlertSystem/Sum	CMSfwScheduler Notification for checkP	TRUE	CAME	TRUE		2017.03.12 11:32:07.756	
E	cms gem dcs 1:CMSGEM Di PT6128		Alert: too low pressure	WENT	1.0573873519	!!!	2017.03.12 11:44:41.390	

Figure 4.24 – Screenshot of the alarm screen of the 2018 slice test DCS [78]

The DCS also contains an alarm panel, shown in Figure 4.24. Here, the DOC could monitor any misbehaviors of the system, including LV and HV power trips, incorrect gas flow (either too low/high flow from the gas system, or too low/high pressure within the chambers), too high temperature of the front-end readout electronics, and other system-critical problems. In Figure 4.24, alarms which have not yet been acknowledged by the DOC are seen with three red exclamation marks in the “Ack” column, while alarms which have been acknowledged are marked with three xxx’s. These errors will be seen by the central control technical shifter, and must be acknowledged, and the problem fixed, before

the DOC can continue with operations. Errors at the CMS rather than the GEM level which do not need to be handled by the GEM operator are shown with a blank grey box in the “Ack” column.

Assuming stable gas flow and stable HV and LV, and assuming no active alarms, the DOC would proceed with daily operations as follows:

After taking note of the HV, LV, and gas flow values in an elog, the DOC would proceed to check and, if necessary, reset the communication. This includes:

- Checking the CTP7 uptime. In the event of a crash, reloading the firmware to recover the CTP7. In the event this recovery does not work, contacting the GEM DAQ expert on-call to help solve the problem.
- Checking the TTC state, and following the same CTP7 recovery procedure described above if the state is bad. Due to a fault in the electronics of the slice test detectors, the TTC would go into a bad state following every change of the LHC clock. This problem has since been fixed for GE1/1.
- Checking the AMC status, and ensuring the SCA communication is good. If not, performing an SCA reset. This reset would notably need to be performed each time the LV of a detector had to be cycled.

Once communication was established, the DOC would configure the chambers, using the following command.

```
sudo -u gempro -i confAllChambers.py -s 3 --ztrim=<ztrim value> --vt1bump=<vt1bump value>
--config 2>&1 | tee -a conf_log.txt
```

Here, they would input the desired value for the ztrim and vt1bump, as explained in Section 3.4.1. In general, ztrim was kept at the default value of 4, while vt1bump was alternated between 0 and 10. For each value of vt1bump, the DOC would then perform a set of S-Curves for all chambers of interest - that is, all chambers which were on, well-behaved, and not being excluded for any reason. If a chamber was to be excluded, the DOC would use the following command to comment out that chamber with a #, as shown for example on GEMINIm01 and GEMINIm30L2:

```
sudo -u gempro -i <editor of choice> /nfshome0/gempro/src/gem-plotting-tools/mapping/
chamberInfo.py
```

```
chamber_config = {
    # 0: "GEMINIm01L1",
    # 1: "GEMINIm01L2",
    2: "GEMINIm27L1",
    3: "GEMINIm27L2",
    4: "GEMINIm28L1",
    5: "GEMINIm28L2",
    6: "GEMINIm29L1",
    7: "GEMINIm29L2",
    8: "GEMINIm30L1",
    # 9: "GEMINIm30L2"
}
```

Once the correct chamber configuration was saved, the S-Curves would be taken with the following command:

```
sudo -u gempro -i run_scans.py --tool=ultraScurve.py -s3 2>&1 | tee -a
scurve_log.txt
```

These scans could only be performed when there was not active beam in the LHC, as the presence of beam would skew the results, and were to be performed at least once per day at each vt1bump value.

Once the scans were completed, the DOC would then analyze them in order to produce human-readable plots, using the following command, where the date was given in the format yyyy.MM.dd.HH.mm, and can be found in /gemdata/GEMINIm<#>L<#>/scurve/:

```
sudo -u gempro -i ana_scans.py --scandate=<date> --ztrim=<val> --anaType=scurve 2>&1 | tee
-a ana_scurve_log.txt
```

These plots would then be uploaded to an elog, so as to be available to the larger CMS GEM operations community. These daily S-Curve scans provided a crucial look into the health and behavior of the slice test detectors, as we will discuss in much greater detail in Chapter 6.

4.4.2 Slice Test Results

Ultimately, the slice test succeeded in its goal as a demonstrator system. While the majority of 2017 was spent commissioning the detectors and working with them in local operating mode, on November 9th, 2017, the slice test detectors were successfully integrated into CMS global running. This milestone was marked by Figure 4.25, a screenshot of the CMS information screen from that time.

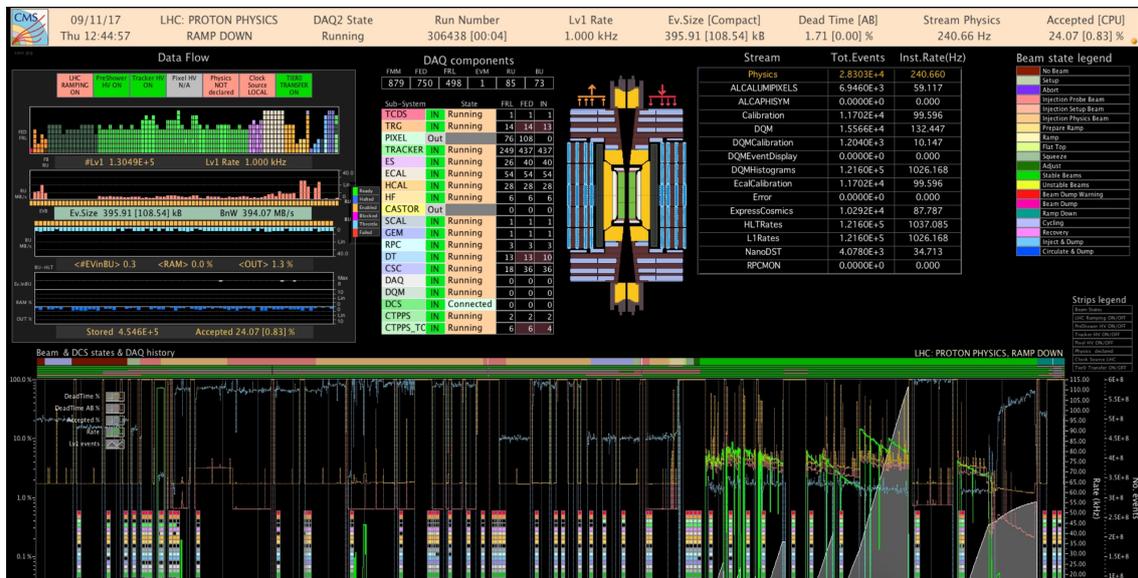


Figure 4.25 – A screenshot of the CMS information screen from November 9th, 2017, showing the inclusion of the slice test GEMs in global CMS operations for the first time since their installation. Data taking mode was cosmic muons during an LHC interfill period

While Figure 4.25 shows the inclusion of the slice test GEMs during cosmic running, Figure 4.26 represents a further achievement: the inclusion of the slice test GEMs in CMS global running during p-p collisions.

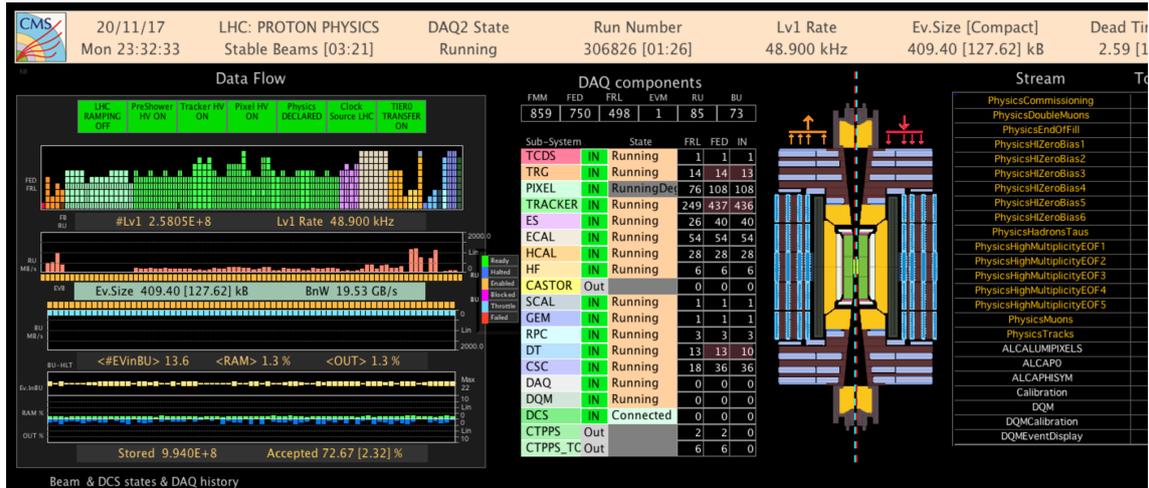


Figure 4.26 – A screenshot of the CMS information screen from November 20th, 2017, showing the inclusion of the slice test GEMs in global CMS operations during p-p collisions for the first time since their installation

And, indeed, we can see that the slice test GEMs were correctly configured in the CMS data taking system from this point. Figure 4.27 shows the full chain of data processing within the CMS Online Data Quality Management (DQM) system (see: References [79] and [80]) for all eight single-channel power supply slice test detectors. For each detector, we see the raw data, the unpacking of that data, the digitization of that data into “digis”, and finally the reconstruction of those digis to their physical position on the chamber. In CMS parlance, the digitized data are called “digis” and the reconstructed hits are called RecHits. In particular, the top left histogram for each chamber shows the digi strips fired, the top right histogram shows the VFATs vs. cluster sizes, the bottom left histogram shows the RecHits collected in each VFAT, and the bottom right histogram shows the RecHit distribution in the local position.

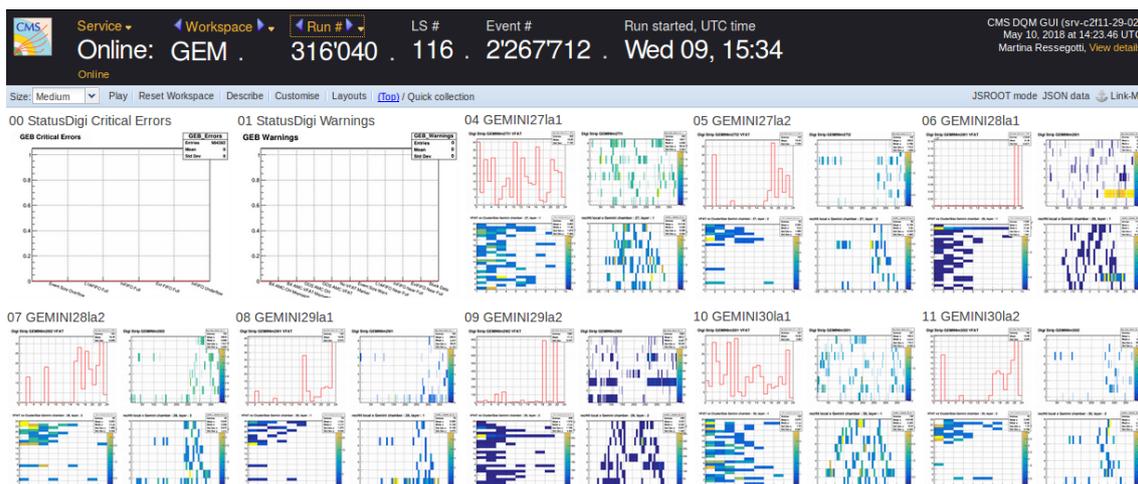


Figure 4.27 – A screenshot of the CMS online DQM system from cosmic run #316040 on May 9th, 2018 showing the various stages of data processing for each of the eight single-channel power supply slice test detectors

With the GEMs operating correctly in CMS global running, we are able to finally see collision events as they are reconstructed using the CMS GEMs. One such event can be

seen in Figure 4.28, as viewed using the CMS iSpy event display (see: Reference [81] for more details). Here, we see a collision event within the CMS detector from July 8th, 2018. The single-channel power supply slice test detectors are highlighted in their respective positions in lavender. Two muons, represented by the red lines, from a reconstructed J/Ψ candidate, passed through the chambers, with their associated hits in the GEM chambers shown in purple. This represents a monumental achievement for the CMS GEM project as a whole.

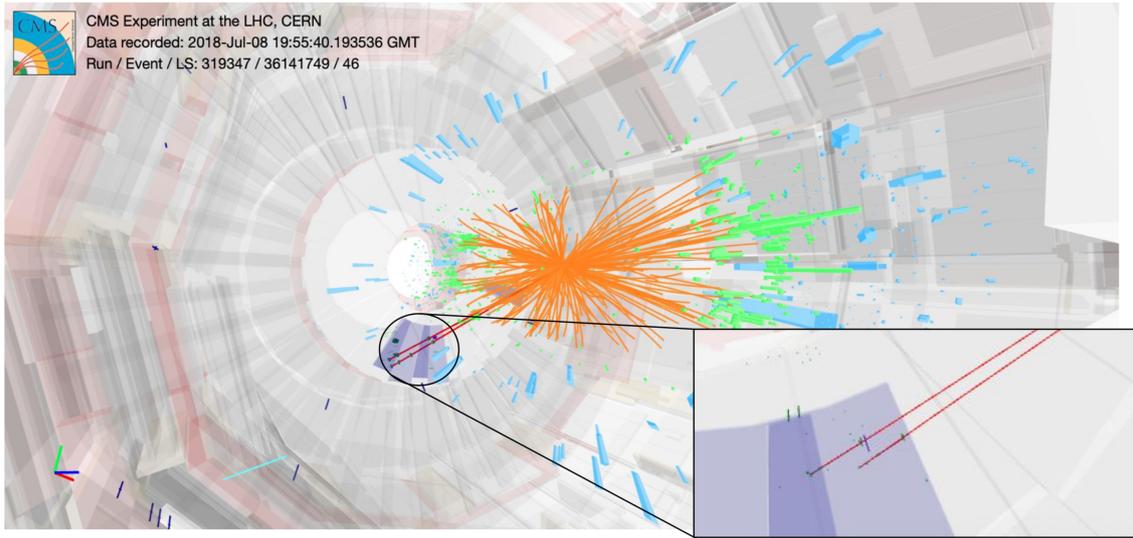
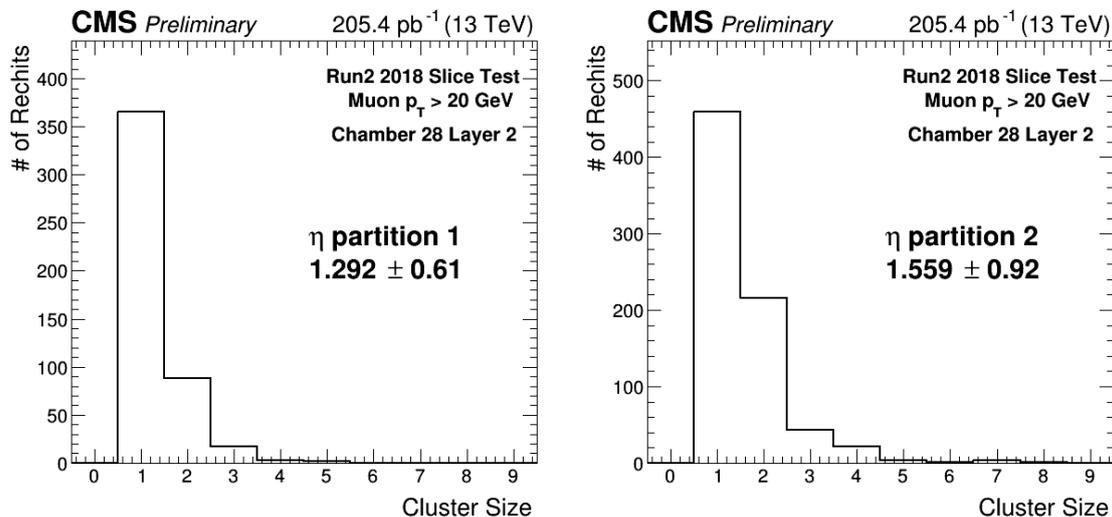


Figure 4.28 – An LHC proton-proton collision event display, with a close-up view on the slice test GEMs (lavender trapezoids) showing two muons (the red lines) and their associated hits in the slice test GEMs (purple) and the CSCs (green). The antimuon ($p_T = 30.11$ GeV, $\eta = -1.956$) and muon ($p_T = 53.60$ GeV, $\eta = -1.993$) have a combined invariant mass of 3.01 GeV. [82]

Taking GEMINIm28L2 as an example, as it will be shown in Chapter 6 to have the best overall sustained health, we can use the data taken during these collisions to calculate the efficiency of the slice test detectors. Figures 4.29a through 4.29h show, for 2018 collision data, the cluster size of the RecHits in each η partition of the GEMINIm28L2 detector.



(a) Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 $\eta = 1$ (b) Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 $\eta = 2$

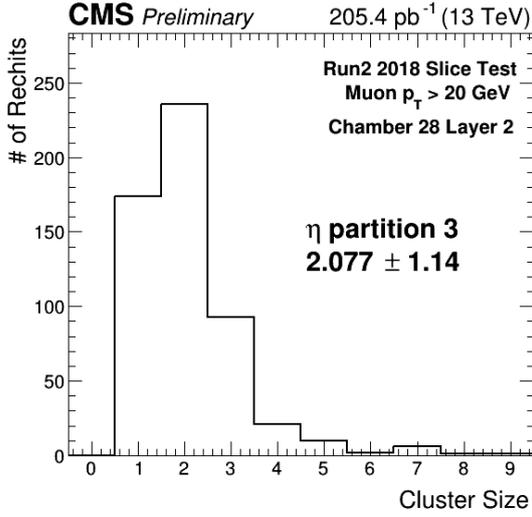
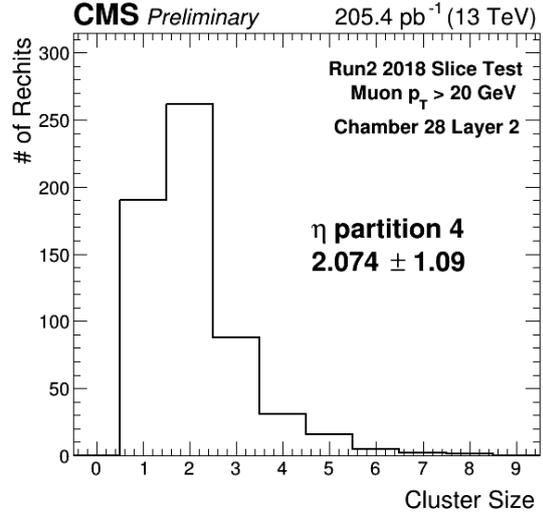
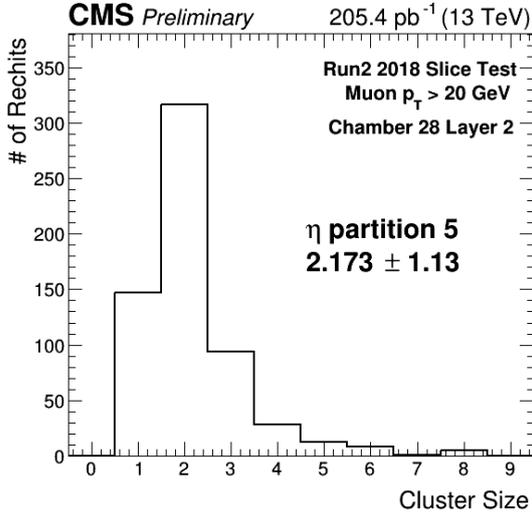
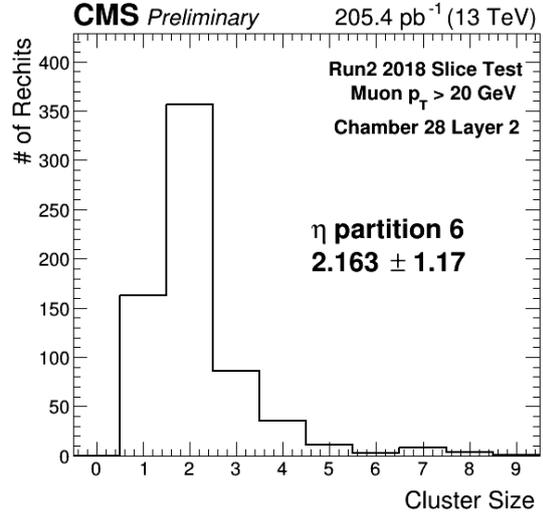
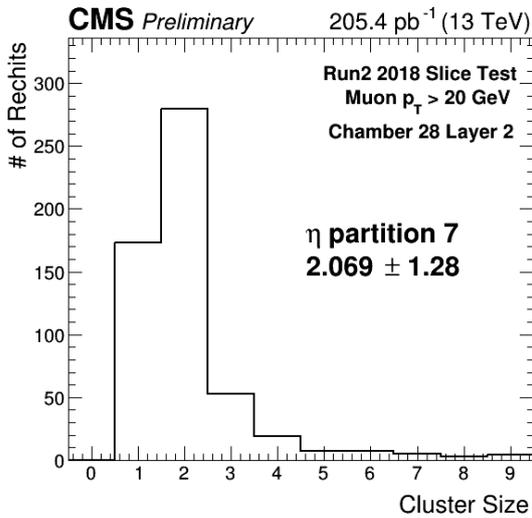
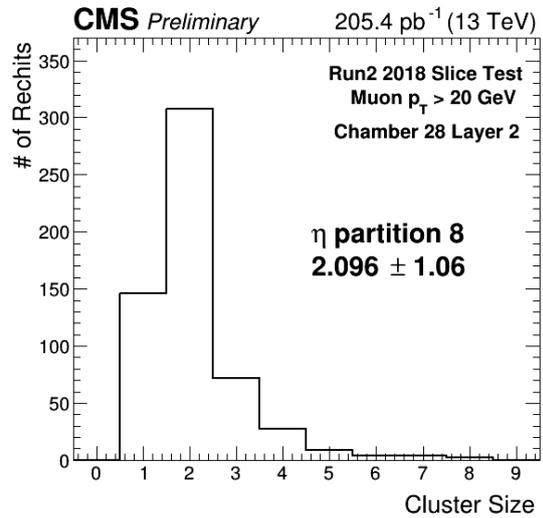
(c) Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 $\eta = 3$ (d) Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 $\eta = 4$ (e) Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 $\eta = 5$ (f) Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 $\eta = 6$ (g) Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 $\eta = 7$ (h) Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 $\eta = 8$

Figure 4.29 – Cluster size of RecHits to global muons with $p_T > 20$ within GEMINIm28L2 per η sector from 205.4 pb^{-1} of 2018 p-p collision data at a center-of-mass energy of 13 TeV [82]

In particular, the muons used in this study were well-identified muons which were reconstructed only from well-established non-GEM detectors (i.e. CSCs), but from which at least one GEM RecHit could be reconstructed. These muons were required to conform to strict identification criteria, with $p_T > 20$. Once identified, these muons were then propagated from their final hit within a non-GEM detector to the region of the slice test detectors. In order to exclude inefficiencies due to propagation uncertainty, only those muons which were successfully propagated to a point within the GEM detector not within 20 channels of the detector's edge were considered. If the GEM RecHit was found within 5 cm of the propagation point, it was considered as being due to that muon [83].

The RecHits, on the y-axis, represent the number of times there was a particle hit within the detector, and the cluster size, on the x-axis, represents the number of readout strips which saw a signal charge as a result of that particle interaction. An ideal case would show the vast majority of the RecHits as having a small cluster size, and a higher amount of larger cluster sizes may be indicative of noise within the chamber.

For $\eta = 1$, at the wide end of the detector, the average cluster size is closest to 1, at 1.292 ± 0.61 . And as you go up the detector towards the narrow end, increasing in η sector number, this average cluster size increases as well, with the majority of the η sectors having an average cluster size closer to two strips than one. In this case, the increase in cluster size is not in fact due to noise, as we have seen in Section 3.4.1 that the noise is actually greatest at lower-valued η partitions closer to the optohybrid, but rather may be due to the size of the strips themselves. As the width of the detector narrows, so too does the width of the readout strips present there. Thus, a Townsend avalanche which might only be read out on one readout strip in η sector 8 might cover two of the narrower strips of η sector 1. So, unfortunately, we cannot necessarily compare these plots against each other in order to get a sense of the performance of the detector. Instead, we can look at the efficiency of the chamber by η partition, as shown in Figure 4.30, and consider some additional factors.

Here, the efficiency is derived from the extrapolation of global muons with $p_T >$ to the GEM chamber in order to associate hits within GEMINIm28L2, once again using 205.4 pb^{-1} of 2018 p-p collision data at a center-of-mass energy of 13 TeV. These extrapolated muons were required to be 20 strips away from the edges of the chambers.

In this plot we list both the efficiency of the chamber per η sector (black marker) and the fraction of inactive strips (white marker) per η sector, without dividing the latter into the reasons why the strips may be inactive (See Section 6.1 for more detail on channel mask reasons).

Here, we see, as we might have expected, that the lowest efficiency (73%) is seen in the $\eta = 1$ sector and the highest efficiency (98%) is seen in the $\eta = 8$ sector. However, while the noise in the detector undoubtedly plays some role, here the inefficiency is determined more by the fraction of inactive channels in the detector. While GEMINIm28L1 was the best-performing chamber from a channel loss perspective, it still had a non-negligible amount of inactive channels. In particular, η partitions 1 and 2, which had the lowest efficiencies, had channel inactivity at 2.1% and 1.3% of their total channels, respectively. Of these inactive channels, 9.4% and 19.0% were due to the higher thresholds applied to the VFATs in those chambers, which results in the loss of some signal charge. The average VFAT thresholds for each η sector can be seen in Table 4.7.

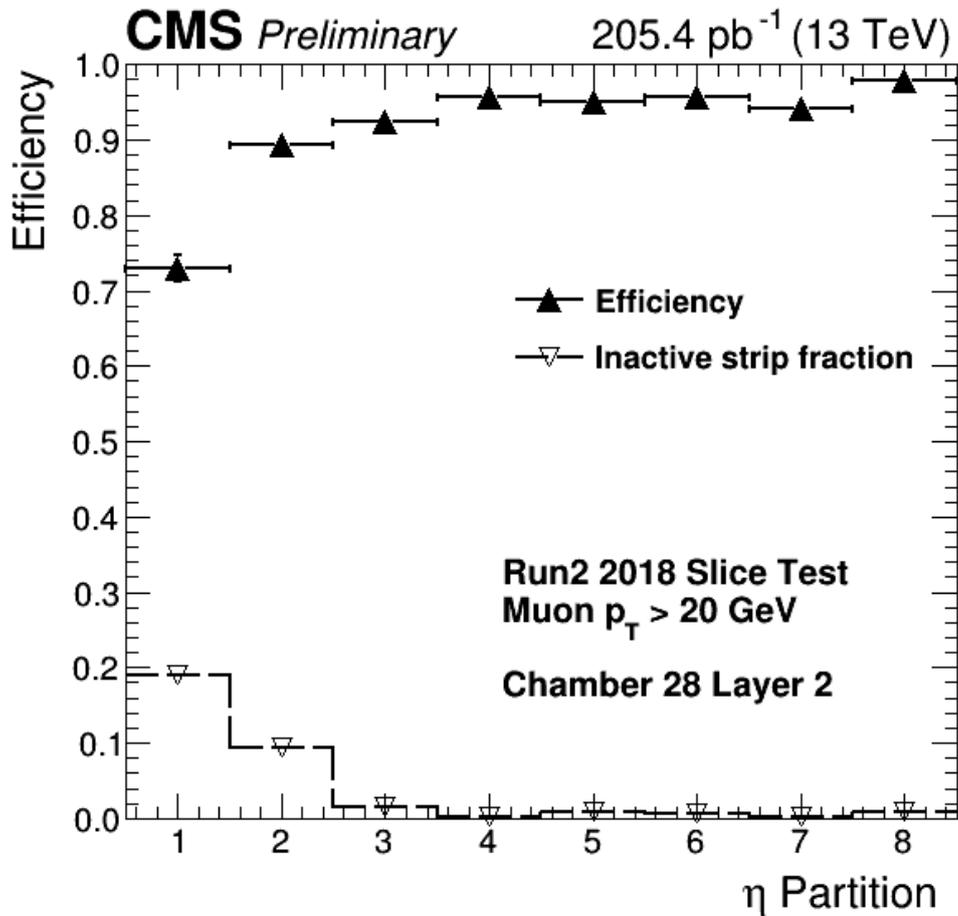


Figure 4.30 – Efficiency of GEMINIm28L2 as a function of η partition, as derived from the extrapolation of global muons with $p_T >$ to the GEM chamber in order to associate hits within GEMINIm28L2, using 205.4 pb^{-1} of 2018 p-p collision data at a center-of-mass energy of 13 TeV [82]

η	Average VFAT2 threshold (fC)
1	8.8 ± 2.6
2	4.0 ± 0.9
3	3.5 ± 0.3
4	3.2 ± 0.6
5	3.0 ± 0.6
6	2.5 ± 0.3
7	3.1 ± 0.5
8	2.8 ± 0.3

Table 4.7 – Average VFAT2 threshold in fC per η sector

Knowing this, when we turn back to the cluster size plots of Figure 4.29, we can infer, in addition to the proposal of the strip width playing some role, that there may be some bias from this loss of signal charge as a result of the high thresholds, in which the true signal multiplicity is lost in some cases for the lower-valued η sectors.

In addition, some of the problem may also stem from the response nonuniformity of the detector. Figure 4.31 shows the normalized response uniformity over the detector, in mm^2

bins. Here, we see a significant non-uniformity in the detector, with the higher η sectors showing a lower response uniformity and portions of the lower η sectors showing a higher response uniformity. Compensating for this required voltage to be “pushed” to the rest of the detector, which may have also contributed to the inefficiencies.

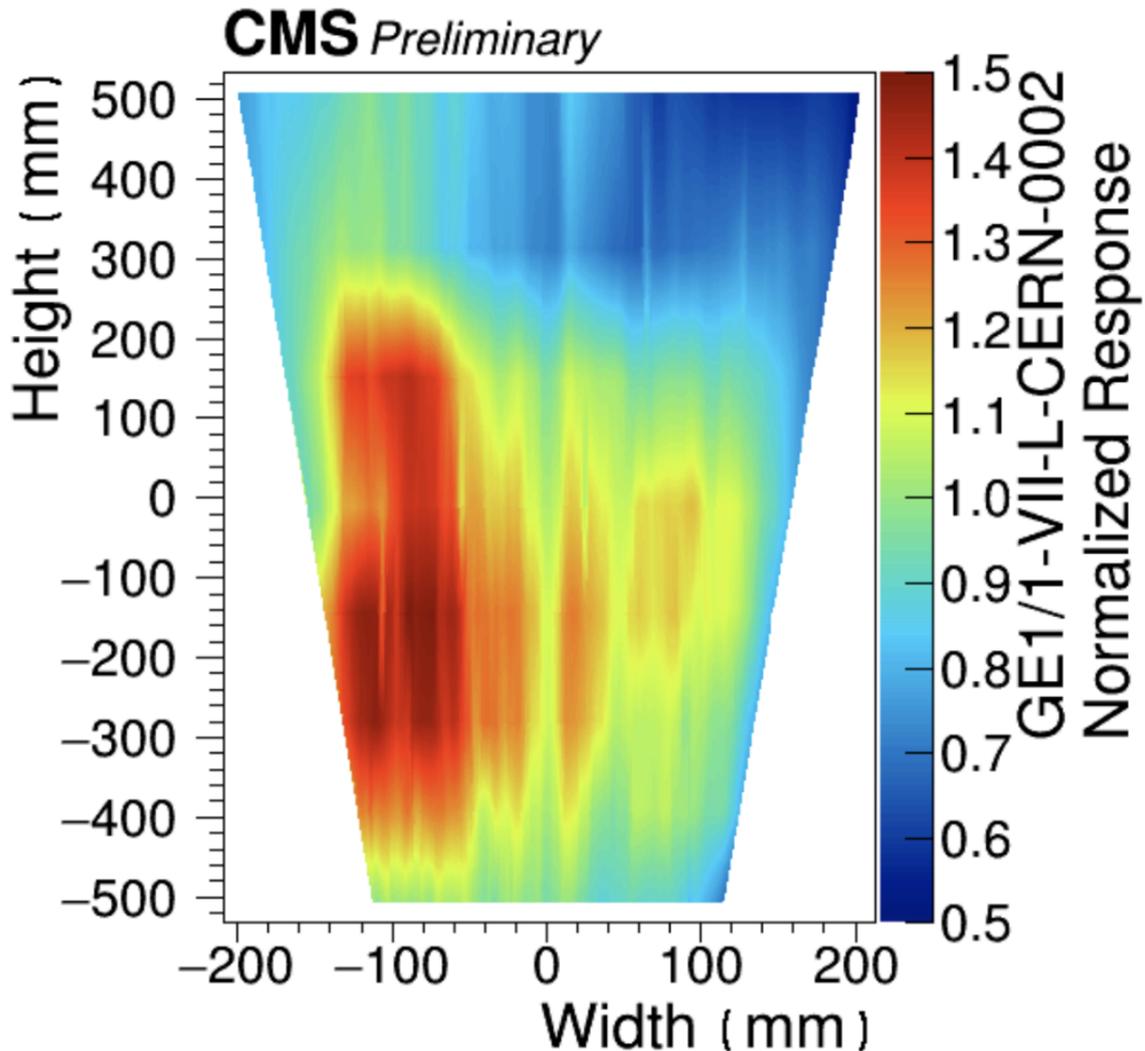


Figure 4.31 – Normalized response uniformity over the physical area of the detector

Nevertheless, despite these two problematic sectors, the overall efficiency of the chamber remains above 90%, which is acceptable given the very experimental nature of the slice test detectors. And, indeed, we can say that the slice test as a whole was a success, as we were able to successfully integrate into CMS global running in a stable manner for both cosmic and proton-proton collision runs, and in these global runs we were able to take real muon data, reconstructing tracks with both the slice test GEMs and ME1/1 CSCs.

However, despite the solid successes of the slice test, one goal in particular was not met: the ability to tolerate an integrated charge of $\sim 18 \frac{mC}{cm^2}$ without suffering damage of performance loss, and the presence of a discharge rate which does not impede performance or operations. But as we will see in the following chapters, the manner in which the slice test detectors failed to meet this goal actually proved to be one of the most important learning experiences gained during the slice test, as it would impact not just the GE1/1 system, but the GE2/1 and ME0 systems as well.

Chapter 5

Discharge Damage within Gaseous Detectors

Before discussing the health and longevity studies that were performed with the slice test detectors in Chapter 6, it is crucial to first establish the main cause of damage to gaseous detectors: electrical discharges within the active gas volume.

5.1 Theory of Discharges within Gaseous Detectors

When discussing discharges within gaseous detectors, we are ultimately discussing the electrical breakdown of the gas within them, in which the dielectric strength of the gas is exceeded and electrical conduction occurs.

One characteristic of the breakdown within a gaseous detector is that the voltage across the gas gap drops in a process which results in a high conductivity between the cathode and anode. This takes place over a relatively short time period, such that the gap capacity cannot be reloaded quickly, resulting in the breakdown of the voltage across the gap [29]. However, the process by which this high conductivity is built up is not identical for all discharges, resulting in the classification of discharges by the mechanism by which this occurs: either “slow” breakdowns or “fast” breakdowns.

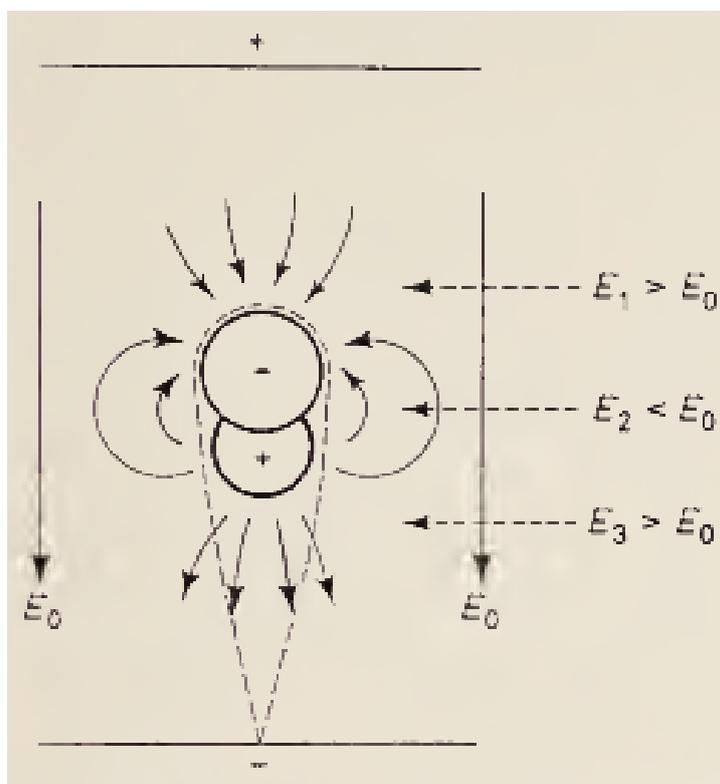


Figure 5.1 – The effect of the space charges of a high-amplification avalanche on the electric field E_0 [29]

Referring back to Section 2.2, we can consider the body of the Townsend avalanche within a gaseous detector as a space charge, in which excess electric charge is treated as a continuum of charge distributed over a given volume rather than as distinct point charges. The electric field around such an avalanche can be seen in Figure 5.1, where the field behind and in front of the head of the avalanche is increased by the space charge and reduced between the electron and ion cloud.

In both cases, the space charge produces the decisive current growth within the gas gap that leads to the characteristic voltage drop, though the mechanism proceeds differently.

5.1.1 Slow Breakdowns

The first mechanism is referred to as a “slow breakdown” as it is a relatively slow one, developing over a period of $>10 \mu\text{s}$ over several, up to several dozens, secondary processes. Here, the space charge of the slow positive ions, which are produced by the rapidly-succeeding avalanches, generally favors the ionization conditions for the electrons and produces a rapid current growth. In all cases of slow breakdown, the area in which these secondary avalanches appear rapidly expands, but whether the breakdown ends in a spark or manifests as a type of glow discharge depends heavily on gas composition, pressure, and detector geometry [84].

Slow breakdowns are usually seen within a single-wire counter with a thin anode wire or in detectors which are combined with highly-efficient photocathodes. This breakdown type has been studied intensely over the last half of a century, and is of lesser interest for us in this Section.

5.1.2 Fast Breakdowns

The second method is referred to as a “fast breakdown” as it is far more rapid than the first method, developing within a 1 - 10 ns timescale, in which the first avalanche develops directly to breakdown. Here, the space charge approaches 10^8 electrons, a value known as the Raether limit [85], which is the physical limiting value of the multiplication factor (M) in an ionization avalanche process (see Section 2.2). This limit is described by Equation 5.1, where A is the gas gain of the detector and n_0 is the number of primary electrons created in the drift region of the detector. This value of $M \approx 10^8$ was first proposed by Heinz Raether in “Electron Avalanches and Breakdown in Gases” and then experimentally determined in “The study and optimization of new micropattern gaseous detectors for high-rate applications” by V. Peskov et al. [86]

$$M = An_0 \approx 10^8 \quad (5.1)$$

At this value, the space charge is great enough that it disturbs the external electric field. As a result, photoelectrons created by avalanches in the surrounding gas volume begin to drift toward the positive ions which remain from the initial avalanche. This results in a high-conductivity plasma streamer developing towards the anode from the avalanche, known as an “anode-directed streamer”. This anode-directed streamer was observed by Raether to have a velocity of $8 \cdot 10^5 \frac{\text{m}}{\text{s}}$ in air at $E/p \approx 40$, while the electron drift velocity in the same conditions had a velocity of $1.3 \cdot 10^5 \frac{\text{m}}{\text{s}}$. This discrepancy can be understood to be on account of the higher electric field on the anode side of the avalanche head, as shown in Figure 5.1. In addition, this increased velocity can be the result of the cooperation of a gas-ionizing radiation with the space charge field, where the radiation produces electrons

around the avalanche head. Those electrons which are created in the increased space charge field between the avalanche head and anode can ionize more efficiently. Thus, new radiation is excited and new electrons are produced further ahead, and ion pairs are produced in a certain distance without the electrons proceeding so far. As a result, the electron cloud can therefore reach the anode more quickly than under the influence of the space charge field alone [29]. This concept is illustrated in Figure 5.2.

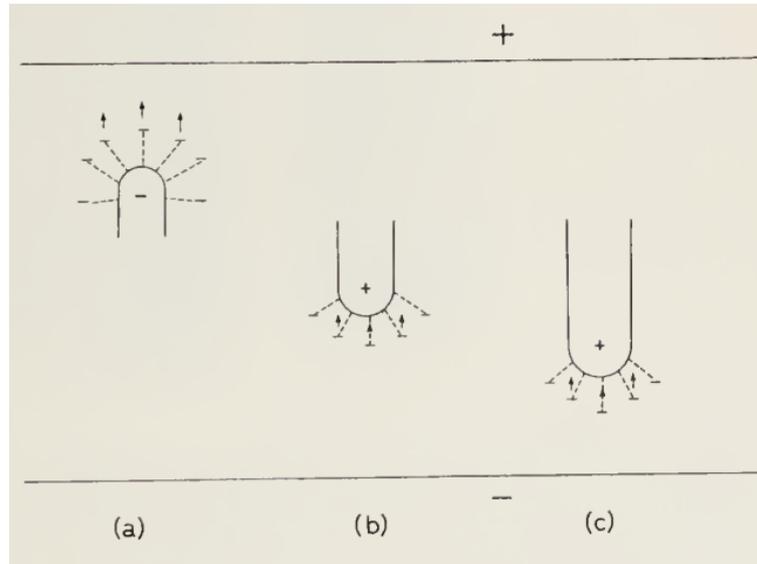


Figure 5.2 – Development of plasma streamers via gas-ionizing radiation [29]

This gas-ionizing radiation can also be used to explain the development of an opposing-direction plasma streamer, known as the “cathode-directed streamer”. As with the anode-directed streamer, the electrons produced via this ionizing radiation fall into the high space charge field very rapidly and excite new radiation, thus propagating the streamer. This radiation is shown as the photons of $h\nu$ directed from the head of the streamer in Figure 5.3, interspersed with the secondary avalanches that result.

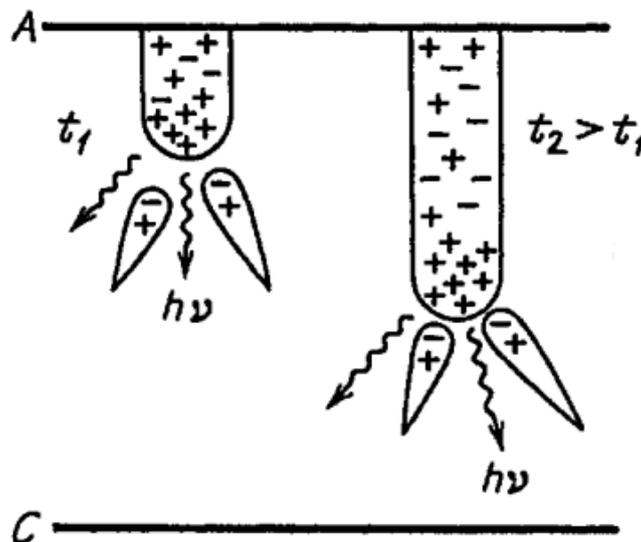
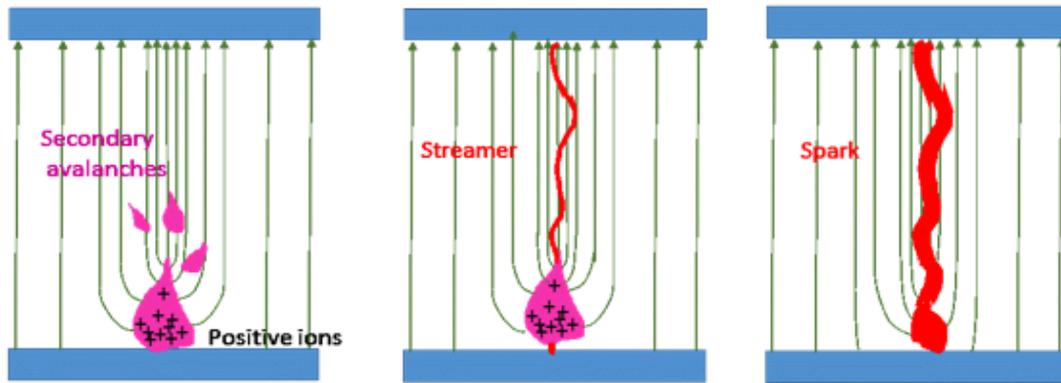


Figure 5.3 – Cathode-directed streamer at two consecutive moments in time, showing ionizing photons which propagate the streamer’s formation [29]

Once both the anode-directed and cathode-directed streamers have made contact with their respective electrodes, thus creating an electrical short between the two, the streamer channel then becomes the spark channel of the electrical discharge. The resulting spark is visible to the naked eye, with a channel radius of up to $r \approx 1$ cm, and can consist of currents greater than 1 A (up to 10^4 A), in comparison to the preceding streamers which allow for a current flow only on the order of 10^{-2} A. The flow of this high discharge current releases a large amount of heat which increases the temperature of the gas, with the spark channel reaching temperatures of up to 20,000 K. This temperature increase in turn causes a sharp increase in gas pressure, which results in a shock wave, which produces a characteristic “snap” sound [87].

The full process from avalanche to discharge is shown in Figure 5.4. Here, we see the electric field lines closest to the avalanche experiencing a focusing effect, which can lead to the start of secondary avalanches and result in the formation of the streamer. When this streamer reaches both electrodes, a spark occurs in the same channel the streamer occupied.



(a) Focusing of the electric field lines around an avalanche and resulting secondary avalanches (b) Formation of an anode-directed and cathode-directed plasma streamer (c) Conduction of an electrical spark through the streamer channel

Figure 5.4 – Process of a fast breakdown resulting in a spark [88]

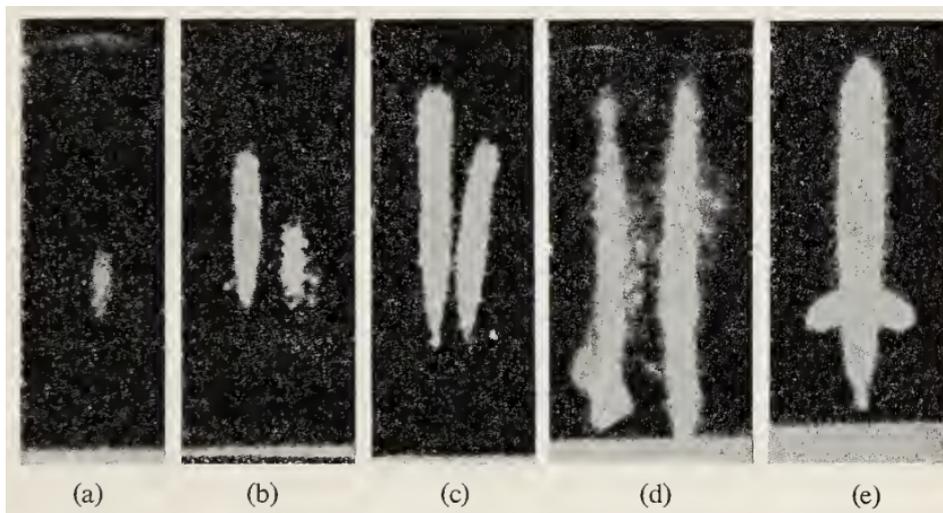


Figure 5.5 – Fast breakdown resulting in a plasma streamer and spark from an electron avalanche in a cloud chamber [89]

This process was imaged in a cloud chamber (air, 270 Torr) by Heinz Raether, seen in Figure 5.5. In (a), the expansion ratio was reduced such that only the head of the avalanche, as the region of the highest ion density, was visible as a track. In (b) and (c), the voltage was slightly raised at a constant pulse duration, such that an anode-directed streamer developed out of the avalanche head. In (d), the voltage was increased further such that a cathode-directed streamer developed out of the avalanche head as well, creating a plasma channel which bridged the two electrodes. And finally, in (e), an electrical spark flows through the channel created by the plasma streamer's shorting of the two electrodes.

Although fast breakdowns are usually seen in parallel plate avalanche counters (PPACs) or MPGDs [88], the study of electrical discharges within MPGDs such as GEMs is a fairly recent one, as the near-instantaneous occurrence of the primary electrical discharge makes measurements difficult.

5.2 Propagating Discharge Mechanism in GEM Detectors

The condition of the avalanche's space charge overcoming the Raether limit is generally only met at the third GEM foil, G3, where the charge density is highest. Under these conditions, the electron avalanche develops into a plasma filament, which causes a short circuit between the top and bottom GEM electrodes, as described in Section 5.1.2. This short circuit is referred to as a "primary discharge".

This primary discharge creates a hot spot on the copper of the GEM foil near the rim of the hole where the discharge occurred, at a temperature of $>3500^\circ\text{C}$. These hot spots cause damage to the GEM foil holes in as little as one discharge, as seen in Figure 5.6b, widening the diameter of the hole with each successive discharge and causing splatter damage around the circumference of the hole.

The temperature of $>3500^\circ\text{C}$ is a key component here, as this is high enough to initiate, along with the presence of the local electric field, the thermionic emission of electrons via the Schottky Effect (also called field-enhanced thermionic emission), which is given by Equations 5.2 and 5.3. Here, J is the emission current density, F is the magnitude of the electric field at the emitter's surface, T is the temperature of the metal, W is the work function of the metal, k is the Boltzmann constant, q_e is the elementary charge, ϵ_0 is the vacuum permittivity, and A_G is the product of a universal constant A_0 multiplied by a material-specific correction factor λ_R which is typically on the order of 0.5.

$$J(F, T, W) = A_G T^2 e^{-\frac{(W-\Delta W)}{kT}} \quad (5.2)$$

$$\Delta W = \sqrt{\frac{q_e^3 F}{4\pi\epsilon_0}} \quad (5.3)$$

These equations come as a modification of the Richardson equation, by replacing W by $(W - \Delta W)$. Here, the thermionic emitter of electrons is negatively biased in relation to its surroundings, which creates an electric field of magnitude F at the surface of the emitter. Without this field, the surface barrier which would be seen by an escaping Fermi-level electron would have a height W which is equal to the local work function, which is defined as the minimum thermodynamic work required to remove an electron from the surface of the metal to a point just outside of the surface, at a distance which is far on an atomic scale but still too close to the metal to be influenced by any ambient external electric fields.

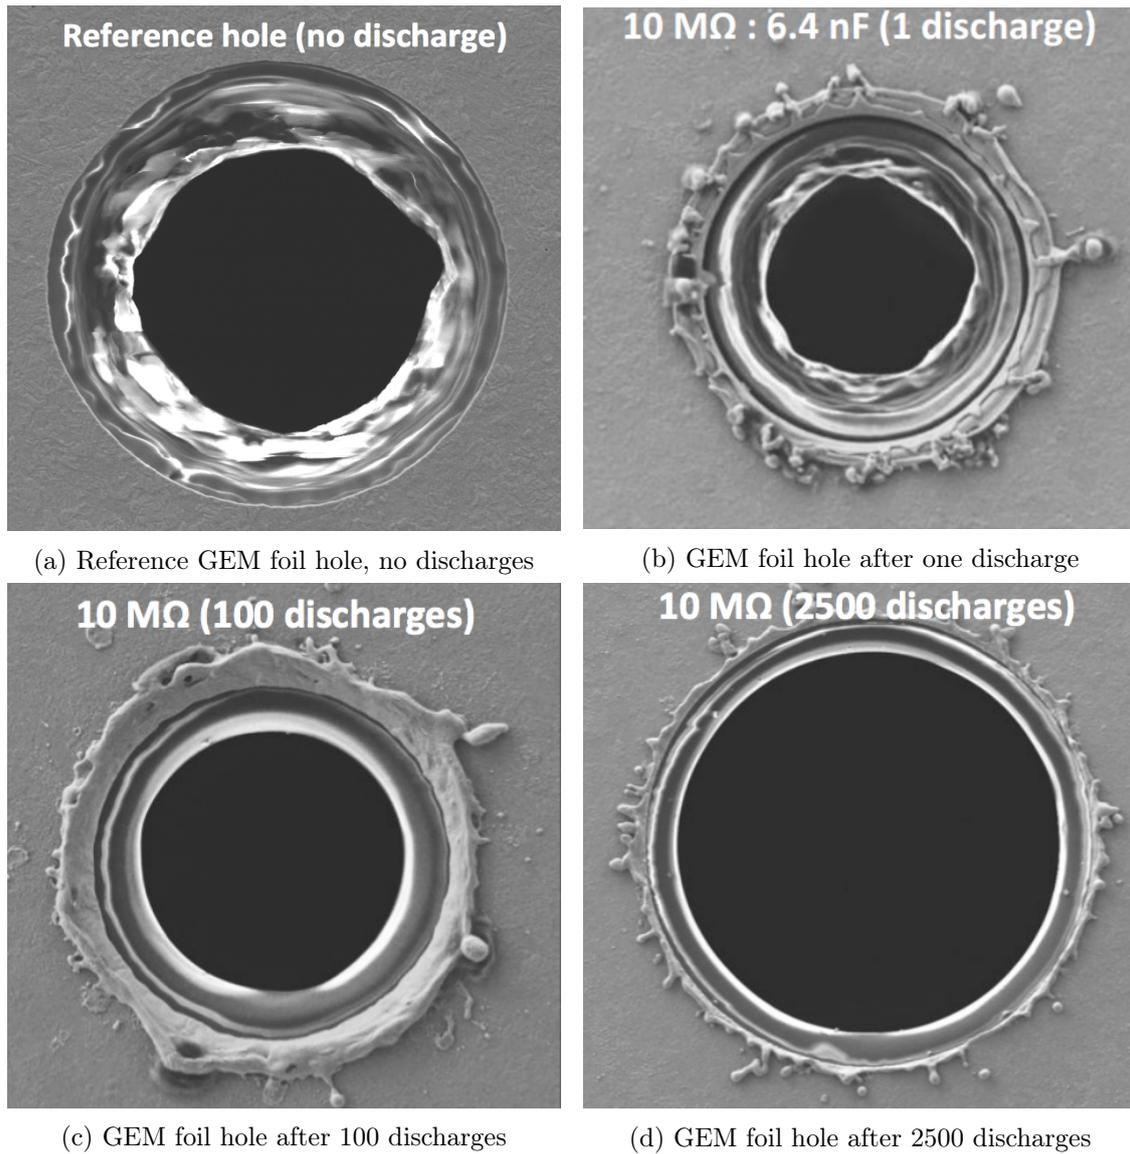


Figure 5.6 – SEM images of the bottom of a GEM foil hole after a number of discharges with an energy of 2.0 mJ [90]



Figure 5.7 – Thermionic emission present in a THGEM with $E_{ind} = 6.5 \frac{kV}{cm}$ and $\Delta V_{THGEM} = 1000 \text{ V}$ [91]

However, the presence of the electric field lowers the surface barrier by an amount ΔW and increases the emission current. Schottky's modifications of the Richardson Equation take this lowering into account to create a more accurate model. This emission was imaged in a thick GEM (THGEM) by Antonija Utrobičić et. al from the University of Zagreb [92], as seen in Figure 5.7.

At this point, it is possible for a precursor current to develop. It is possible to induce a self-quenching of this precursor current, by introducing a drain resistor, also known as a bleeder resistor in some literature, to the detector. This drain resistor is inserted between the readout strips and the ground plane of the detector, labeled as “Termination 100 k Ω ” in Figure 5.8, in order to temporarily reduce the induction field present after the primary discharge.

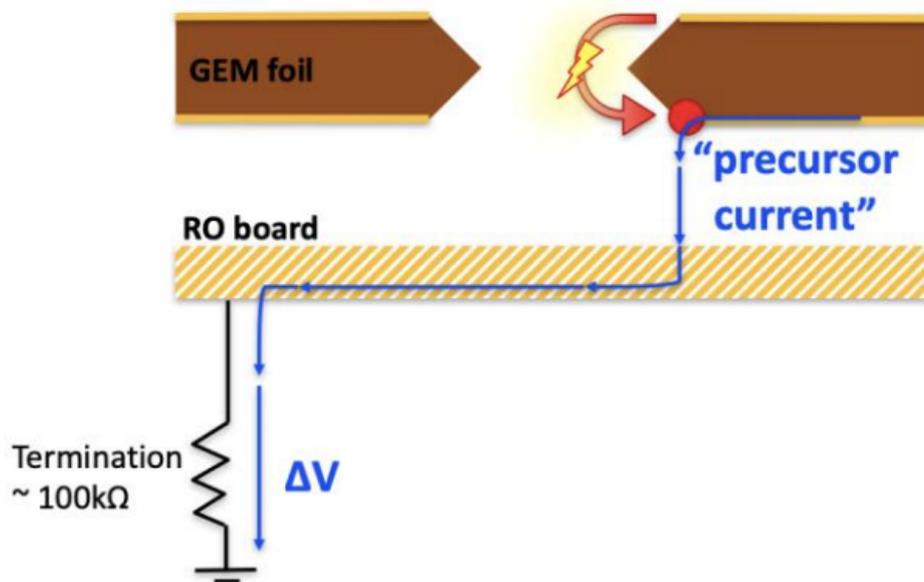
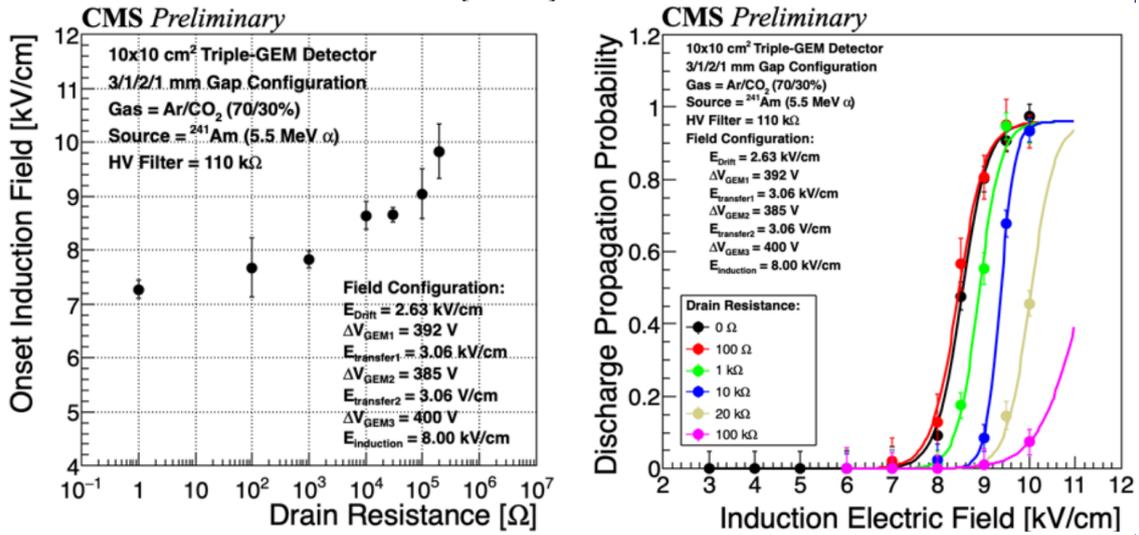


Figure 5.8 – Example of a 100 k Ω drain resistor on a GEM detector readout board [93]

This induction field reduction, shown via the onset induction field as a function of the drain resistance under a constant electric field configuration, is shown in Figure 5.9a for a 10x10 cm² triple-GEM detector with a 3/1/2/1 mm gap configuration in 70% Argon, 30% CO₂. Here, the “onset induction field” is defined as¹ the induction field at which the discharge propagation probability is equal or greater than 1%. As a result of this field reduction, the precursor current cannot grow and develop into a true streamer, regardless of the amount of energy available on the GEM foil.

¹Sophie defined it as “———4” before shutting the author’s laptop off, so her definition is final.



(a) Induction field as a function of drain resistance in a 10x10 cm² triple GEM detector in 70% Ar, 30% CO₂ (b) Probability of discharge propagation as a function of drain resistance and induction electric field

Figure 5.9 – Effects of drain resistance on onset induction electric field and discharge propagation in a 10x10 cm² triple GEM detector [93]

In the absence of any quenching, this precursor current can then further develop into a true streamer, causing a second discharge to take place. This discharge is then referred to as a propagating discharge. The probability of developing into a propagating discharge, in the absence of any quenching as described previously, is dependent in part on the induction capacitance.

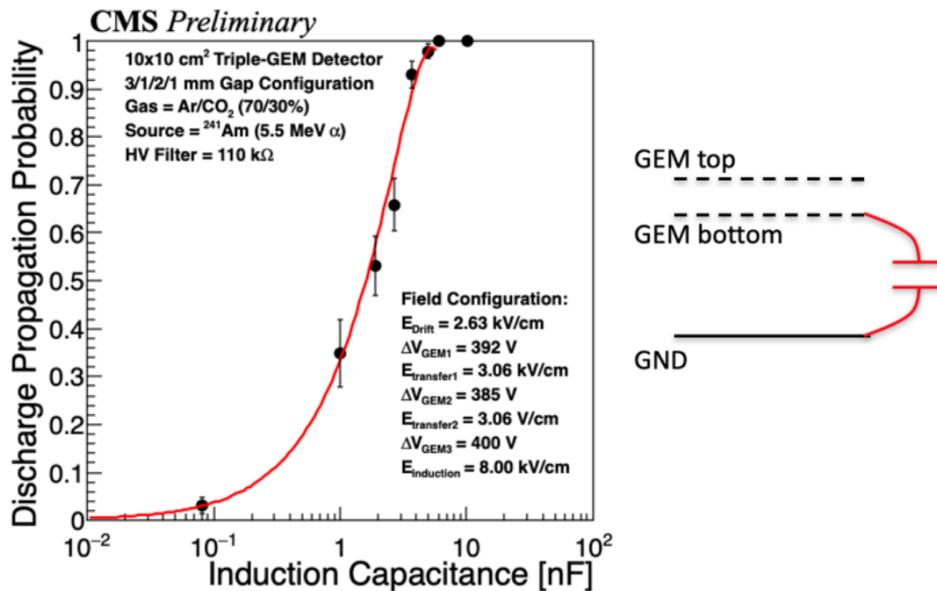


Figure 5.10 – Discharge propagation probability as a function of induction capacitance between G_{3bottom} and ground [93]

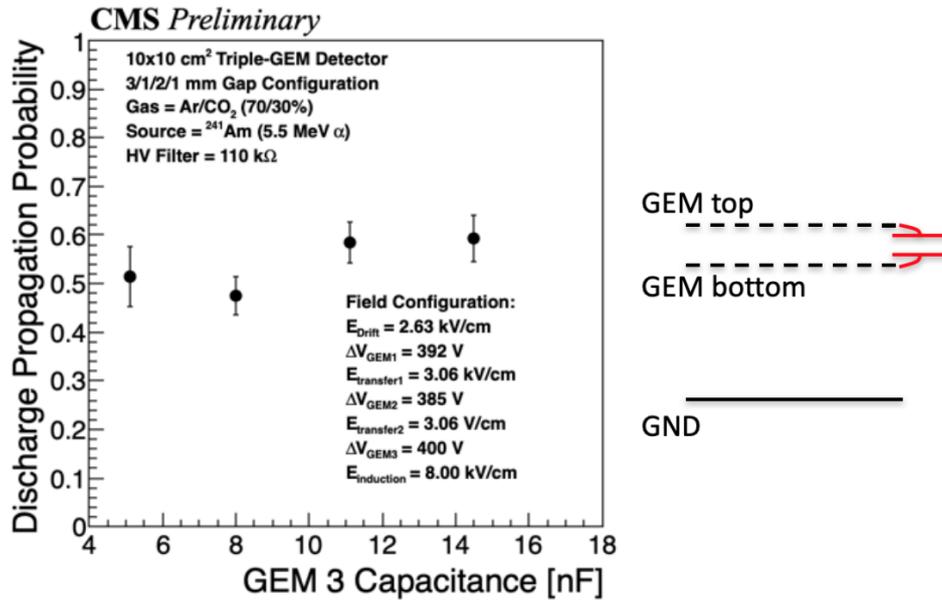


Figure 5.11 – Discharge propagation probability as a function of the capacitance of the G3 foil [93]

We see in Figure 5.10 that for a 10x10 cm² triple-GEM detector with a 3/1/2/1 mm gap configuration in 70% Argon, 30% CO₂, as the induction capacitance between G3_{bottom} and ground increases, so too does the probability of discharge propagation. This is due to the fact that as the induction capacitance increases, the amount of energy on the foil available to feed the precursor current and trigger discharge propagation increases as well. Note that this change of induction capacitance is achieved via the addition of an external capacitor between G3_{bottom} and ground, as shown on the right of Figure 5.10. Interestingly, the capacitance of the foil itself does not appear to have any effect on the primary discharge energy, as shown in Figure 5.11 for the same detector. Here, the data point at 4.8 nF represents the bare 100 cm² foil, and the remaining three data points represent the baseline 4.8 nF plus some additional external capacitor.

Figure 5.9b, then, shows how an increased drain resistance on the GEM detector, by virtue of quenching the precursor current, can therefore stop this development into a true discharge streamer and therefore lower the discharge propagation probability regardless of the amount of energy available on the foil.

The full process from primary discharge to propagation is illustrated in Figure 5.12, starting with the primary discharge, then the hot spot formation, followed by the precursor current and finally the discharge which propagates to the readout board.

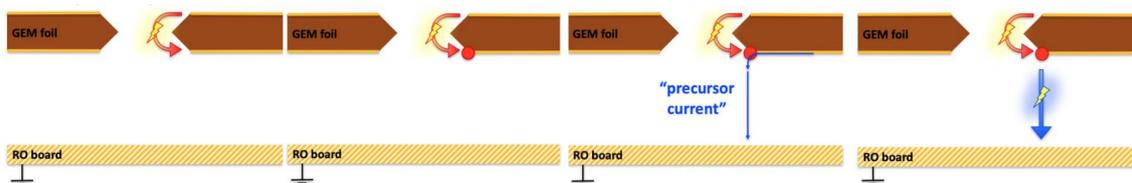


Figure 5.12 – Process by which a discharge propagates in a GEM detector [85]

In small (ex: 10x10 cm²) GEM detectors, this process is simple and localized, as seen in Figure 5.13a.

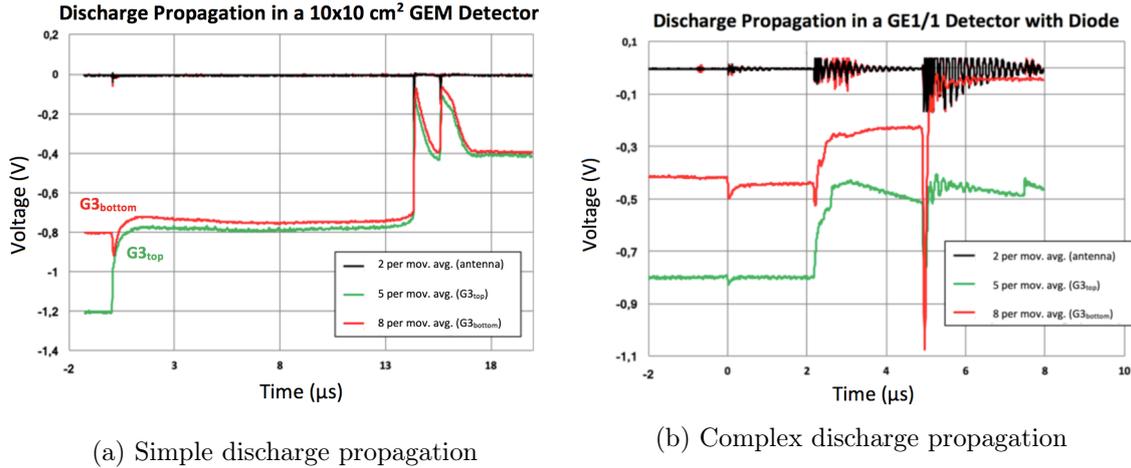


Figure 5.13 – Examples of simple and complex propagating discharges in different-sized GEM detectors [85]

Here, the top and bottom G3 electrodes have been connected to an oscilloscope, seen in green and red respectively. The collected data has been smoothed out using a moving average of five points and eight points, respectively. An antenna has been placed at the HV pads of the GEM detector in order to amplify any discharge signal, and is seen in black, with the collected data plotted using a moving average of two points. At $\sim 0 \mu s$, we see a primary discharge occurring in G3, marked by the first downtick in the antenna line where the voltages of the top and bottom foil electrodes equalize due to the short. At $\sim 16 \mu s$, we see the first instance of discharge propagation, marked with the second downtick in the antenna line, with both G3 electrodes going to 0 V as they short with the readout board and ground plane. Finally, at $\sim 17 \mu s$, we see a reignition of the propagation, marked by the third and final downtick in the antenna line and the zeroing of the G3 electrode voltages again.

However, in large detectors such as the GE1/1 detectors, the discharge propagation process is significantly more complex, as seen in Figure 5.13b. Here, after the primary discharge in G3 at $\sim -0.75 \mu s$ (marked by the first downtick in the antenna line), there is backwards propagation to G2 at $\sim 0 \mu s$ (marked by the second downtick in the antenna line), followed by a forward propagation back to G3 before propagating onwards to the readout board at $\sim 2-3 \mu s$, and finally a reignition of the propagation beginning at $\sim 5 \mu s$. As the discharge travels backwards and forwards, it accumulates more energy, which can be seen in the increasing amplitudes of the oscillating structures seen on the antenna line in Figure 5.13b.

Although Figure 5.13b is provided as a direct analogue to Figure 5.13a, what is shown is not, by itself, enough to understand the propagation of discharges within a large GEM chamber. The mechanisms involved are complex enough that in order to understand the full picture, one must look at the top and bottom of each GEM electrode. This fuller picture is given in Figures 5.14 and 5.15. Figure 5.14 is labeled at each stage of the discharge process according to where the propagation is occurring within the framework of the detector, as shown in Figure 5.15.

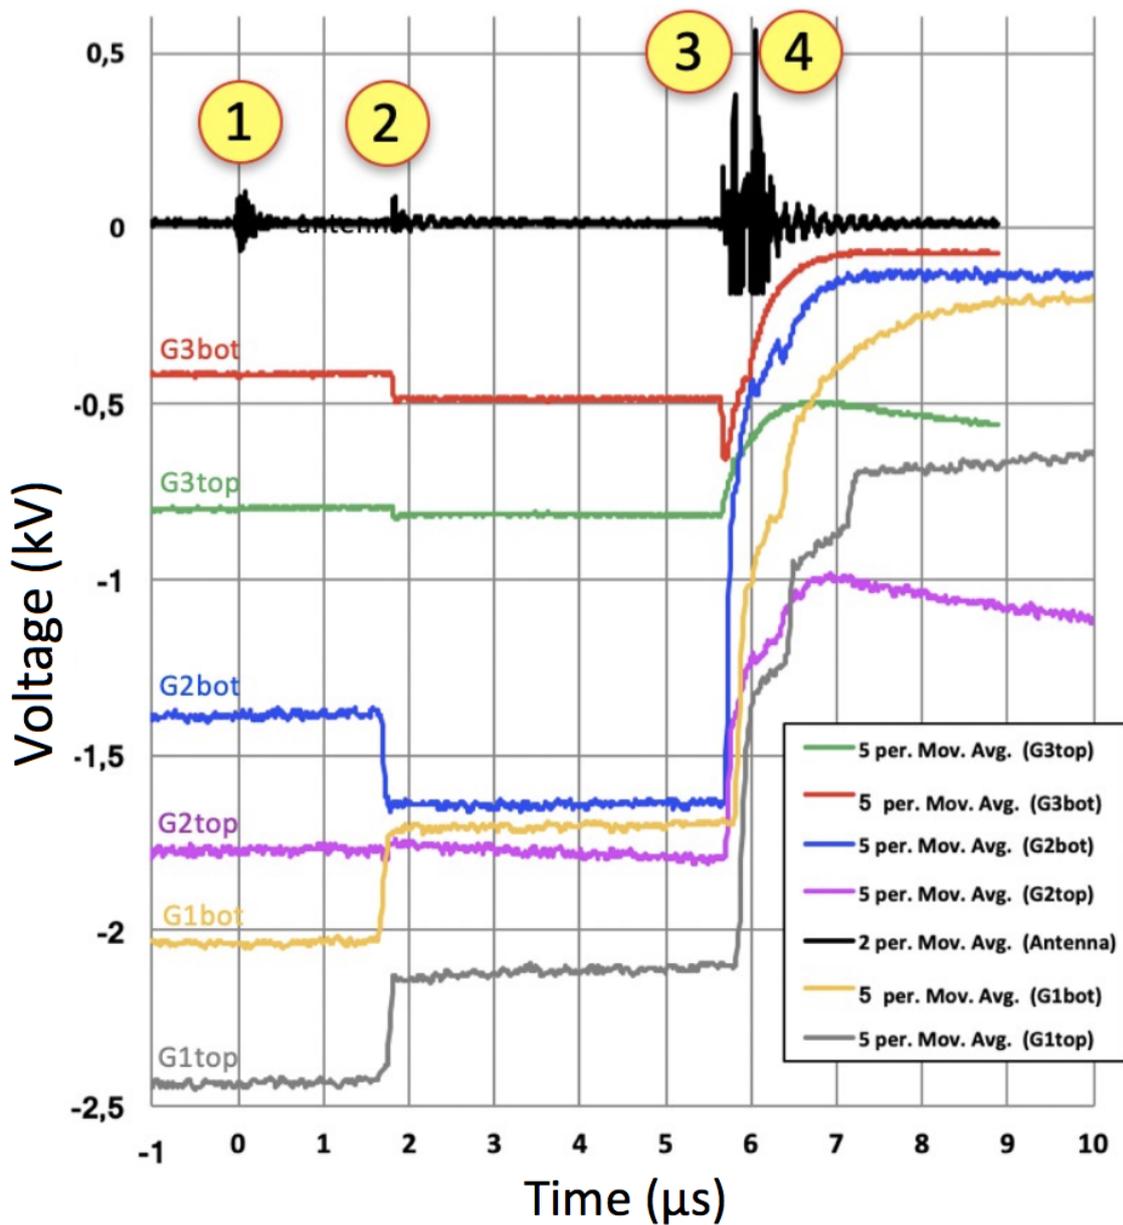


Figure 5.14 – Measurement of the tops and bottoms of all electrodes within a GE1/1 detector during a propagating discharge [94]

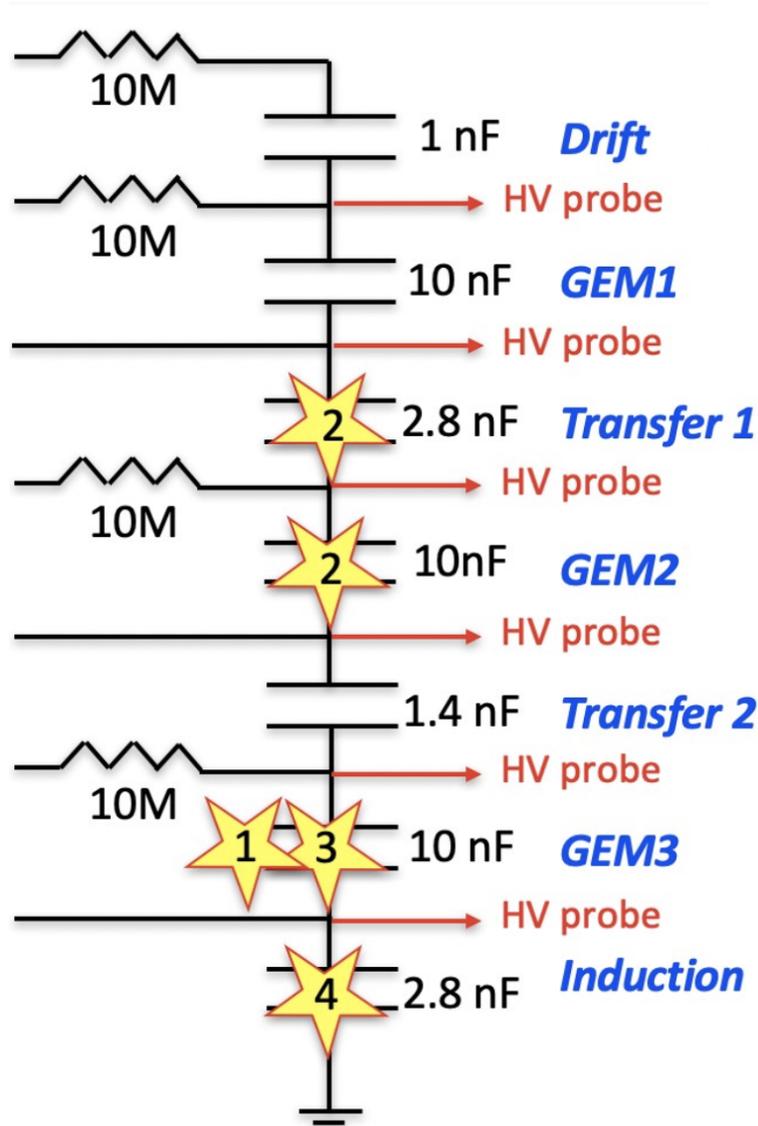


Figure 5.15 – Complementary to Figure 5.14, the locations of each discharge event within the electrical framework of the detector [94]

Here, the initial discharge occurs at the first disturbance in the antenna line, marked with a “1”. As previously described earlier in this section, this primary discharge occurs in G3, as this is the only place within the GEM detector that the Raether limit is exceeded. In this case, however, this cannot be directly deduced from Figure 5.14, but must be inferred from this prior knowledge. The reason for this is largely due to an effect that occurs when the capacitance within a GEM detector is large. In this case, the typical voltage fluctuations are not measured during a given discharge event, despite the discharge occurring as expected. An example of this is shown in Figures 5.16 and 5.17, where a $10 \times 10 \text{ cm}^2$ GEM detector was equipped with a 10 nF induction capacitor and then probed at the top and bottom G3 electrodes and the readout board, with and without this induction capacitor. In cases of both propagating and non-propagating discharges, the typical voltage structures on $G3_{top}$ and $G3_{bottom}$ were masked for measurements on the detector when the 10 nF induction capacitor was present. The physical reasons for this phenomenon are still under study: see Reference [94].

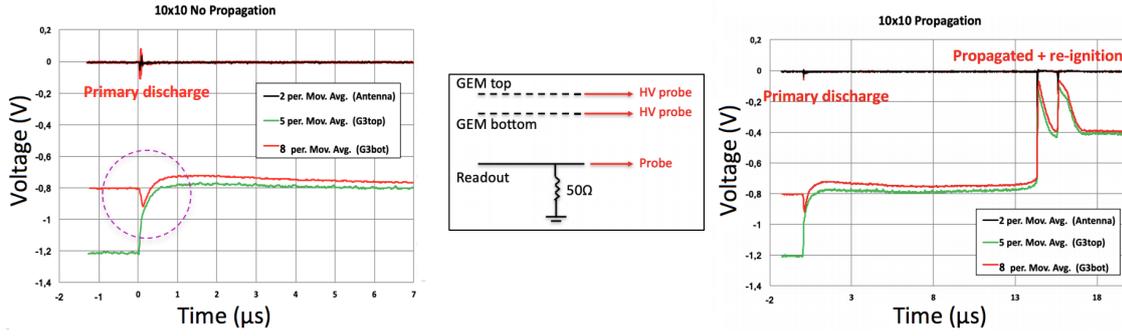


Figure 5.16 – Non-propagating and propagating discharges in a $10 \times 10 \text{ cm}^2$ GEM detector without additional induction capacitance, showing the usual HV fluctuations associated with the discharges [94]

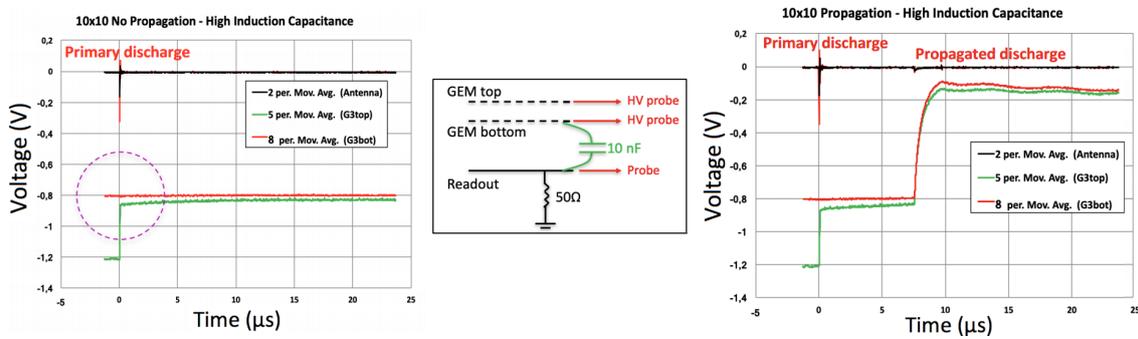


Figure 5.17 – Non-propagating and propagating discharges in a $10 \times 10 \text{ cm}^2$ GEM detector with 10 nF additional induction capacitance, lacking the usual HV fluctuations associated with the discharges [94]

Following the initial primary discharge, there is a backwards propagation to G2 and to the transfer gap between G1 and G2, marked with a “2”. In this case, we can deduce that it propagates to G2 by the slight increase in voltage on $G3_{top}$ and $G3_{bottom}$ (noting that the voltage polarity in all cases except for the antenna is negative) and the greater increase in voltage on $G2_{bottom}$, and that it reaches the transfer gap by the decrease in voltage on $G1_{top}$ and $G1_{bottom}$, with $G1_{bottom}$ nearly equalizing with $G2_{bottom}$. No change is seen on $G2_{top}$, though this is likely again due to the masking effect of the foil’s capacitance.

There then occurs a forwards propagation to G3, marked with a “3”. As the previous backwards propagation to G2 allowed the discharge to build up energy, the disturbance in the antenna line is of a much greater amplitude than the initial primary discharge within G3. Again, we can deduce this forwards propagation from the increase in voltage on $G3_{bottom}$, assuming the increase on $G3_{top}$ is masked by the capacitance effect, and the equalization of this voltage with that of $G3_{bottom}$.

Finally, soon after this forwards propagation, we see a propagation of the discharge to the readout board, with the voltage on all electrodes sharply decreasing towards zero.

Although every electrical discharge within the GEM detector has the potential to cause damage, as shown in Figure 5.6b, it is these propagating discharges which are of the greatest concern, as they possess not only the ability to cause damage to the GEM foil, but also the possibility to cause damage to the front-end readout electronics in such a way as to cause channel loss such as that seen in the slice test.

Chapter 6

Sustained Operations Investigations

6.1 Introduction

As the lifetimes of the CMS GEM systems are foreseen to be on the order of decades, one of the most crucial tasks of the slice test was to ensure that the GE1/1 detectors would be able to function long-term in a stable and efficient manner. In order to do that, a task force was formed in mid-2017, the “Sustained Operations” group, to monitor the health and behavior of the slice test detectors within CMS. The primary way of monitoring this was through the creation of “time series S-curves”, which plotted the daily S-curves taken by the DOC shifters in chronological order, allowing for an easy overview of the behavior of the chamber during the chosen time period. In particular, they allowed for the monitoring of any significant changes in noise levels, persistent loss of communication, and problems with individual channels such as unexpected masking due to poor channel health.

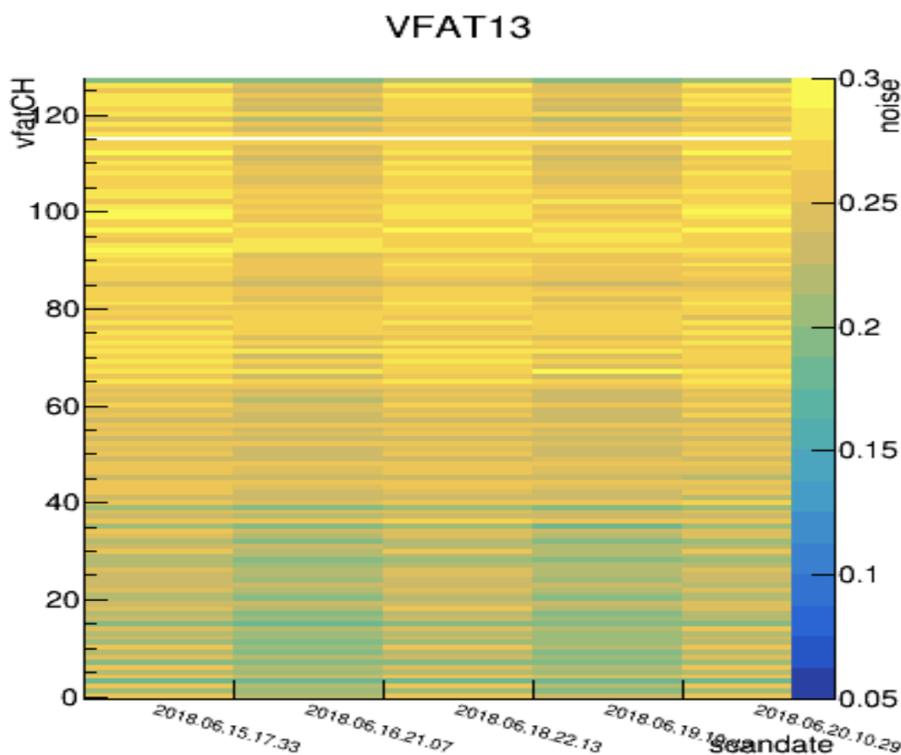


Figure 6.1 – Example of a time series S-curve plot

Figure 6.1 is an example of such a time series plot. This figure is the time series plot of VFAT13 of GEMINIm30L1 the week of June 15th, 2018, taken with settings $ztrim=4$ and $vt1bump=0$. Each of the five daily S-curves is presented as a column of data, with the scandate represented on the x-axis in format yyyy.MM.dd.mm.hh, and the VFAT channel number (or, alternatively, strip number) shown on the y-axis. The z-axis is the noise of each channel in each S-curve, ranging from 0.05 fC to 0.3 fC. Although this plot was chosen largely at random by the author from a host of many such S-curves, it can serve as a prime

example for what the sustained operations task force was looking for.

As an example, over the five days of S-curves within the plot, we can see that a pattern emerged, with fluctuations in noise between days, higher on days 1, 3, and 5 than 2 and 4. As such a fluctuation is ephemeral rather than a permanent change in the noise levels of the chamber, this could indicate that the performance of the chamber was affected by conditions within the CMS detector at the time of the scan, or could indicate a change in detector operations (such as HV on or off). This is further validated by the existence of this fluctuation in all VFATs of both chambers in the GEMINIm30 superchamber during this time period. If, on the other hand, the change was a permanent one, with the noise increasing significantly between scans and remaining so for all future scans in the time period, that would be a cause for concern that the detector has been negatively affected or damaged in some way.

One other crucial piece of information that can be gleaned from this plot is the masking of VFAT channel #115, as indicated by the white strip across the entire time series plot. This masking is defined at the time of the S-curve analysis, using the class *MaskReason* within the analysis script *AnaInfo.py*. A select portion of this script is reproduced below¹:

```
#: The default values for the cuts that determine the scurve fit quality masks
maxEffPedPercentDefault=0.02
highNoiseCutDefault=1.5
deadChanCutLowDefault=0
deadChanCutHighDefault=0

...

class MaskReason:
    #: The channel is not masked
    NotMasked = 0x0

    #: The channel was identified as an outlier using the MAD algorithm
    HotChannel = 0x01

    #: The S-curve fit of the channel failed.
    FitFailed = 0x02

    #: The channel has a burned or disconnected input.
    DeadChannel = 0x04

    #: The channel has an scurve sigma above the cut value.
    HighNoise = 0x08

    #: The channel has an effective pedestal above the cut value.
    HighEffPed = 0x10
```

In the first portion of the above code, we see that certain definitions have been made for the default cuts on S-curve masks. The maximum effective pedestal may not exceed 0.02, or 2%, and the noise of a channel may not exceed 1.5 fC. If these limits are exceeded, then the class *MaskReason* is invoked. As an example, if the effective pedestal for a channel exceeds the default maximum value of 0.02, it would be masked with the mask reason value 0x10, or 0b10000. If the noise of a channel exceeds the default value of 1.5 fC, then it would be masked with the mask reason value 0x08, or 0b1000. The class is additive,

¹The full script can be found at <https://tinyurl.com/anaInfo-py>.

as well, such that if a given single channel exceeded both the effective pedestal and noise limits, it would give a mask reason value of 0x18, or 0b11000.

Within this class is a slightly more complicated but no less important mask reason: HotChannel. This particular mask reason, unlike the others, does not rely on a simple hard-coded cut value. Rather, it is the result of an algorithm which detects outlier values in S-curves via the “mean absolute deviation” (MAD) method.

In the MAD method, the maximum threshold (VT1) value is computed for each channel by scanning from DAC values of 255 to 0, and taking $VT1_{max}$ to be the first non-zero bin which is encountered. Thus, for a distribution of $X = \{x_1, x_2, x_3, \dots, x_N\}$, we can compute the median value of this set via Equation 6.1, and then define the MAD via Equation 6.2.

$$\tilde{m} = median\{x_i\} \quad (6.1)$$

And then define the

$$MAD = median\{|x_i - \tilde{m}|\} \quad (6.2)$$

Then, a given channel’s S-curve is considered an outlier if it satisfies Equation 6.3, where the variable zScore is given as the multiple value of the noise, σ , which the user has chosen to consider (Generally 4 by default during slice test operations).

$$x_i > \tilde{m} + (0.6745 \cdot zScore \cdot MAD) \quad (6.3)$$

Of all of the mask reasons, however, the reason of greatest concern is 0x04, “DeadChannel”. As stated in the comments of the code above, this indicates that the input of the channel has either been burned or disconnected. Either case presents a serious physical issue that cannot be remedied while the detector is installed within CMS, and thus presents a significant problem for maintaining the stable lifetime of the detector.

Thus, if a masked channel such as the one in Figure 6.1 is detected in a given time series plot, we can examine the mask reason that was given and take the appropriate actions from there.

6.2 Environmental Effects on Detector Health

In addition to understanding when a channel’s health changes (the scandate), and how it changes (the mask reason), we can understand what conditions caused it to change in the first place. As previously mentioned, Figure 6.1 shows a fluctuation in noise levels that could be attributed to differing conditions within the CMS detector at the time of the scan. As well, in the particular case of the “DeadChannel” mask reason, the environmental conditions may directly lead to the cause of the channel’s death, whether via discharges or otherwise.

6.2.1 Beam Luminosity

As discussed in Section 5.2, electrical discharges within the active volume of a gas detector present a real and significant potential for causing lasting damage to the detector, both to the holes in the foils and, if they propagate to the readout board, to the readout electronics as well. This is what is referred to as a “burned input” in the MaskReason class, as the electrical discharge causes physical burn damage to the input of the VFAT. From Section

5.1, we know that a discharge occurs when the space charge of an electron avalanche overcomes the Raether limit, thus developing from an avalanche to a plasma filament and causing a primary discharge within the detector. Overcoming this limit requires a high charge density within a small area of the detector, and thus a discharge may be more likely to occur when the detector is under such conditions that would favor such a high charge density - for example, conditions of high beam luminosity which might create highly-ionizing particles, or a high concurrent rate of other lesser ionizing particles, which interact with the detector.

6.2.2 Detector Operational Conditions

Of course, regardless of the luminosity present within the area of the slice test, a discharge cannot grow from a Townsend avalanche if the conditions are not right for such a growth. Thus, when looking for a discharge which may be correlated with detector damage, the only sensible times to consider are those where the detector was powered on. Without voltage being applied to the foils, and without the electric fields present between and within them, no electron multiplication will occur, and thus there is no possibility for electrical damage. Beyond the simple question of on or off, higher applied voltages and electric fields, and thus higher gain settings, can increase the likelihood of a discharge event occurring.

6.2.3 Magnetic Field

Although discharge damage will not occur while the detectors are off, this does not mean that these periods are to be entirely discounted in the search for detector damage. Indeed, in exploring the “disconnected” channel portion of the “DeadChannel” mask reason, we can also consider the physical condition of the detector at any given time. Consider, for example, the harsh magnetic field conditions of the CMS solenoid magnet ($|\vec{B}| = 3.8 \text{ T}$). Although every effort is put into using non-magnetic materials in the production of the GEM detectors, the presence of this relatively large magnetic field can still pose a problem. In one instance, for example, it was noticed that the standoffs used in the GEM detector production are very slightly magnetic. And although all of the hardware which is installed in UXC55 is required to be functional in the long-term under these harsh magnetic field conditions, that does not mean that it is entirely unaffected by it. The endcap muon detectors of CMS are interleaved with a twelve-sided, three-layered iron return yoke which surrounds the magnet’s coils and contains and guides the field. Being ferromagnetic, this yoke is vulnerable to flexing when the solenoid magnet is ramped up or down and the magnetic field subsequently changes. Although this flexing is comparatively small, it is not negligible, and thus should not be removed from consideration. If one of the Panasonic connectors was not properly connected, and thus maintained a loose electrical connection between the readout board strip and the VFAT and/or GEB, this flexing of the disk could cause that tenuous connection to completely disconnect, causing the affected channel(s) to appear dead. Although this is unlikely, given the 128 pins present in the small footprint of the Panasonic connectors, if the dead channels were connected to the input pins on either outer edge of the connector, this could be a potential reason as to why.

Similarly, if a piece of debris (dust or other small contamination) were to have made it past the quality control steps, the flexing of the disk could cause the debris to move, potential blocking the signal path and making the channel appear dead. Although the strict quality controls in place make this similarly unlikely, it may serve as an explanation

for channels which are masked for being “dead” and then subsequently recover, as we will see in Section 6.5.6.1.

6.3 GEMINIm01L1_{v2}

The GEMINIm01 superchamber presented a unique opportunity for the sustained operations group, being the only superchamber which was removed at the midpoint of the slice test. This allowed the sustained operations team to physically inspect the chamber for damage, while still having time remaining in the duration of the slice test to make and test changes with the other detectors if necessary.

During the slice test, VFAT4 of GEMINIm01L1 exhibited odd behavior. As seen in Figure 6.2, four groupings of readout strips had significantly narrower S-curves than the rest of the strips. While narrow S-curves are normally indicative of a good signal-to-noise ratio, these groupings are uncharacteristic for the VFAT and abnormally narrow. More specifically, they represent a noise level of between 0.0414 and 0.109 fC, a range which indicates that, rather than being low-noise, these strips were disconnected from the input of the VFAT in some way, such that the capacitance of those strips was not being properly seen.

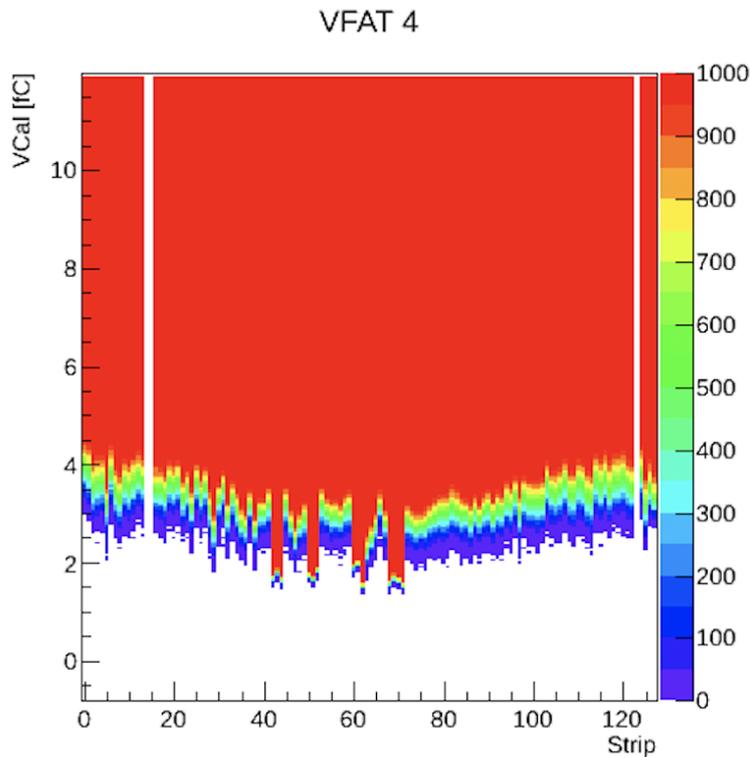


Figure 6.2 – S-curve plots of GEMINIm01L1_{v2} VFAT4, taken within the CMS endcap

Thus, when the chamber was brought back to the lab, that VFAT was removed for investigation. When placed under a microscope, several instances of physical damage were observed, which are shown highlighted in pink in Figure 6.3.

In three cases, flecks of burnt material were observed as being present on the lower portion of the VFAT2 chip, and in the top righthand corner, solder from a group of bonding wire pads has flowed off of the pads and onto the VFAT chip itself. These four instances of visible damage are consistent with the four groups of narrow S-curves. For reference,

Figure 6.4 shows the location of the microscopic view on the VFAT chip relative to its location on the v2 hybrid PCB.

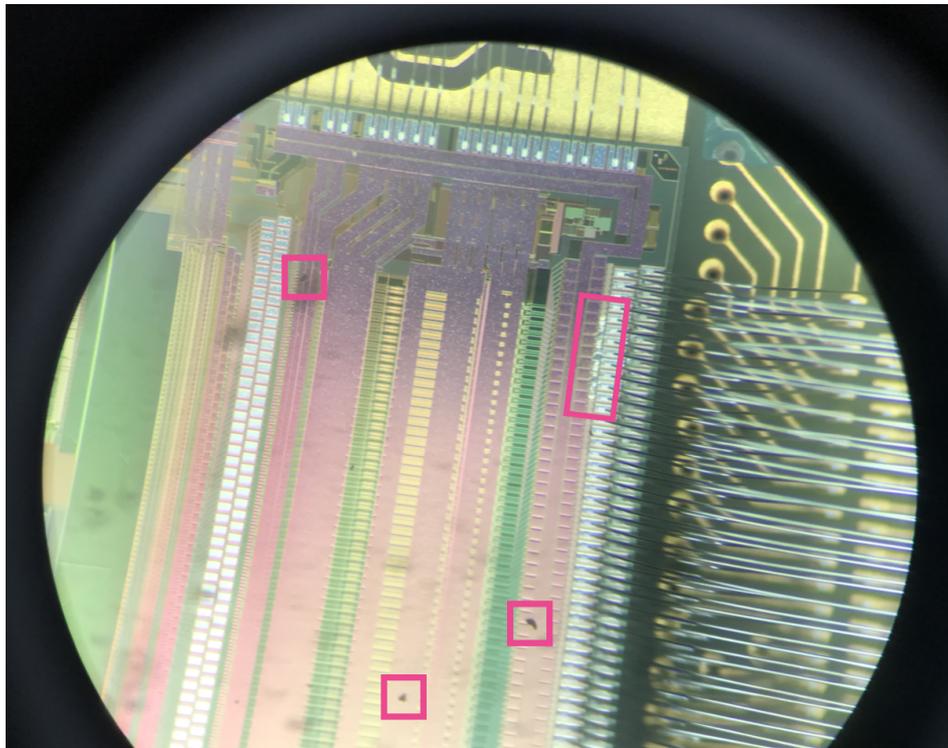


Figure 6.3 – Photo of GEMINIm01L1_{v2} VFAT4 under a microscope, showing instances of burn damage and solder flow

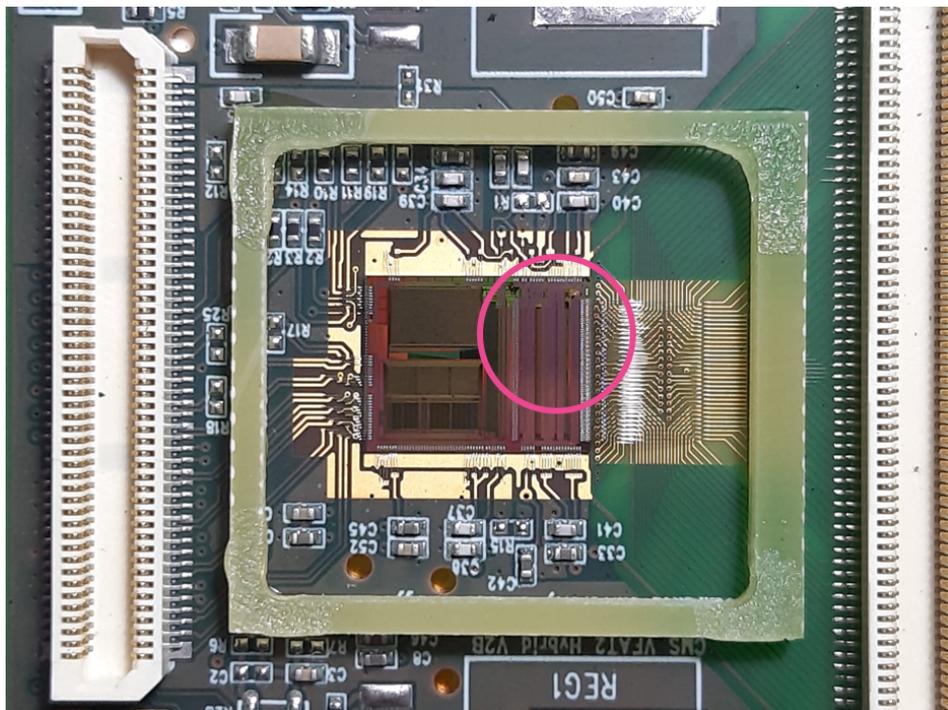


Figure 6.4 – Photo of a VFAT2 chip installed on a v2b hybrid PCB, showing the approximate location of the microscopic view of Figure 6.3. Note that this is not the same chip as in Figure 6.3 but is of the same type

As the detectors are kept under a cooling circuit at all times during their operation, the presence of burn damage and melted solder indicates that one or more electrical discharges must have taken place, which caused enough heat to melt the solder and burn off some of the silicon of the VFAT. The GEMINIm01 detectors present the unique opportunity to search for these discharges, as they were equipped with a CAEN A1515TG multi-channel power supply which allowed for the individual monitoring of the current and voltages present on each of the GEM foil electrodes. If a discharge occurred during the operation of the detector, it would be easily viewed through these monitoring capabilities.

Unfortunately, this opportunity was lost in the case of GEMINIm01L1. Figure 6.5 shows the time series S-curves for GEMINIm01L1 for the entirety of its operational presence in the CMS endcap, beginning on April 25th, 2017 and ending on November 3rd, 2017.

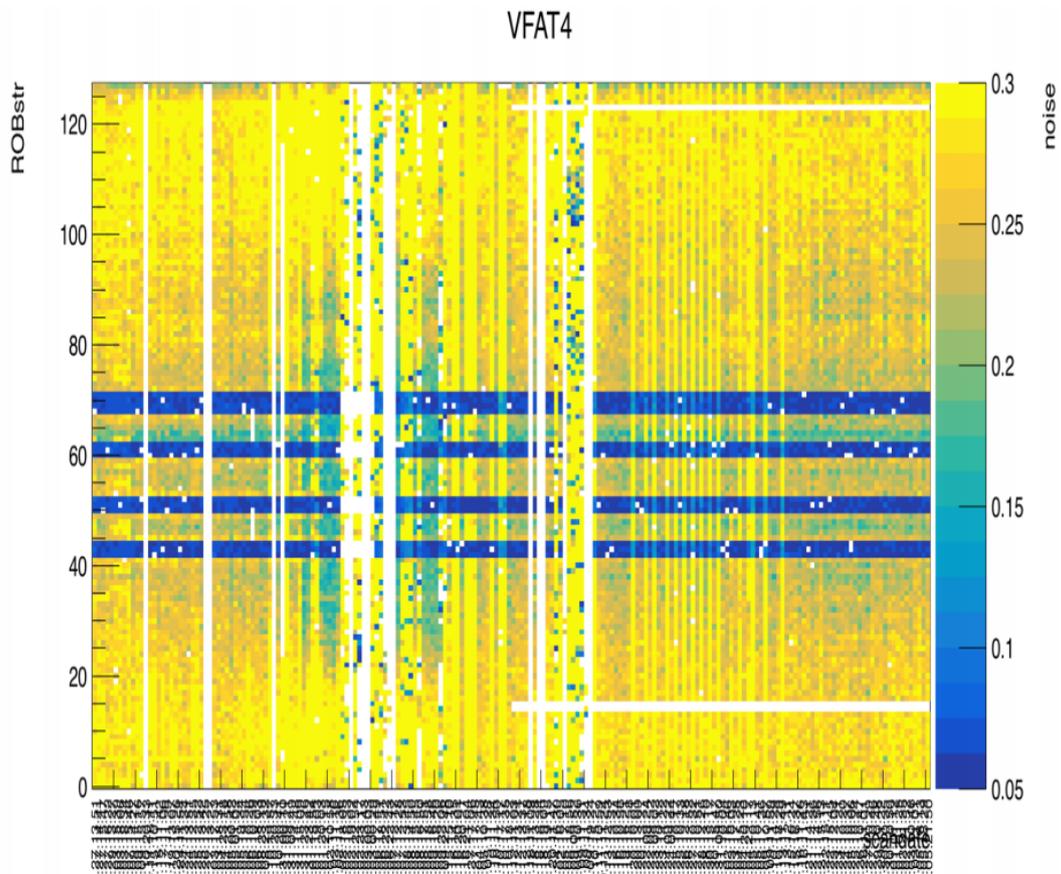


Figure 6.5 – Time series S-curve of GEMINIm01L1_{v2} VFAT4 from April 25th, 2017 through December 7th, 2017

Note that the y-axis is presented as readout board strip (ROBstr) number rather than VFAT channel number in order to get a better idea for the physical location of the damage, as the VFAT channel mapping is not necessarily the most intuitive. Here, we see, as expected from the S-curves shown in Figure 6.2, the four groupings of damaged channels: ROBstr #42 - 44, 50 - 52, 60 - 62, and 68 - 71. Unfortunately, the damage to these channels was present from the very first S-curve taken during the slice test, indicating that the damage occurred either in the QC8 stand during qualification procedures, where, unlike at P5, no monitoring software or database was present. Thus, no determination can be made in regards to the alleged discharges which caused the observed damage.

The plot also shows two groupings of masked strips, ROBstr #14 - 15 and 123, which

were previously healthy and became masked from scandate 2017.07.10.11.34 until the end of the detector’s operational lifetime within the slice test. Should these strips have been masked using the “DeadChannel” mask reason, these would provide the desired opportunity to study the channel loss. However, this was not the case - for all three masked strips, the mask reasons provided were HighEffPed, indicating the channels had an effective pedestal over the cut value of 0.02, and HotChannel, indicating the channels were identified as outliers using the “MAD” detection method, described in Section 6.1. Thus, although these channels were problematic for their own reasons, they were not damaged via an electrical discharge, and are thus of less interest to this study.

All is not lost, however, as we turn to GEMINIm01L1’s sister chamber, GEMINIm01L2.

6.4 GEMINIm01L2_{v2}

6.4.1 Channel Loss

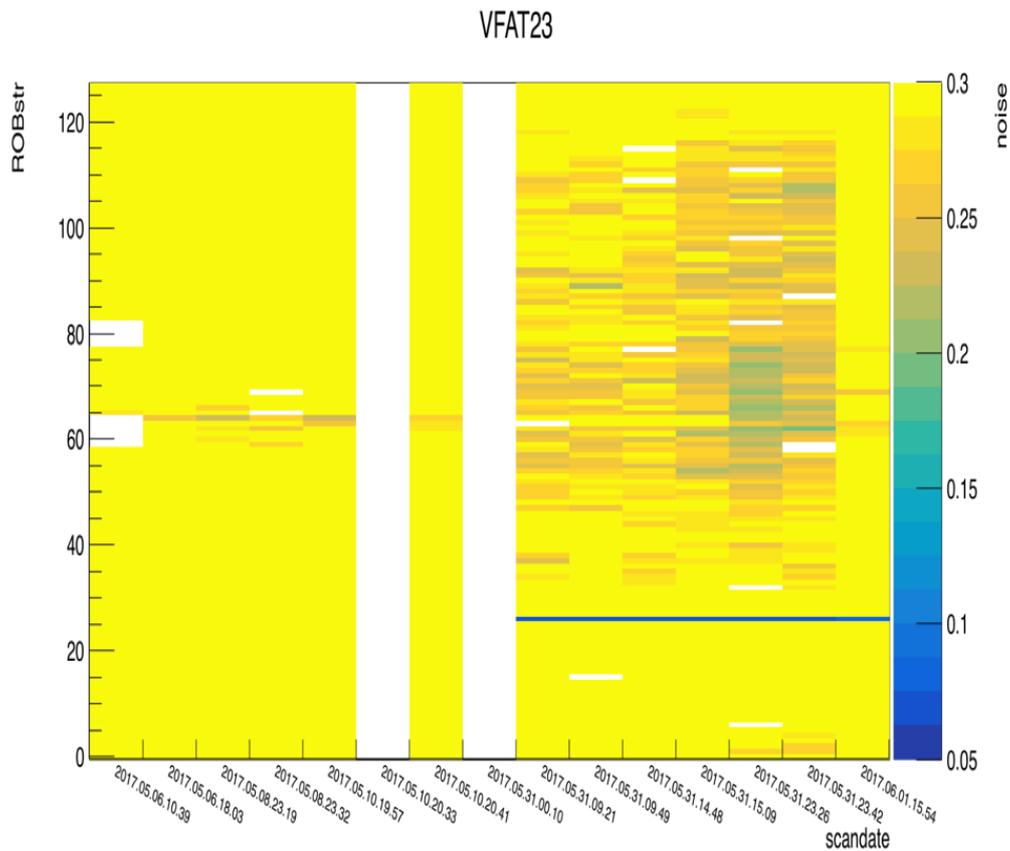


Figure 6.6 – Channel Loss in GEMINIm01L2_{v2} VFAT23

During the period of operation of the v2 GEMINIm01 superchamber, April 25th, 2017 through December 7th, 2017, there was only one occurrence of a previously functioning channel burning during operation. This channel, corresponding to readout board strip #26 on VFAT23 of GEMINIm01L2, burned at some point during the two-week period of 2017.05.10.20.41 and 2017.05.31.09.21. The time series S-curve illustrating this can be seen in Figure 6.6. Prior to this occurrence, it was not standard protocol for the DOC to take daily S-curves, leading to the three-week period in which there were no S-curves

taken. This period unfortunately corresponded with the moment of channel loss, making pinpointing the exact moment of damage difficult to discern.

However, despite this large window of possibility for the damage-causing event, the presence of the multi-channel power supply and monitoring software nevertheless provided an opportunity to study the time period in detail in hopes of uncovering the discharge.

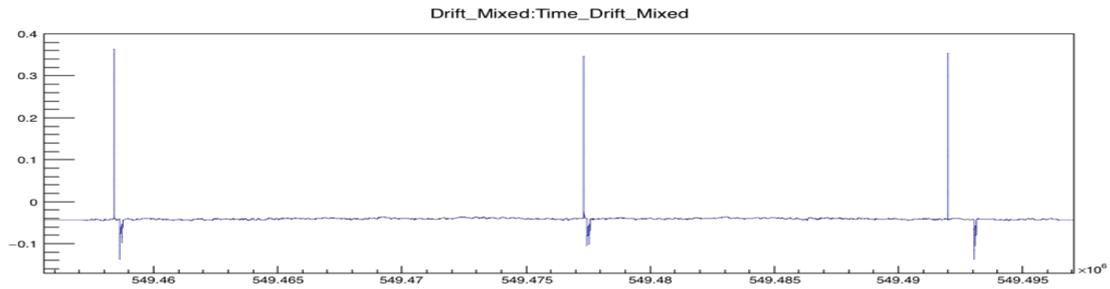
6.4.2 Detector Operation Periods

This large window of time can first be narrowed down by considering only time periods in which the detector was powered on, as per Section 6.2.2. From 2017.05.10.20.41 through 2017.05.31.09.21, GEMINIm01L2 was on a total of nine times, as listed in Table 6.1. Of these nine periods, only one, 2017.05.30.11.59 to 2017.05.30.22.54, was during the presence of LHC beam. Thus we can narrow our search even further, as per Section 6.2.1.

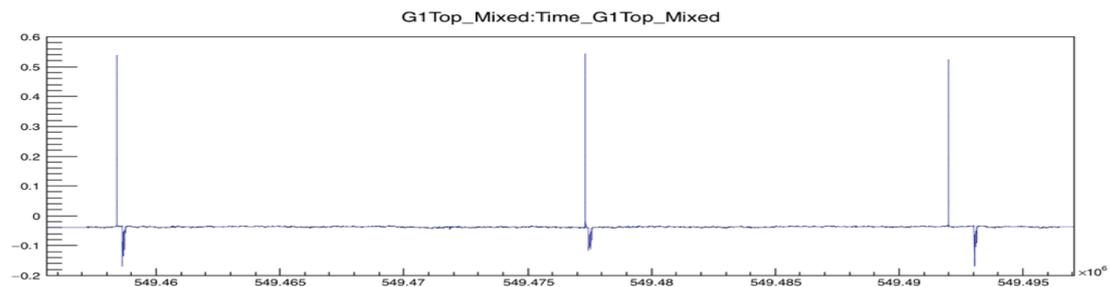
Turned on	Turned Off	Beam?
2017.05.10.13.55	2017.05.11.00.51	No
2017.05.11.09.36	2017.05.11.22.41	No
2017.05.12.08.17	2017.05.12.17.00	No
2017.05.12.18.01	2017.05.14.01.45	No
2017.05.14.09.35	2017.05.15.12.31	No
2017.05.29.17.50	2017.05.29.22.47	No
2017.05.30.11.59	2017.05.30.22.54	Yes
2017.05.31.08.56	2017.05.31.13.52	No
2017.05.31.14.28	2017.05.31.20.58	No

Table 6.1 – Periods of operation for GEMINIm01L2 during the possible window of discharge damage to VFAT23

6.4.3 Potential Damage-Causing Events



(a) Current measurements of GEMINIm01L2 Drift



(b) Current measurements of GEMINIm01L2 $G1_{top}$

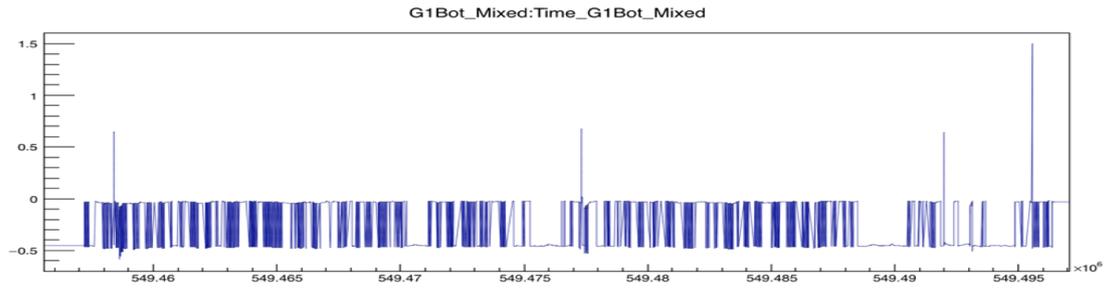
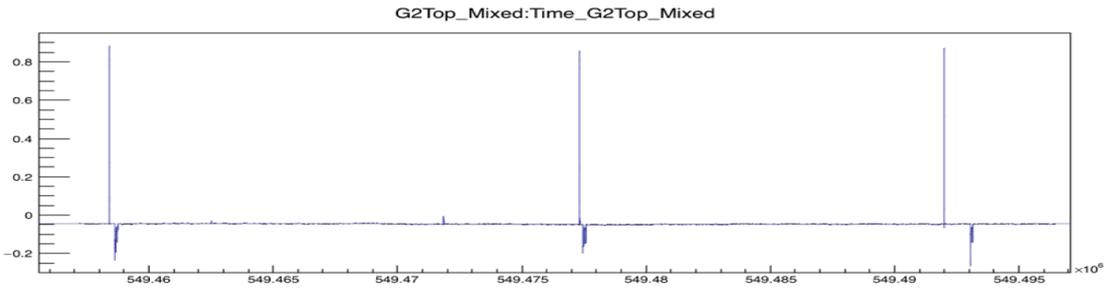
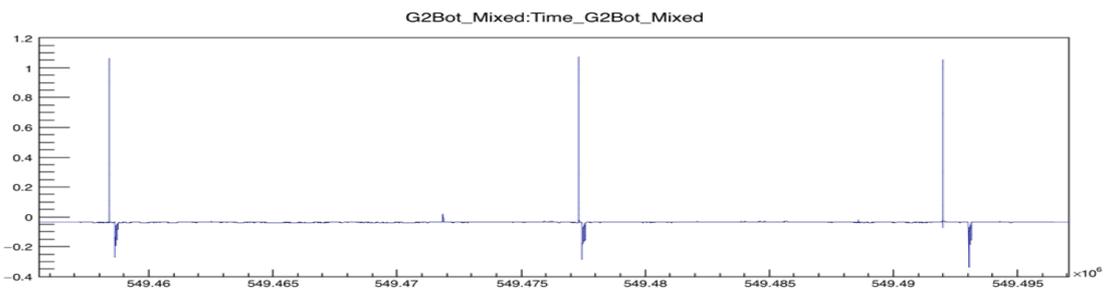
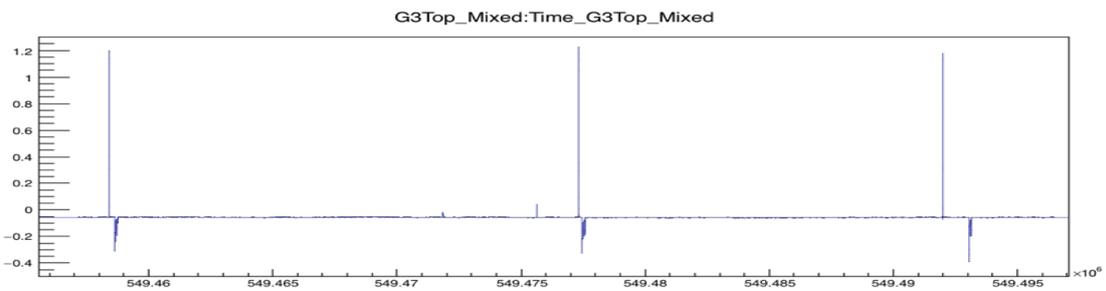
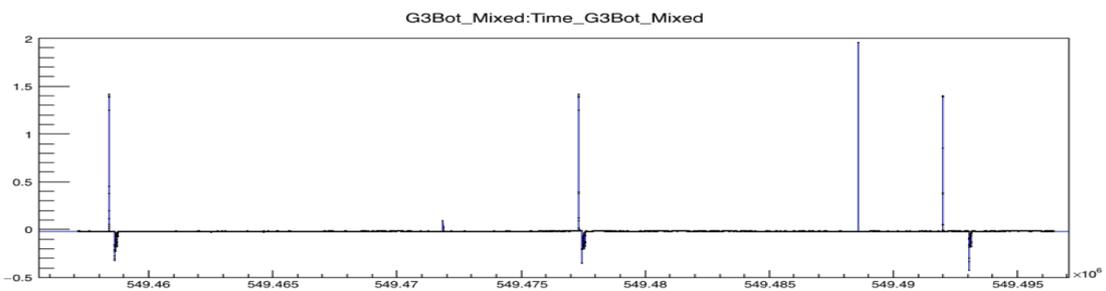
(c) Current measurements of GEMINIm01L2 $G1_{bottom}$ (d) Current measurements of GEMINIm01L2 $G2_{top}$ (e) Current measurements of GEMINIm01L2 $G2_{bottom}$ (f) Current measurements of GEMINIm01L2 $G3_{top}$ (g) Current measurements of GEMINIm01L2 $G3_{bottom}$

Figure 6.7 – Current measurements of all seven monitored electrodes of GEMINIm01L2 during the period of 2017.05.30.11.59 to 2017.05.30.22.54

Figure 6.7 shows the current measurements on the tops and bottoms of all three foils, as well as the drift board, during this time period. In all seven plots, the y-axis is the current measurement in μA and the x-axis is the date in modified Unix epoch time, which is the Unix epoch timestamp minus 946,681,200 seconds.

- Unix epoch time is the number of seconds that have elapsed since January 1st, 1970 (midnight UTC/GMT), not counting leap seconds
- Modified Unix epoch time is the number of seconds that have elapsed since January 1st, 2000 (midnight GMT+01:00), not counting leap seconds

Thus, a timestamp such as 2017.05.30.11.59 = 1,496,138,340 - 946,681,200 = 549,547,140.

Here, we see several current anomalies which may be of interest. To begin, $G1_{bottom}$ shows anomalous behavior which is not present in any of the other graphs. Throughout the time period under study, the current measurement on this electrode fluctuated between approximately $0 \mu\text{A}$ and $-0.05 \mu\text{A}$. Although the time between fluctuations varies, the maximum and minimum values do not. This fact leads us to the conclusion that this is not a physical current fluctuation at all. Rather, this is more likely due to the discretized nature of the ADC within the current monitoring system. If the actual current is at the edge of two ADC bins, where the least-sensitive bit cannot resolve it, the current will appear to fluctuate between two discrete values, as it does in Figure 6.7c. This particular behavior was observed on several other electrodes during time periods in which no damage or other detector misbehavior occurred, bolstering this explanation.

With this anomaly accounted for, we can look at the next largest anomaly - namely, three distinct sets of current spikes which span all seven electrodes. For the sake of conciseness, all seven electrodes are presented in one plot in Figure 6.8, with these three events labeled as “A”, “B”, and “C”.

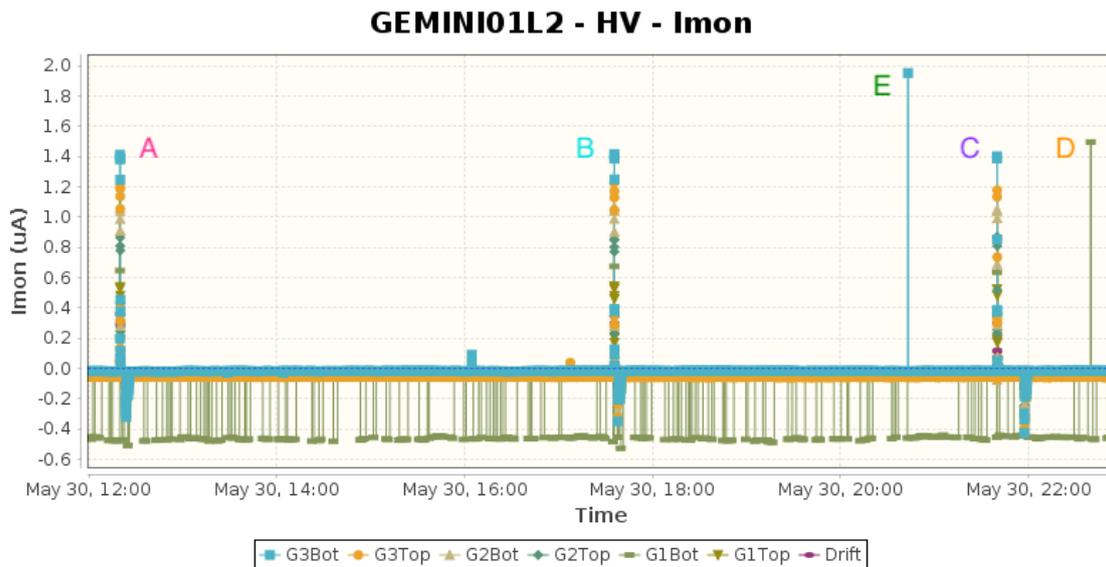


Figure 6.8 – Current measurements of all seven monitored electrodes of GEMINIm01L2 during the period of 2017.05.30.11.59 to 2017.05.30.22.54, with each event of interest labeled A through E

Each set of current spikes has a set of features in common: The current spikes on each electrode are the same polarity, with an initial large positive spike followed shortly after

by a smaller negative spike. This in and of itself suggests that these are also not physical discharge events. From Section 5.2, we know that a discharge on a given GEM detector electrode results in a current spike of one polarity, with an opposite-polarity current spike on the neighboring electrode(s). If A, B, and C were true discharges, one would expect several of the spikes to be in the opposite polarity, which is not the case.

Furthermore, the amplitude of these spikes increase as we move down the layer stack, with the smallest amplitude at the Drift layer and the largest amplitude at the G3_{bottom} layer. Together, this points to these events being evidence of leakage currents, i.e. artifacts of the HV distribution scheme, rather than true discharges within the detector.

We can verify this by looking at the HV distribution scheme for the CAEN A1515TG power supply used for GEMINIm01, shown in Figure 6.9. The majority of the figure is drawn with industry-standard electrical components, with the exception of “SSS”, which represents solid stat switches. When in the OFF position, these switches are active, resulting in their giving electrical continuity in low impedance to the channel chain.

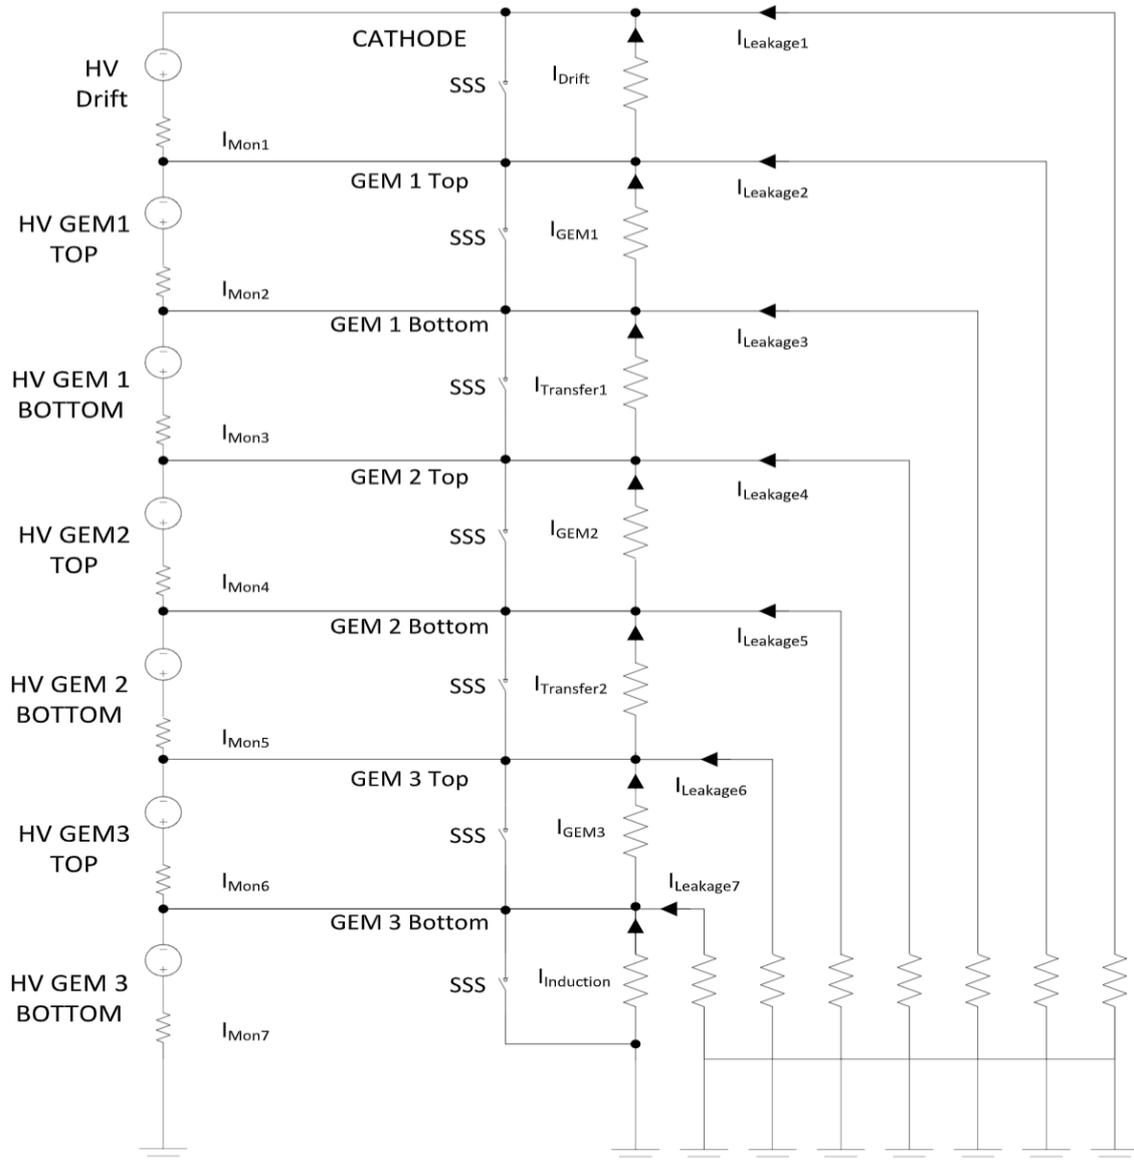


Figure 6.9 – HV distribution scheme for the CAEN A1515TG power supply

The distribution of current within this scheme is then given by Equations 6.4 through 6.10.

$$I_{Mon1} = I_{Leakage1} + I_{Drift} \quad (6.4)$$

$$I_{Mon2} = I_{GEM1} + I_{Leakage2} + I_{Leakage1} \quad (6.5)$$

$$I_{Mon3} = I_{Transfer1} + I_{Leakage3} + I_{Leakage2} + I_{Leakage1} \quad (6.6)$$

$$I_{Mon4} = I_{GEM2} + I_{Leakage4} + I_{Leakage3} + I_{Leakage2} + I_{Leakage1} \quad (6.7)$$

$$I_{Mon5} = I_{Transfer2} + I_{Leakage5} + I_{Leakage4} + I_{Leakage3} + I_{Leakage2} + I_{Leakage1} \quad (6.8)$$

$$I_{Mon6} = I_{GEM3} + I_{Leakage6} + I_{Leakage5} + I_{Leakage4} + I_{Leakage3} + I_{Leakage2} + I_{Leakage1} \quad (6.9)$$

$$I_{Mon7} = I_{Leakage7} + I_{Leakage6} + I_{Leakage5} + I_{Leakage4} + I_{Leakage3} + I_{Leakage2} + I_{Leakage1} \quad (6.10)$$

From these equations, we see that the Drift only sees contributions from $I_{Leakage1}$, whereas $G3_{bottom}$ sees contributions from $I_{Leakage1}$ through $I_{Leakage7}$, and each is of the same polarity. Thus, the increase in amplitude of our unipolar signals from Drift to $G3_{bottom}$ makes logical sense in this context.

Furthermore, if we cross-reference the times at which the spikes occurred with GEMINIm01L1, we see that each spike corresponds with a voltage trip of that chamber. Indeed, in Figure 6.10, we see that the positive spikes in GEMINIm01L2 correspond with the powering down of GEMINIm01L1, and the negative tails of those spikes correspond with GEMINIm01L1 powering back up. This helps illuminate why, in the case of spike C, the negative tail occurred at significantly later than those of spikes A and B, as in that case GEMINIm01L1 took longer to recover from the trip.

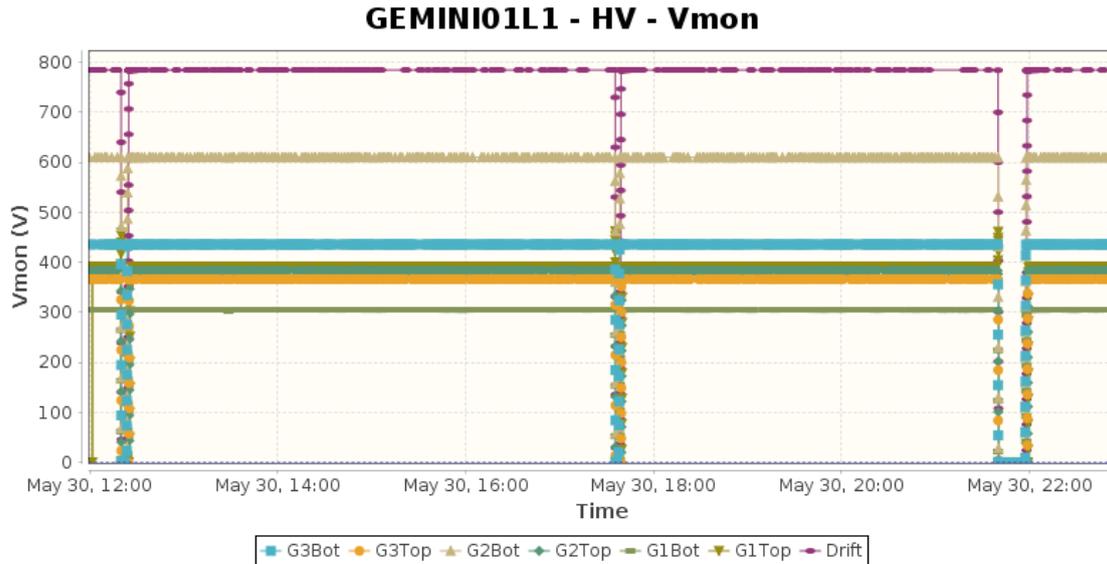


Figure 6.10 – High voltage measurements of all seven monitored electrodes of GEMINI01L1 during the period of 2017.05.30.11.59 to 2017.05.30.22.54

Thus, we can firmly rule these three events out as discharge events, which leaves only two remaining discharge candidates, labeled in Figure 6.8 as D and E. Unlike A, B, and C, neither D nor E are present on all of the electrodes, ruling out the possibility of their being leakage currents. However, this brings with it its own problems - if these are true discharges, then once again taking the information from Section 5.2 into account, one would expect to see an opposite-polarity spike on a neighboring electrode.

Once again, this brings us back to the ADC in the current monitoring chain. The ADC used has a finite sampling rate which is limited by its relatively long integration time of 500 ns. As electrical discharges within gaseous detectors are instantaneous, the missing opposite-polarity spike could have occurred during the “down-time” of the ADC, thus not appearing on the monitoring graph. Assuming this to be the case, we can consider other factors to either confirm or rule out these spikes as discharges.

In the case of D, its presence on $G1_{bottom}$ is in itself strange. From Section 5.2, we know that, in general, G3 is generally the only location within a GEM detector that the Raether limit is exceeded, causing a discharge from a Townsend avalanche. For this to be a real discharge, then, it would require some extenuating circumstances.

Thus, we first turn to the voltage monitoring software, to see what was recorded. At the same time as the recorded current spike, Figure 6.11 shows that the voltage monitoring software recorded a value on $G1_{bottom}$ of 85,899,344 V, which is of course a non-physical value, as the GEM detectors could not survive such a high applied voltage. This implies a glitch, rather than a physical discharge event, is responsible for spike D.

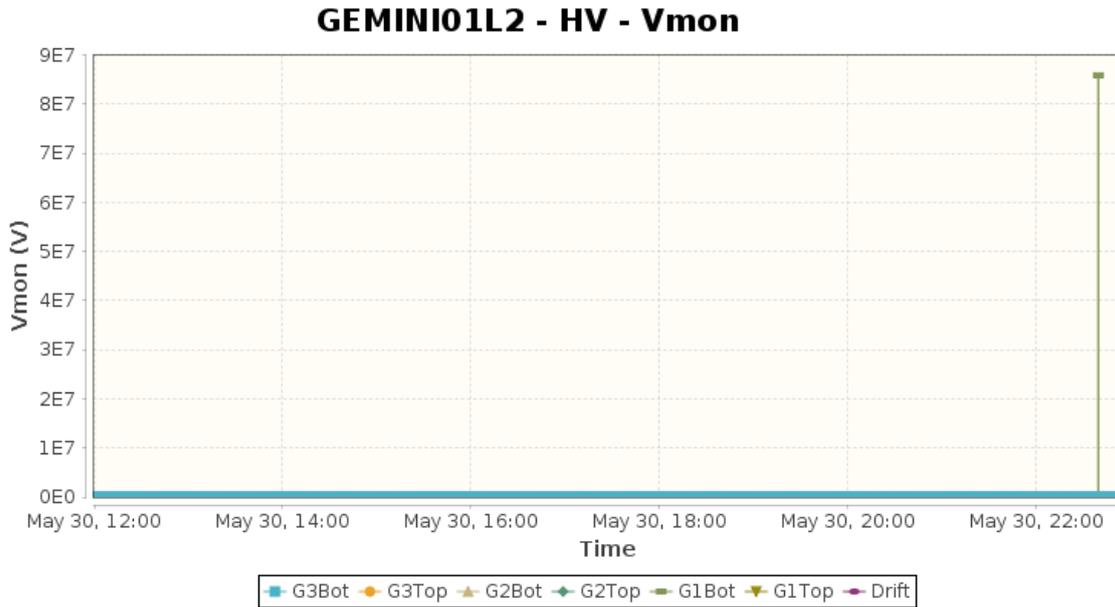


Figure 6.11 – High voltage measurements of all seven monitored electrodes of GEMINI01L2 during the period of 2017.05.30.11.59 to 2017.05.30.22.54

Indeed, when we look at the current monitoring for GEMINI01L2's sister chamber, GEMINI01L1, we see similarly problematic current values recorded at the same time, shown in Figure 6.12 as a 200 μA spike on $G2_{bottom}$ and a 50 μA spike on $G1_{top}$.

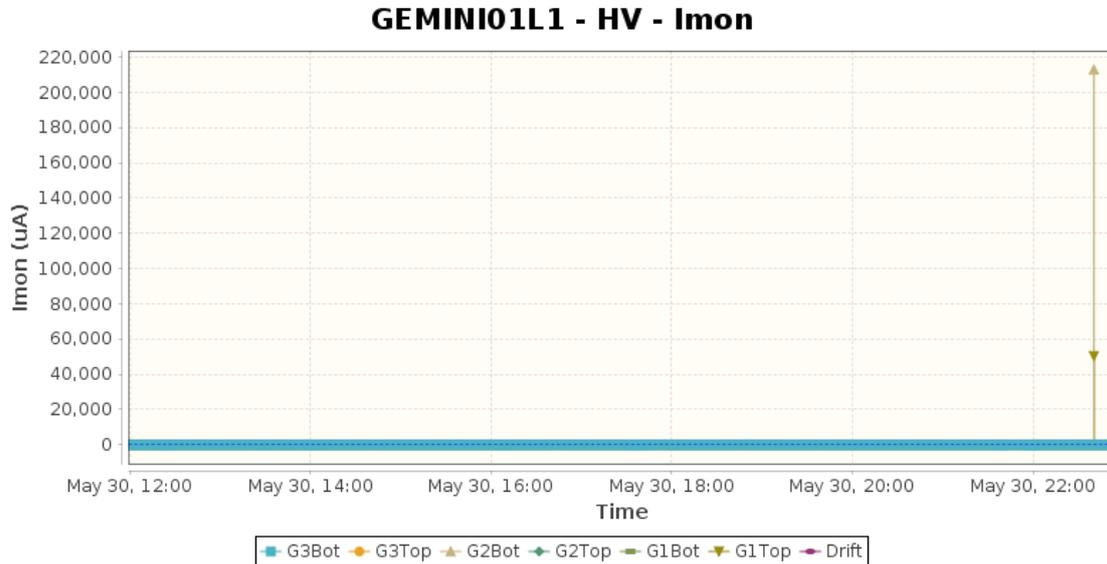


Figure 6.12 – Current measurements of all seven monitored electrodes of GEMINI01L1 during the period of 2017.05.30.11.59 to 2017.05.30.22.54

Given the severe unlikelihood of discharge events happening on both detectors at exactly the same time, and the impossibility of GEMINI01L2 being able to withstand 85.8 MV on a single foil, we can conclude that this observed spike is the result of a glitch within the monitoring software rather than a physical discharge within the detector.

This leaves only one remaining discharge candidate, labeled “E” in Figure 6.8. This candidate is of the most interest because, if a true discharge, its location on $G3_{bottom}$ makes it a prime candidate for the damage-causing event, as this is the electrode closest to the

anode and thus the most likely to be able to propagate to the readout electronics. Its location on $G3_{bottom}$ also means it is within the right area of the detector to represent a discharge caused by an avalanche exceeding the Raether limit.

Although we would again expect to see an opposite-polarity current on $G3_{top}$ at the same time, we can once again operate under the assumption that said current was missed as a result of the integration time of the ADC. And in cross-checking with the HV monitoring, as well as the current and voltage monitoring of GEMINIm01L1, we do not see any non-physical values that would lead us to believe this is the result of a glitch. It is safe, then, to conclude that this is a true discharge, and likely that which caused the loss of VFAT #26's readout channel, as no other such discharge event occurred either during this period of the detector being on with beam present, or any of the periods in which the detector was on but no beam was present.

With this discharge now established as our “smoking gun” for the channel damage, we turn to look at the conditions under which it happened, to see if we can determine what caused the discharge to occur in the first place.

6.4.4 Beam Luminosity

As per Section 6.2.1, we can begin our search with the beam conditions. The luminosity of the beam during the period of 2017.05.30.11.59 to 2017.05.30.22.54 is shown in Figure 6.13. Note that in the plot, the time is shown on the x-axis as the two-digit date and the two-digit hour.

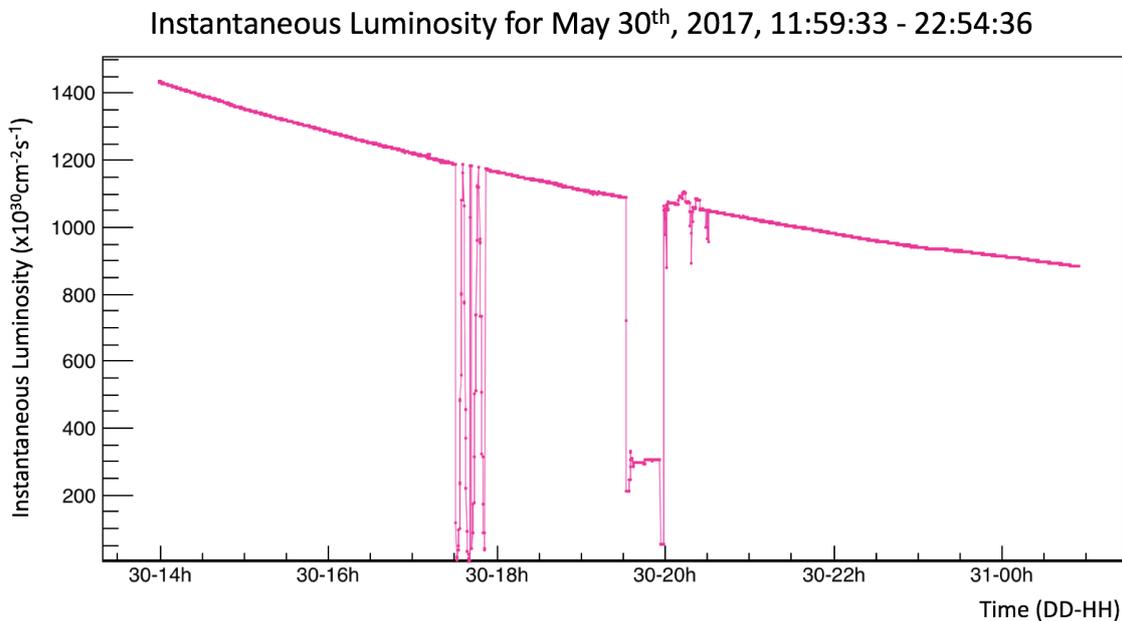


Figure 6.13 – LHC beam luminosity measured by the CMS detector from 2017.05.30.11.59 to 2017.05.30.22.54. The margins on either side of the data represent empty bins, not zero luminosity.

In this plot there are two key periods of time in which the beam luminosity appeared to not be stable. However, it is important to note that this is the luminosity as recorded by CMS, not necessarily the luminosity sent by the LHC. As such, they can potentially be explained by actions taken by the detector.

The first period of observed instability, from 2017.05.30.16.30 to 2017.05.30.16.57, occurred during a Van der Meer (VdM) scan performed by CMS. VdM scans are used to calibrate the luminosity measurement within CMS, and are performed by separating the two proton beams in the transverse plane and then moving them in steps across one another. The scans are taken in orthogonal directions, along the x- and y-axis of the CMS coordinate system. From these two scans, the area of proton density overlap between the two beams is measured as the product of the rate measurements' Gaussian widths as a function of the beam separation [95].

The second period of observed instability, from 2017.05.30.18.28 to 2017.05.30.19.35, occurred at the request of the PIXEL detector, which was taken out of global running in order to perform its own calibration scans via "lumi levelling", in which the two proton beams are displaced with respect to each other in order to adjust the luminosity values at the desired target, in this case PIXEL.

However, although the slice test detectors would have experienced this same drop-off in beam luminosity as the rest of CMS, it is unlikely that these issues in and of themselves would have caused the discharge. Although it might be suspected that a spike in luminosity, and thus in the rate of particles entering the GEM detectors, could potentially lead to the right conditions for an electrical discharge within the detector, a drop in luminosity would not be suspect as such. Even in the case of beam displacement, as in these two conditions, it is highly unlikely that the luminosity seen directly within the GEM detectors would have been significantly higher than during normal fills. And, in any case, the two periods of luminosity instability do not directly correspond with the time of the discharge event.

6.4.5 Magnetic Field

From Section 6.2.3, we can turn next to the magnetic field conditions present during this time. Unfortunately, in this particular case, they do not tell us much. Figure 6.14 is taken directly from CMSOnline and shows the entire period of potential channel loss, from 2017.05.10.13.55 to 2017.05.31.20.58.

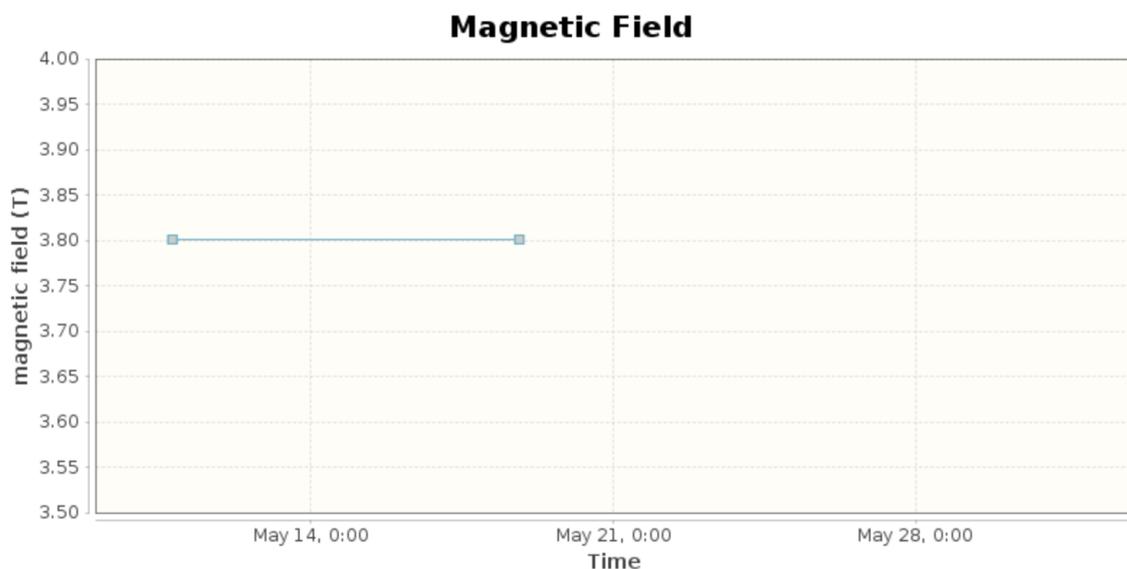


Figure 6.14 – CMS magnetic field from 2017.05.10.13.55 to 2017.05.31.20.58. The margins on either side of the data represent empty bins, not zero magnetic field.

Here, we see that during this three week period, only two data measurements were

made of the magnetic field, both of which were well before the time at which the discharge occurred. Although the CMS magnet can be considered to be quite reliably stable during operations, with the exception of planned ramp ups and ramp downs, and thus we could infer that the magnetic field was similarly at 3.8 T during the time of the discharge, we cannot conclusively draw any conclusions from this.

6.4.6 Environmental Factors

Lastly, we can attempt to discern information from the environmental conditions at the time of the discharge.

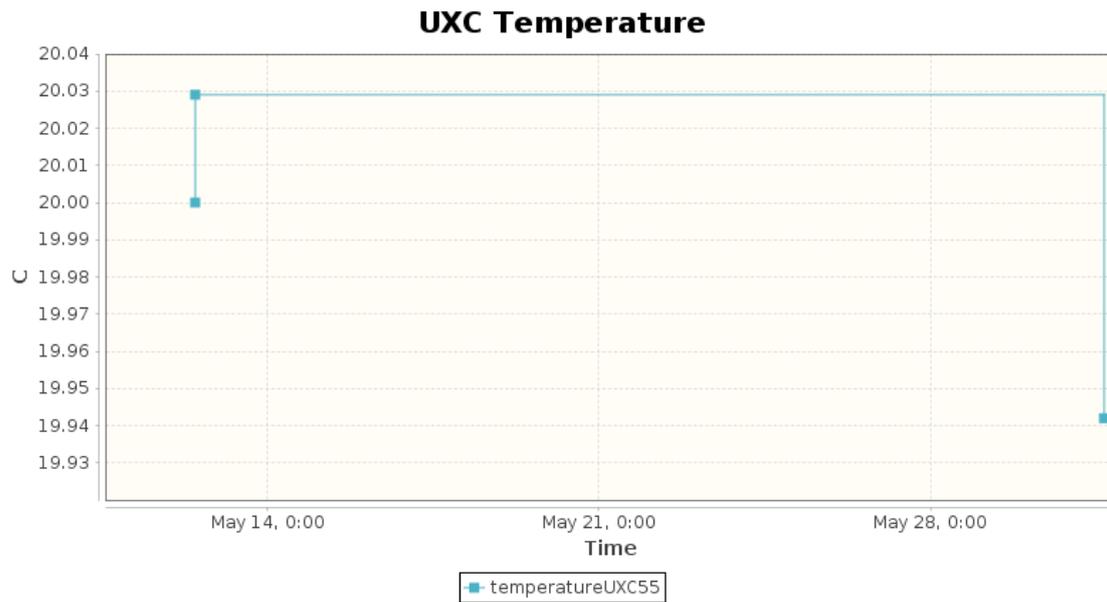


Figure 6.15 – Temperature in UXC55 2017.05.10.13.55 to 2017.05.31.20.58. The margins on either side of the data represent empty bins.

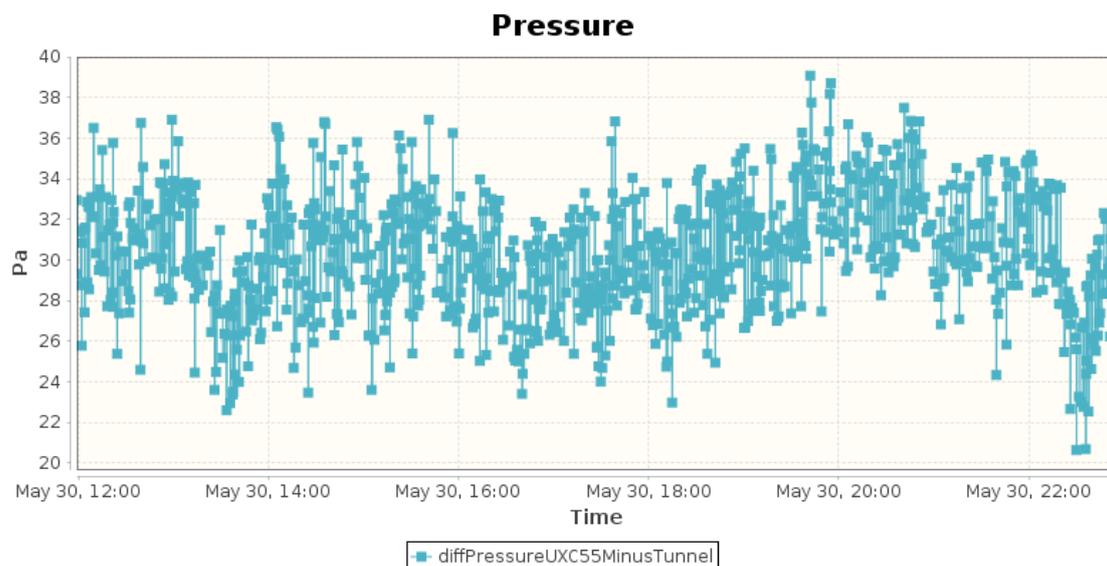


Figure 6.16 – Differential pressure in UXC55 2017.05.30.11.59 to 2017.05.30.22.55.

Figure 6.15 shows the temperature in UXC55 in degrees Celsius during the entire three

week window. As shown, there were only three temperature measurements taken during this period, none of which were at the time of the discharge. However, we can again assume some degree of stability, given the large underground nature of the UXC55 cavern, and infer that it was not a large change in temperature that caused the problem.

Figure 6.16 shows the differential pressure in UXC55 in Pa. Although here we have significantly more data points, and can thus narrow it down to just the period in which the beam was on, the changes in pressure are confined to such a small range that we still cannot draw any significant conclusions from this.

6.4.7 Conclusions

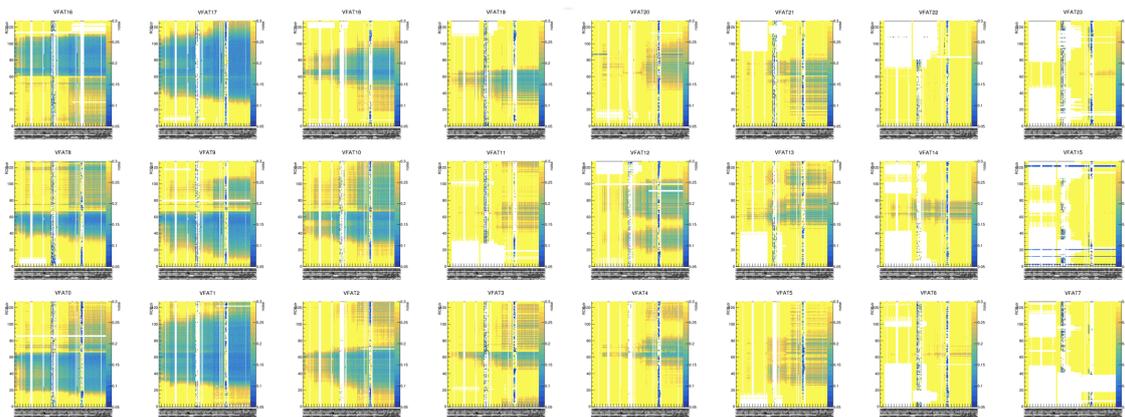
Although we can be relatively certain of the discharge which caused the damage to VFAT23 in GEMINIm01L2, the lack of a clear environmental cause and the sample size of $n=1$ leaves us wanting. Thankfully, there are still eight other slice test detectors, and while the remaining eight do not have the electrode-by-electrode monitoring afforded to GEMINIm01 by the A1515TG power supply, and were not removed for examination midway through the slice test, there is still a great deal of information to be gleaned from them.

6.5 Channel Loss in the HV Divider Chambers

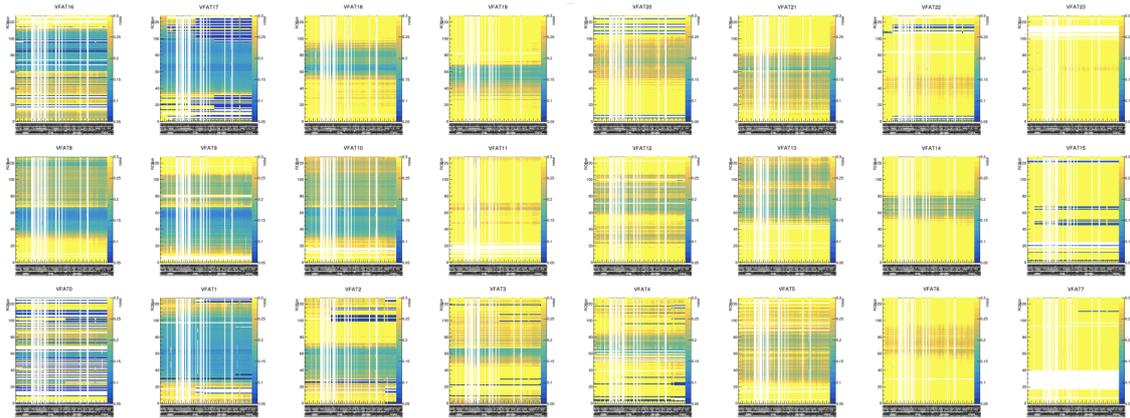
The remaining eight slice test detectors - GEMINIm27, GEMINIm28, GEMINIm29, and GEMINIm30 - first began operation on March 27th, 2017 and remained in operation within the CMS endcap until the end of the slice test on December 7th, 2018. During this time period, S-Curves were taken of all of the chambers frequently, with daily S-curves beginning mid-2017 in order to better track the loss of readout channels.

6.5.1 Time Series S-curves

The time series S-curve plots for all eight HV divider chambers are presented in Figures 6.17 through 6.24 for an overall look at the behavior of the detectors. Note that although the last S-curve was taken for the slice test detectors on 2018.12.07.16.41, these plots end on 2018.10.23.13.23, as no channel loss occurred after this date.



(a) Time series S-curves for GEMINIm27L1 from 2017.03.27.13.51 to 2018.06.28.11.53



(b) Time series S-curves for GEMINIm27L1 from 2018.07.01 to 2018.10.23.13.23

Figure 6.17 – Time series S-curves for GEMINIm27L1 from 2017.03.27.13.51 to 2018.10.23.13.23.13.23, split into two date ranges as a result of a coding error

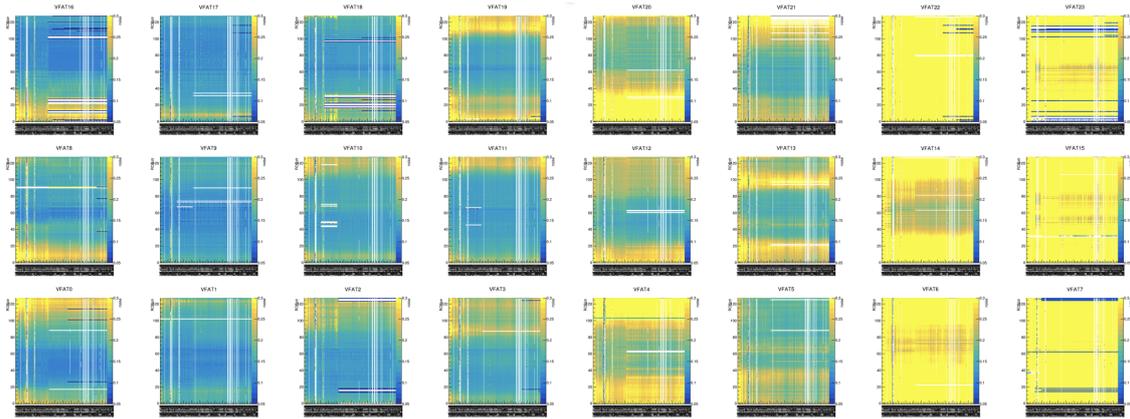


Figure 6.18 – Time series S-curves for GEMINIm27L2 from 2017.03.27.13.51 to 2018.10.23.13.23

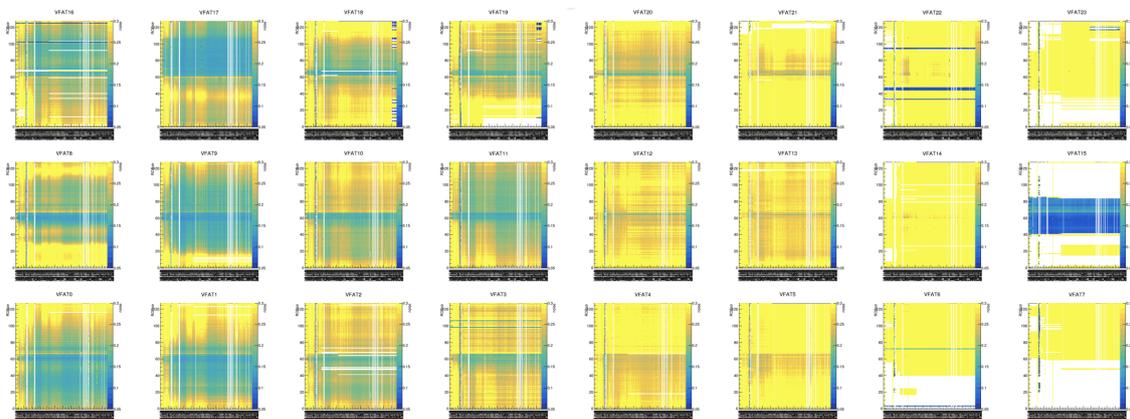


Figure 6.19 – Time series S-curves for GEMINIm28L1 from 2017.03.27.13.51 to 2018.10.23.13.23

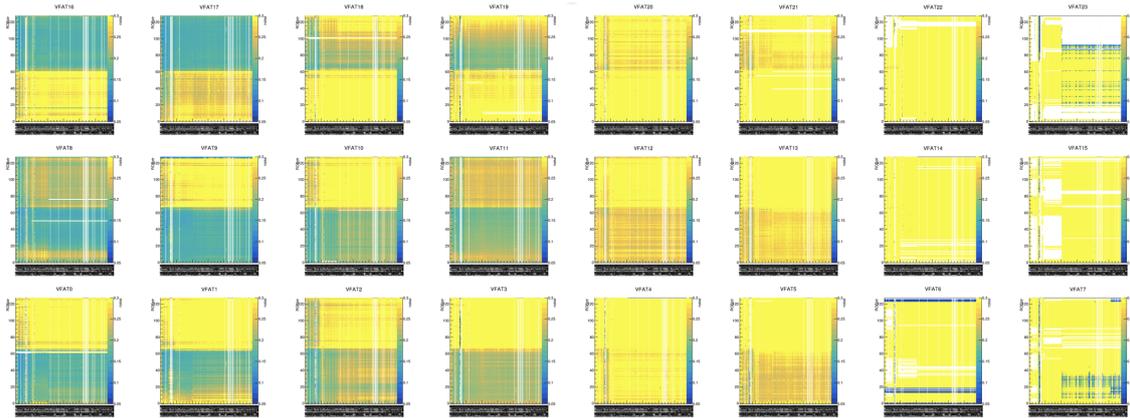


Figure 6.20 – Time series S-curves for GEMINIm28L2 from 2017.03.27.13.51 to 2018.10.23.13.23

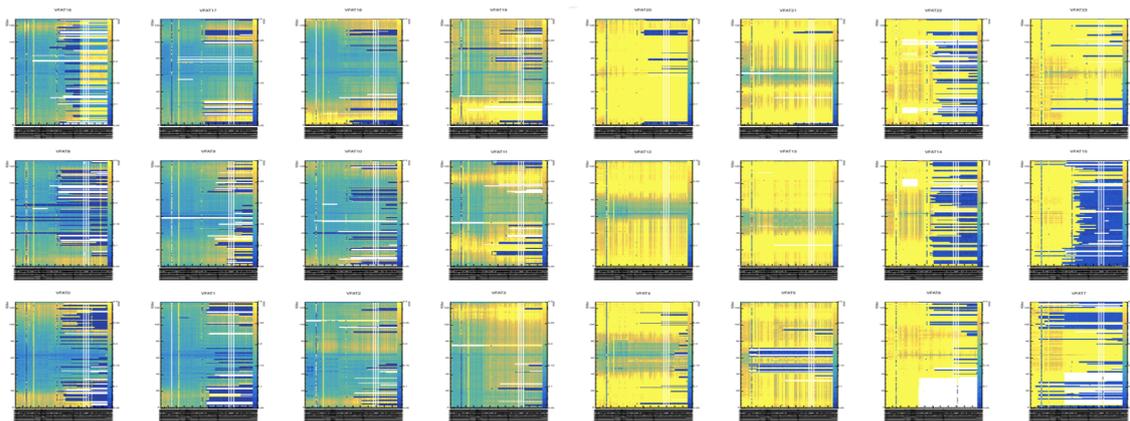


Figure 6.21 – Time series S-curves for GEMINIm29L1 from 2017.03.27.13.51 to 2018.10.23.13.23

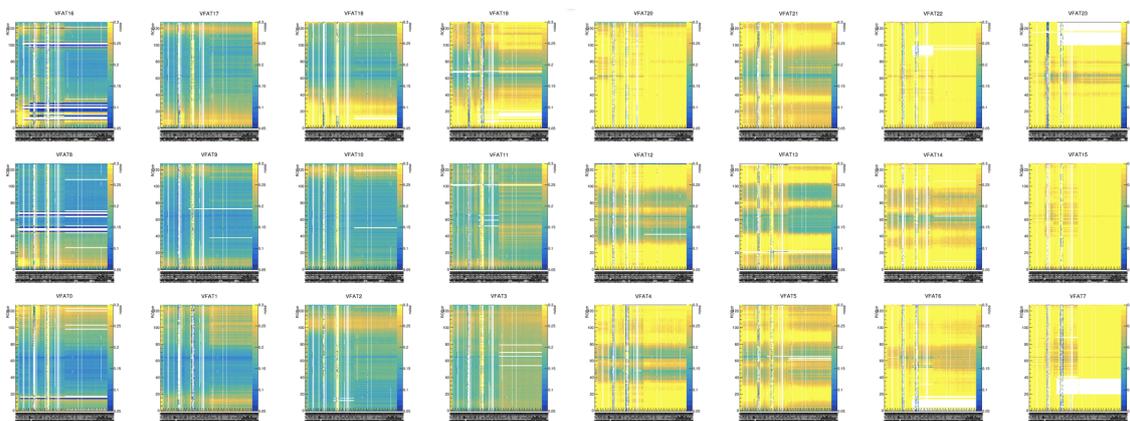


Figure 6.22 – Time series S-curves for GEMINIm29L2 from 2017.03.27.13.51 to 2018.10.23.13.23

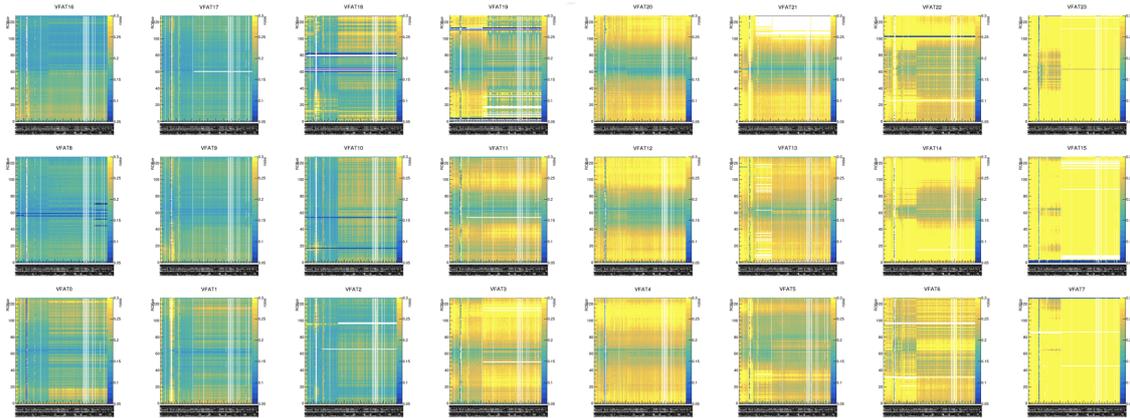


Figure 6.23 – Time series S-curves for GEMINIm30L1 from 2017.03.27.13.51 to 2018.10.23.13.23

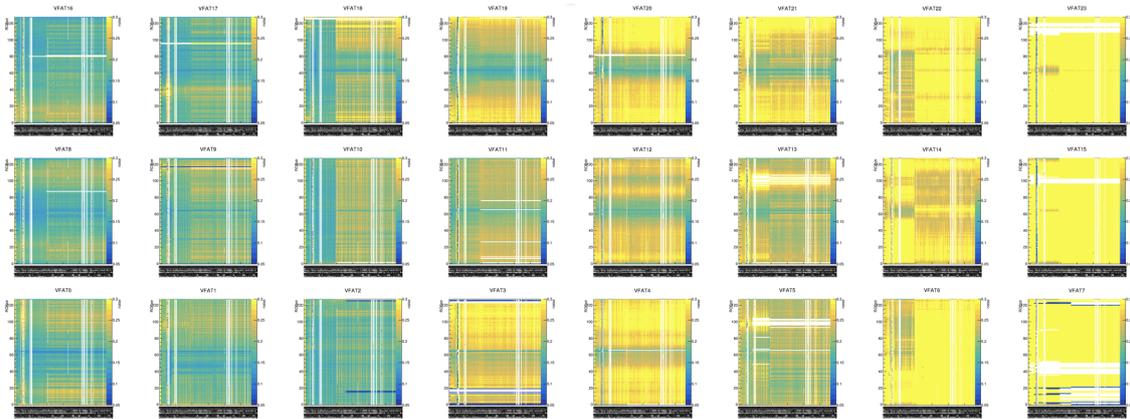


Figure 6.24 – Time series S-curves for GEMINIm30L2 from 2017.03.27.13.51 to 2018.10.23.13.23

6.5.2 Summary Tables

From these S-curves, we were able to track exactly how many readout channels were lost throughout the slice test, in near real-time. Tables 6.2 through 6.9 present a summary of these results. Each table is broken into four columns, which represent the number of channels which were damaged, the percentage of total channels that that number represents, the number of VFATs which those channels are present on, and the distinct scandates which the damage occurred on. This last one is present because, for several detectors, a single scan date incurred dead channels on multiple VFATs, suggesting a detector-wide event and warranting further investigation. These events are then broken into three row categories: channels which were dead from the first scandate of the slice test, channels which burned during slice test operations, and channels which recovered in some way. Recovery of a channel may involve a channel which was dead from the beginning becoming operational, or a channel which died during operation and then recovered some time later, and may include channels which subsequently died again.

Channel Status	Distinct Dates Occurred	Number of VFATs Affected	Number of Channels Affected	Percentage of Total Channels
Recovered	0	0	0	0%
Dead from the start	N/A	1	5	0.16%
Burned during operation	19	11	95	3.09%

Table 6.2 – Summary table of readout channel loss for GEMINIm27L1 from 2017.03.27.13.51 through 2018.12.07.16.41

Channel Status	Distinct Dates Occurred	Number of VFATs Affected	Number of Channels Affected	Percentage of Total Channels
Recovered	2	2	2	0.06%
Dead from the start	N/A	1	1	0.03%
Burned during operation	19	10	66	2.15%

Table 6.3 – Summary table of readout channel loss for GEMINIm27L2 from 2017.03.27.13.51 through 2018.12.07.16.41

Channel Status	Distinct Dates Occurred	Number of VFATs Affected	Number of Channels Affected	Percentage of Total Channels
Recovered	0	0	0	0%
Dead from the start	N/A	4	42	1.37%
Burned during operation	2	3	32	1.04%

Table 6.4 – Summary table of readout channel loss for GEMINIm28L1 from 2017.03.27.13.51 through 2018.12.07.16.41

Channel Status	Distinct Dates Occurred	Number of VFATs Affected	Number of Channels Affected	Percentage of Total Channels
Recovered	0	0	0	0%
Dead from the start	N/A	1	9	0.29%
Burned during operation	1	1	5	0.16%

Table 6.5 – Summary table of readout channel loss for GEMINIm28L2 from 2017.03.27.13.51 through 2018.12.07.16.41

Channel Status	Distinct Dates Occurred	Number of VFATs Affected	Number of Channels Affected	Percentage of Total Channels
Recovered	1	1	1	0.03%
Dead from the start	N/A	1	8	0.26%
Burned during operation	110	22	865	28.16%

Table 6.6 – Summary table of readout channel loss for GEMINIm29L1 from 2017.03.27.13.51 through 2018.12.07.16.41

Channel Status	Distinct Dates Occurred	Number of VFATs Affected	Number of Channels Affected	Percentage of Total Channels
Recovered	0	0	0	0%
Dead from the start	N/A	1	1	0.03%
Burned during operation	9	3	33	1.07%

Table 6.7 – Summary table of readout channel loss for GEMINIm29L2 from 2017.03.27.13.51 through 2018.12.07.16.41

Channel Status	Distinct Dates Occurred	Number of VFATs Affected	Number of Channels Affected	Percentage of Total Channels
Recovered	1	1	1	0.03%
Dead from the start	N/A	6	19	0.62%
Burned during operation	1	1	5	0.16%

Table 6.8 – Summary table of readout channel loss for GEMINIm30L1 from 2017.03.27.13.51 through 2018.12.07.16.41

Channel Status	Distinct Dates Occurred	Number of VFATs Affected	Number of Channels Affected	Percentage of Total Channels
Recovered	0	0	0	0%
Dead from the start	N/A	2	8	0.26%
Burned during operation	4	2	15	0.49%

Table 6.9 – Summary table of readout channel loss for GEMINIm30L2 from 2017.03.27.13.51 through 2018.12.07.16.41

6.5.3 Summary Plot

All of this can be summarized in a single plot which has been presented, published, and shared far and wide over the past two years, and has finally found its rightful place in the dissertation of the woman who put her heart and soul into making it:

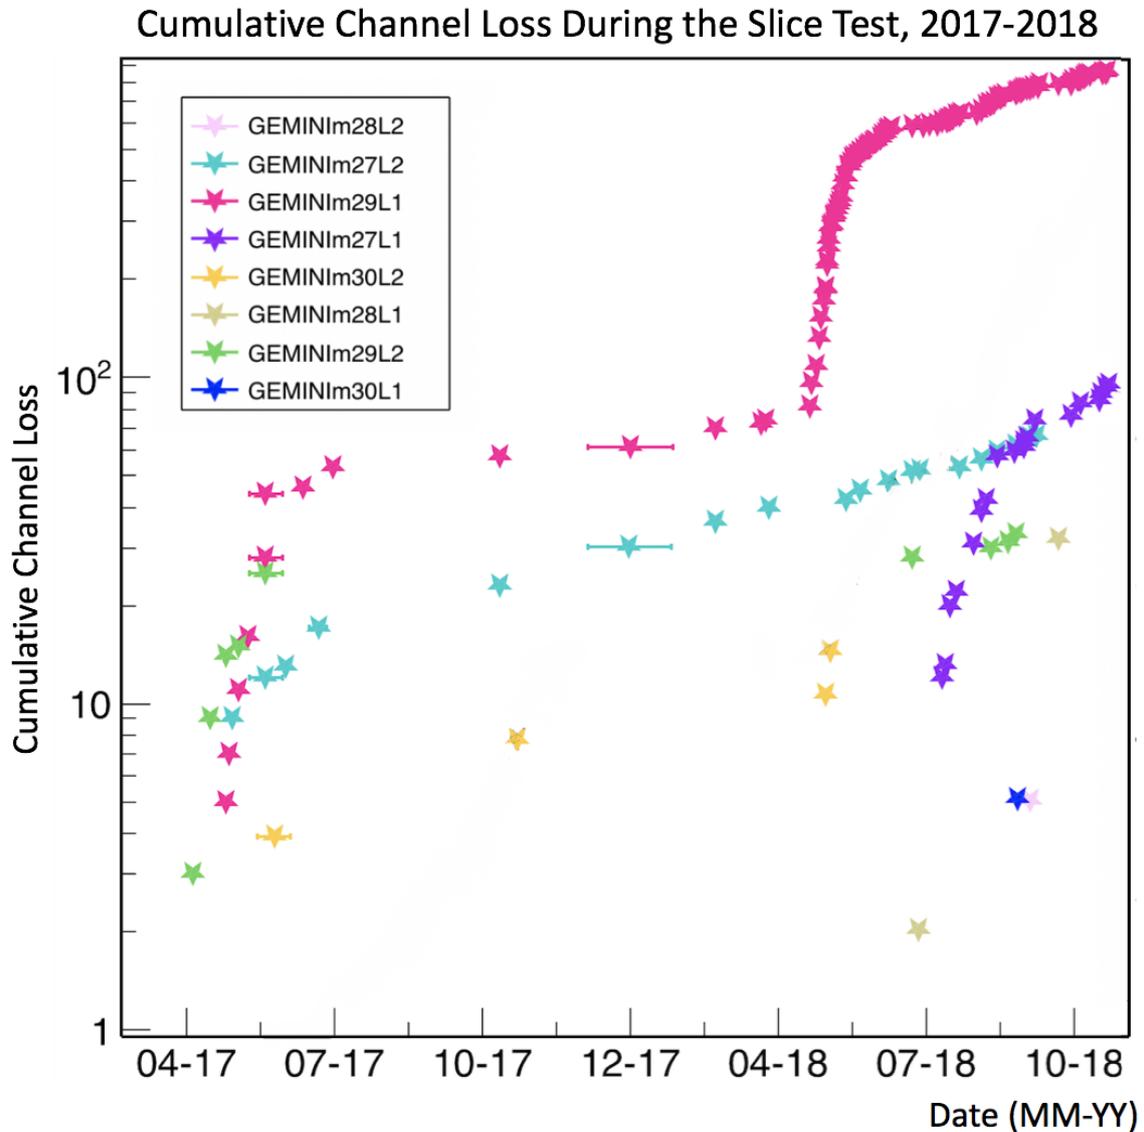


Figure 6.25 – Cumulative channel loss for the duration of the slice test

Immediately, we can begin to discern a pattern in the losses. The three chambers which experienced the most loss - GEMINIm27L1, GEMINIm27L2, and GEMINIm29L1 - are all short chambers. This means that they all share the same geometry, and also share the same orientation within the endcap, which is opposite that of the long chambers. This would warrant an exploration into whether the geometry or orientation has some effect on the discharge rate.

Furthermore, the channel loss seems to be grouped around certain periods of time, with the first major set of channel loss occurring from April to July 2017, the second occurring from April to June 2018, and the third beginning in July 2018 and continuing through the end of October 2018. This would point us to some sort of environmental factor, such as the presence of beam or changing magnetic fields.

6.5.4 High Voltage Instabilities as Discharge Indicators

What makes the HV divider chambers so difficult to pinpoint an exact cause of damage is not only that they lack the electrode-by-electrode monitoring that the GEMINIm01 superchamber had, but also that the HV divider masks the presence of the discharges to a certain degree. And because the discharges are so transient, particularly when compared to the two-second sampling time of the database monitoring software, they often do not trigger an overcurrent or overvoltage in the monitoring software, and thus do not cause the detector to trip.

Let’s take GEMINIm28L1 for example, since it has only two periods of channel loss which are bounded by fairly small gaps between S-curve scans: two channels were lost between 2018.06.27.12.40 and 2018.06.28.11.34, and 30 channels were lost between 2018.09.21.19.27 and 2018.09.22.14.27. During these time periods, no status change such as an overvoltage, overcurrent, or trip was recorded, which could have alerted the DOC to a potential problem. And yet, when we look at the HV values recorded during this time period, we see a very different story.

Figure 6.26 shows the HV current of the GEMINIm28L1 detector between 2018.09.21.19.27 and 2018.09.22.14.27, while Figure 6.27 shows the HV voltage. Each data point represents a new value in either current or voltage, with a minimum of two seconds between each point. If the values do not change, no data point is recorded in the database.

In both plots, we see several periods of ramping up and ramping down over the time period shown. This is the result of the different recipes in the DCS for different modes of operation. The higher values represent the detector in “physics” mode, whereas the lower values represent the detector in “standby” mode. For the most part, when not ramping up or down, the HV current and voltage of the detector is stable, although the HV is more prone to slight fluctuations which result in more data points being recorded at the different plateaus.

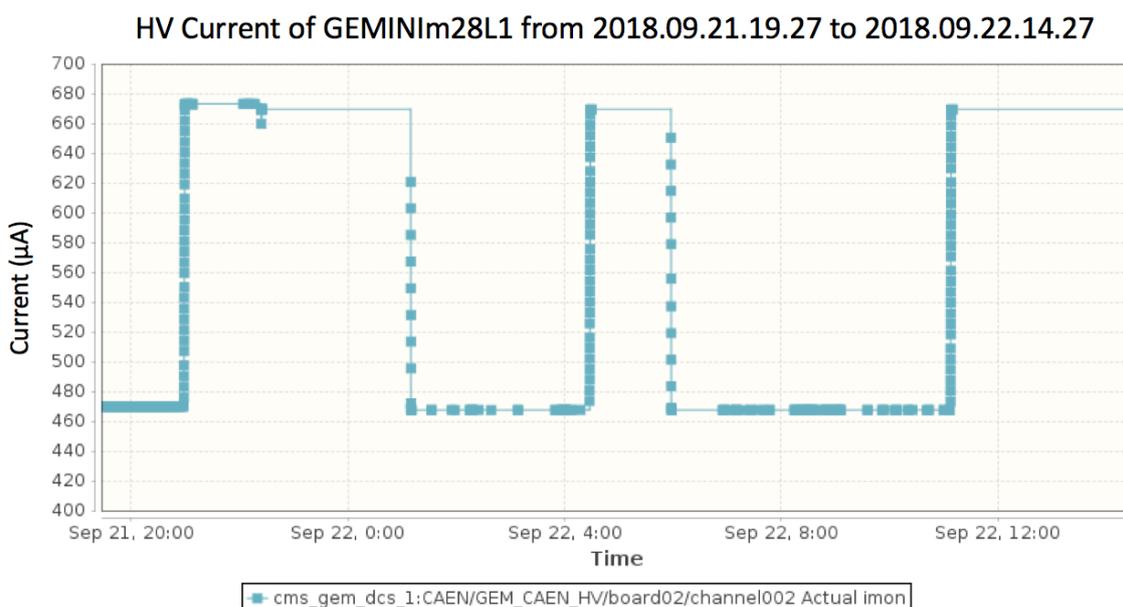


Figure 6.26 – HV current of the GEMINIm28L1 detector between 2018.09.21.19.27 and 2018.09.22.14.27

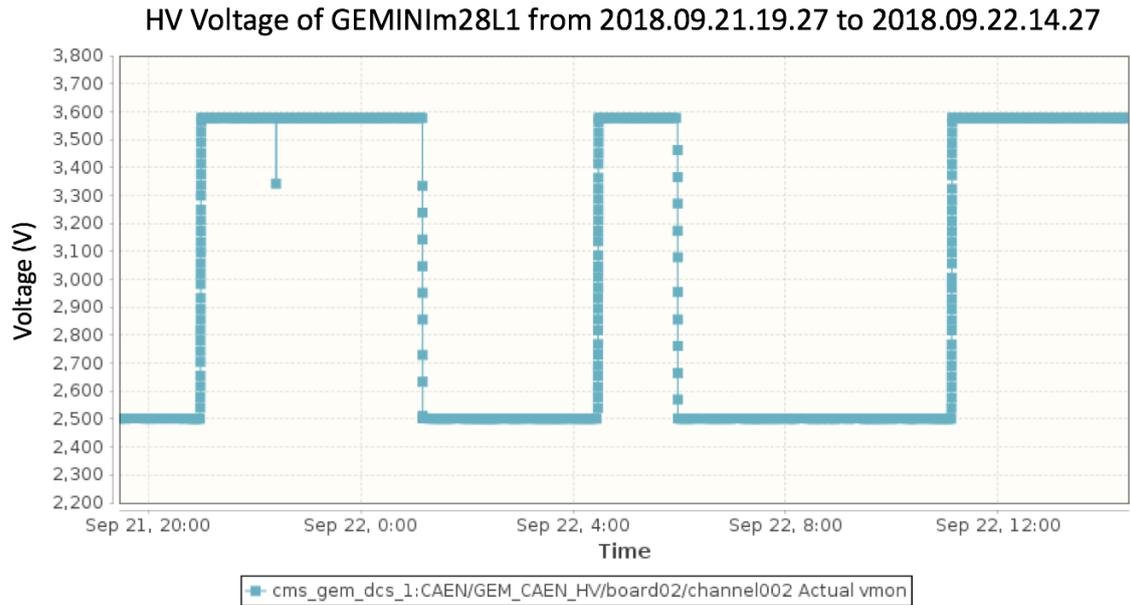


Figure 6.27 – HV voltage of the GEMINIm28L1 detector between 2018.09.21.19.27 and 2018.09.22.14.27

However, on September 21st at 22:24:13, we see a sudden drop in voltage from 3.577 kV to 3.342 kV, which corrects back up to 3.577 kV by the next reading two seconds later. At the same time period, the current suddenly drops from a stable 672.9 μA to 659.6 μA before correcting up to 668.7 μA by the next two second reading. Although the detector did not trip, this is strongly indicative of a discharge occurring within the detector.

Looking now to the other period of channel loss, Figure 6.28 shows both the HV current and voltage in the same plot.

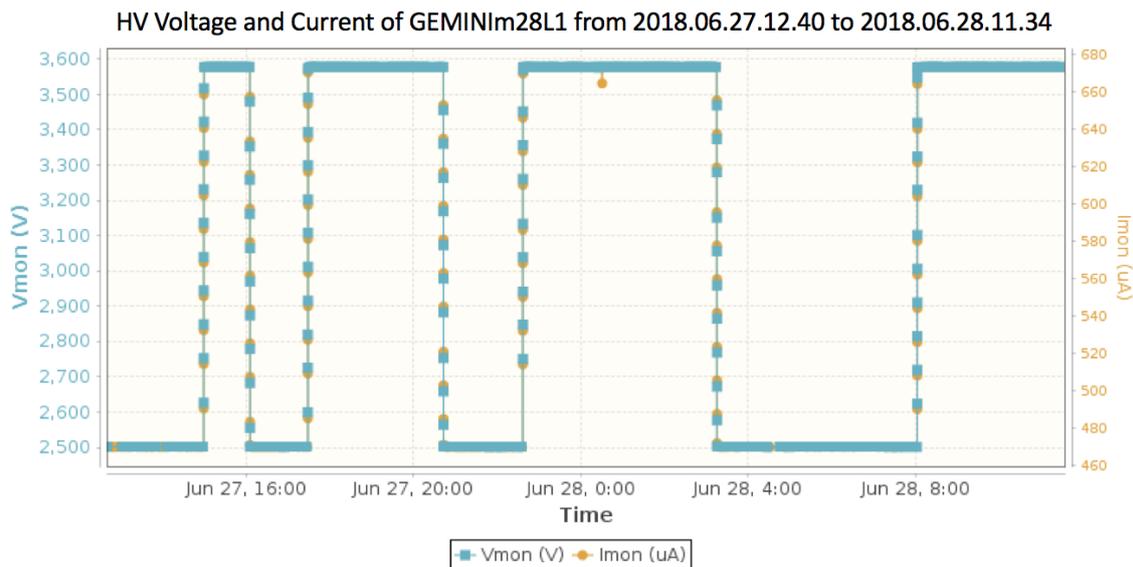


Figure 6.28 – HV voltage and current of the GEMINIm28L1 detector between 2018.06.27.12.40 and 2018.06.28.11.34

Once again, we see several periods of ramping up and ramping down over the time period shown, with general stability in between. And, once again, we see a spike which indicates a discharge took place: on June 28th at 00:30:55, the current suddenly drops from

a stable $672.4 \mu\text{A}$ to $664.2 \mu\text{A}$ before correcting up to $672.7 \mu\text{A}$ by the next two second reading. And, once again, the detector did not trip in response, which would have alerted the DOC to a problem.

Unlike in the previous example, however, the voltage does not record any deviations at the same time as this current spike, and this is what makes uncovering these discharges so difficult. If the discharge was small enough (as this one is certainly smaller than the September one, with a maximum recorded difference of $8.2 \mu\text{A}$ compared to $13.3 \mu\text{A}$), the two-second sampling window might result in the voltage difference not being recorded by the monitoring software (or, vice versa, with the voltage spike being recorded but not the current spike).

Interestingly, while the current of the detector is normally quite stable at the “physics” operating current, in the minutes prior to the September 21st discharge event, we see several fluctuations of less than $0.5 \mu\text{A}$, as shown in Figure 6.29. While these cannot be definitively classified as discharges themselves, they do represent an instability in the detector that could have contributed to the channel loss. This preceding instability was not seen in the June 28th discharge event, and may provide a clue as to why the September event was a particularly damaging one, with the loss of 30 readout channels, compared to just the two channels lost in the June event.

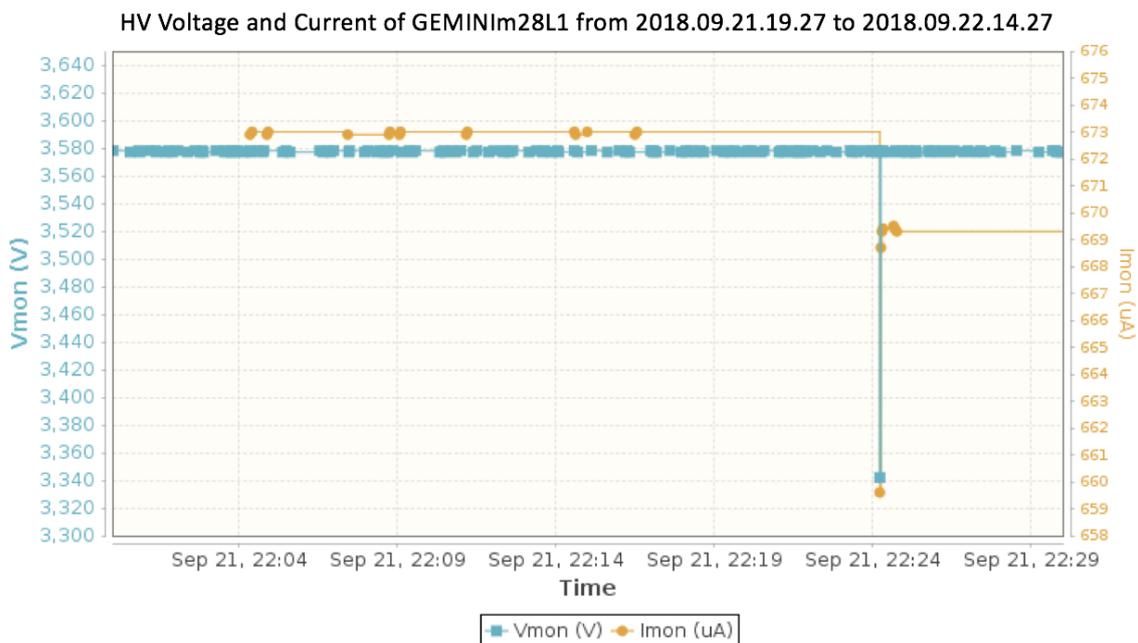


Figure 6.29 – HV voltage and current of the GEMINIm28L1 detector between 2018.09.22.22.00 and 2018.09.22.22.30

6.5.5 Beam Luminosity

As before, we can look to the luminosity of the beam present in CMS during these periods of channel loss to see if there is any distinct correlation between the two.

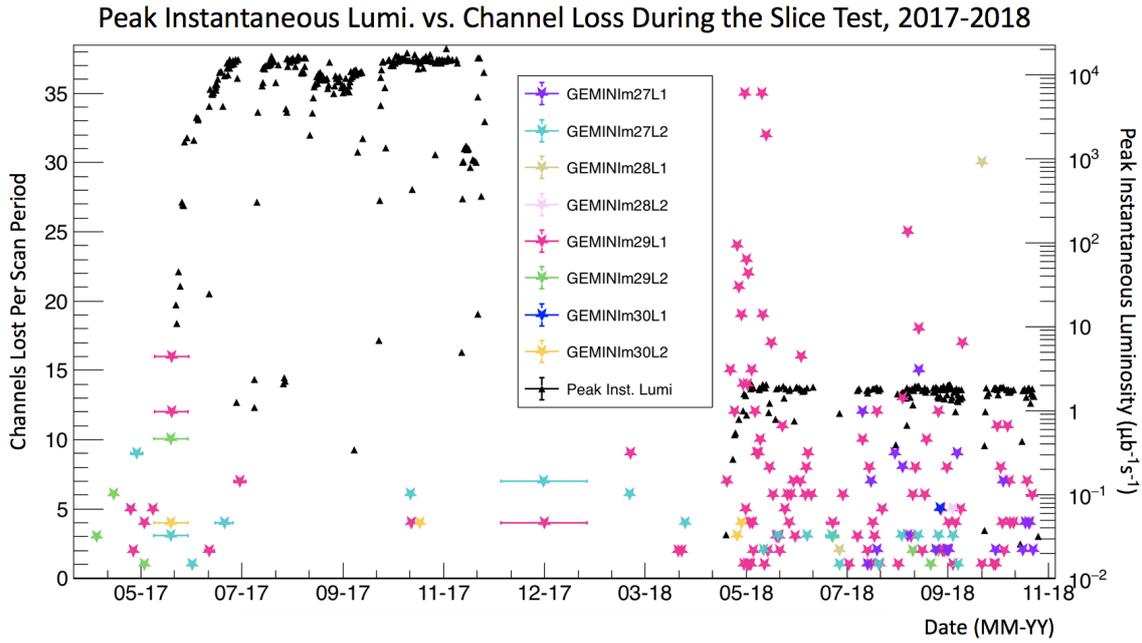


Figure 6.30 – Peak instantaneous luminosity vs. channel loss for the duration of the slice test

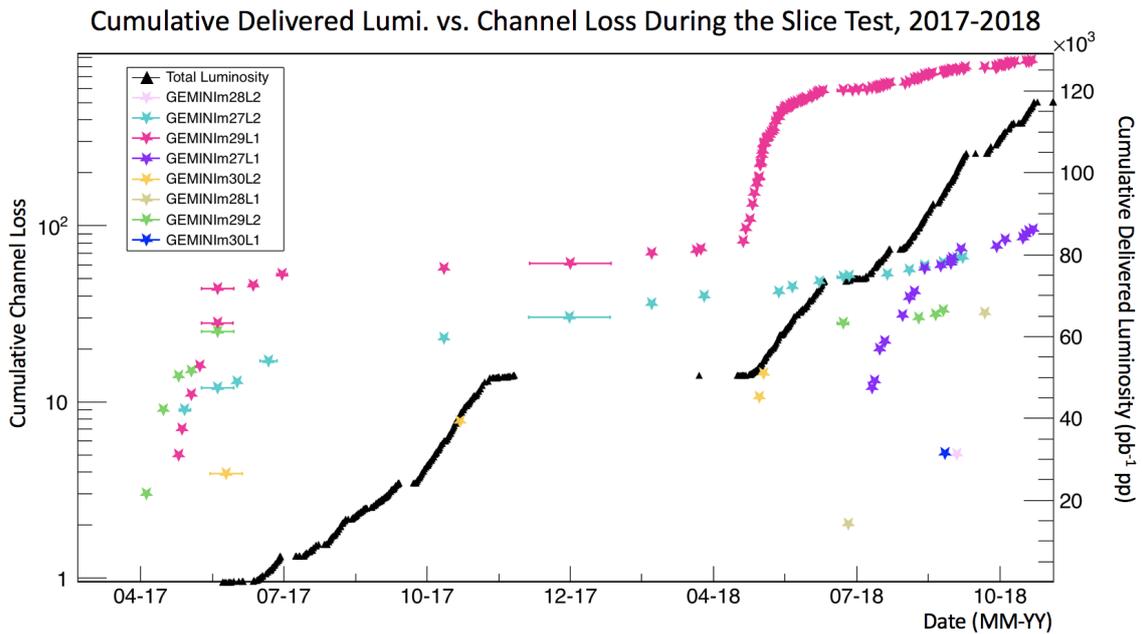


Figure 6.31 – Cumulative delivered luminosity vs. channel loss for the duration of the slice test

Figure 6.30 shows the peak instantaneous luminosity of each run during which a readout channel was lost. The readout channels lost are presented non-cumulatively for each detector, such that the data points represent the number of channels lost per individual scan date. The x-axis is given as the date in month-year format, and spans the entire duration of channel loss, from April 2017 through October 2018.

Figure 6.31, on the other hand, shows the cumulative channel loss per detector against the cumulative (integrated) luminosity seen by the detectors.

In both cases, we find somewhat unexpected results. For both GEMINIm27L2 and

GEMINIm29L1, multiple channels were lost (seven for GEMINIm27L2, four for GEMINIm29L1) during the period of 2017.12.05.21.50 to 2018.01.26.15.56, more than half of which was during the Christmas shutdown and therefore was not exposed to beam, magnetic field, or workers. Several more channels were lost after the 2018 re-opening, still prior to any beam. It can be assumed, then, that these channels were lost as a result of mechanical interventions around the chambers, which may have caused debris to become dislodged (see Sections 6.5.6.1 and 6.5.7).

Another unexpected result is the lack of direct correlation between the presence of beam and channel loss. While the majority of the total channel loss occurred in 2018 when there was beam, the greater part of the channel loss which occurred in 2017 was when there was no beam present.

Nor is there a direct correlation between the luminosity of the beam and the channels lost, as some detectors (such as GEMINIm29L1) lost significantly more channels under the same beam conditions as their neighboring chambers or even the sister chamber of the same superchamber (e.g. GEMINIm29L2). And, indeed, the same chamber might lose > 30 channels one day and < 5 channels the next day despite nearly-identical beam conditions.

Clearly, then, it is not as simple as inferring that a greater concentration of highly-ionizing particles is wholly responsible for the discharges within the detectors and therefore the damage to the readout electronics, and we must continue investigating other potential causes.

6.5.6 Magnetic Field

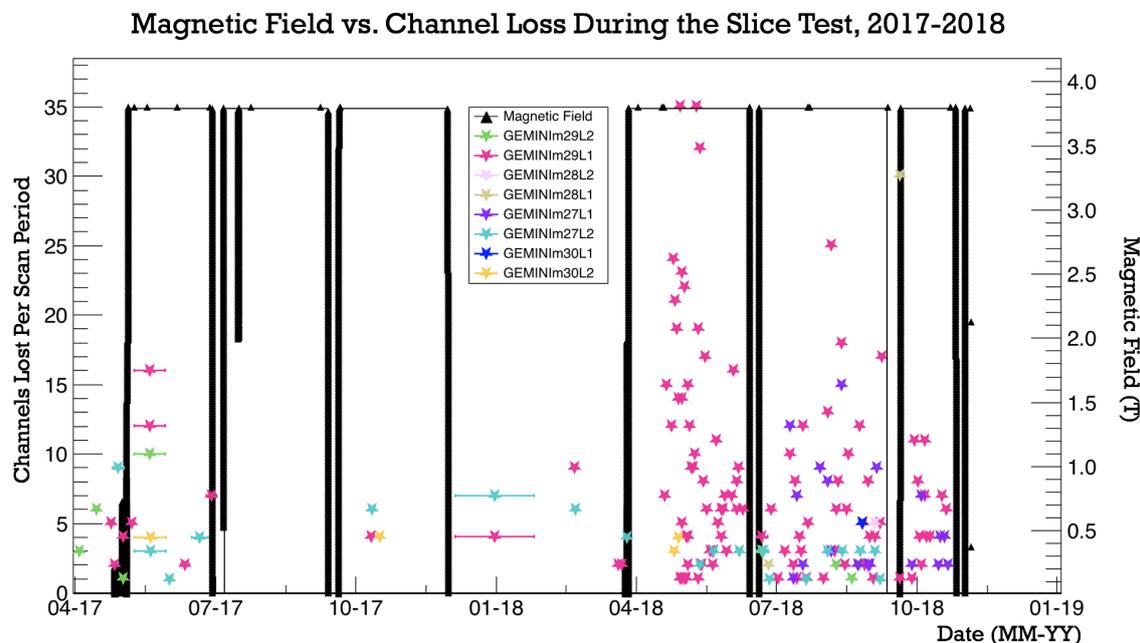


Figure 6.32 – Magnetic field vs. channel loss for the duration of the slice test

Figure 6.32 shows the non-cumulative channel loss for all slice test detectors in comparison to the magnetic field present within the CMS detector.

As is normal for CMS operation, the solenoid magnetic field was kept at a fairly steady 3.8 T, and only ramped down for the year-end technical stop and a handful of planned technical stops during the year. Thus, because the magnetic field is, for the most part, the

same, statistically we expect to see the majority of channel loss during periods of stable magnetic field. And, indeed, this is what we see.

However, this does not mean that the magnetic field data is of no use to us. Rather, we can examine the few periods of unstable magnetic field to see if there is any behavior within the detectors that is particular to these conditions.

6.5.6.1 Channel Recovery

Figure 6.33 shows the time series S-curve for VFAT7 of GEMINIm27L2.

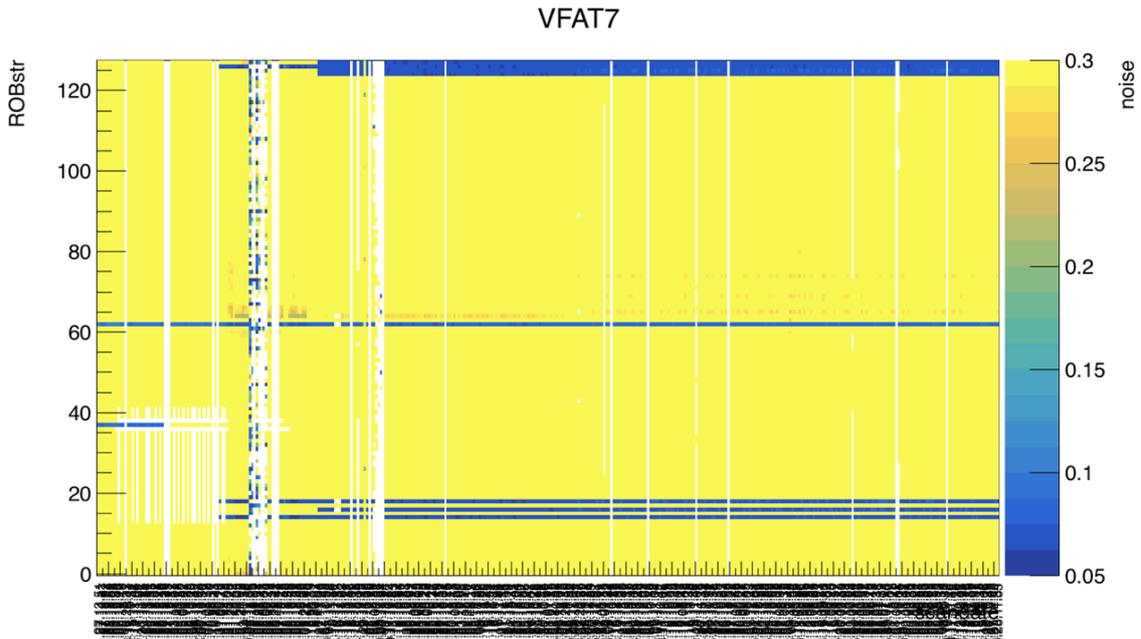


Figure 6.33 – Time series S-curve for VFAT7 of GEMINIm27L2, with readout board strip #37 exhibiting channel recovery

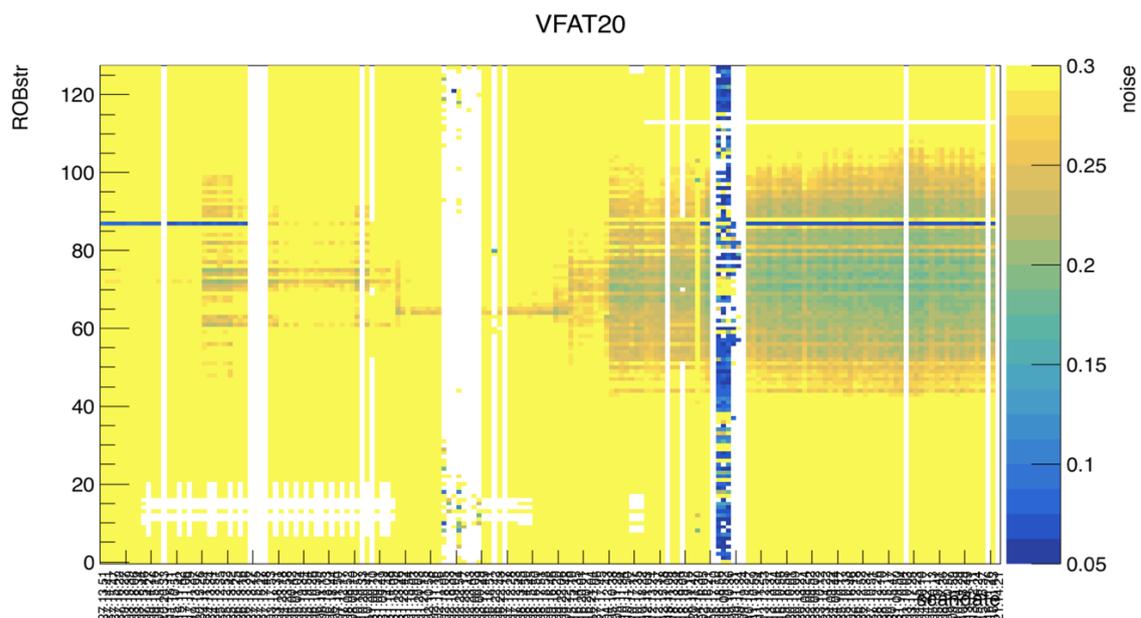


Figure 6.34 – Time series S-curve for VFAT20 of GEMINIm27L1, with readout board strip #87 exhibiting channel recovery

Here, we see unusual behavior in one of the readout channels. Although this channel fit the noise characteristics of a channel which was dead from the point of installation, it recovered to normal operating parameters between 2017.04.26.13.36 and 2017.05.03.17.53.

Figure 6.34, which shows the time series S-curve for VFAT20 of GEMINIm27L1, also has a channel which was dead from the point of installation but recovered to normal operating parameters between 2017.04.26.13.36 and 2017.05.03.17.53. In this case, however, the channel returned to its dead state some time afterwards, between 2017.07.20.11.47 and 2017.10.02.16.05.

Figure 6.35 shows the magnetic field present at the time of the two channels' recovery.

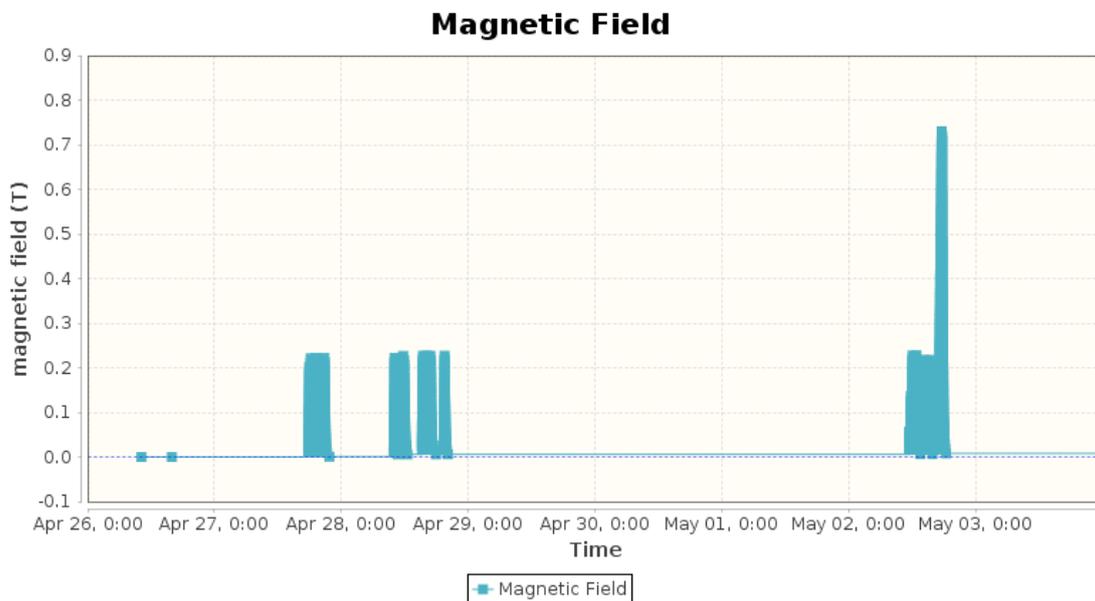


Figure 6.35 – CMS magnetic field from 2017.04.26.13.36 to 2017.05.03.17.53. The margins on either side of the data represent empty bins, not zero magnetic field.

As shown, this time period was in fact one of the rare periods of magnetic field instability, being the first time the CMS magnet was ramped up to any degree after Long Shutdown I. This may imply that there was some physical debris within the detectors that was blocking the input, making it appear as though the channel was dead when in reality the electrical signals were simply blocked. If the ramping of the magnet caused the disk of the endcap to flex, it could provide enough movement to the GEM detectors to dislodge the debris. Alternatively, if the debris was itself magnetic, this ramping could affect it enough to move it from its blocking position, though this is the less likely explanation due to the lack of magnetic materials used in the construction of the GEM detectors.

To strengthen this hypothesis, we can look at the magnetic field during the period of time in which readout board strip #87 of GEMINIm27L1's VFAT20 became dead again.

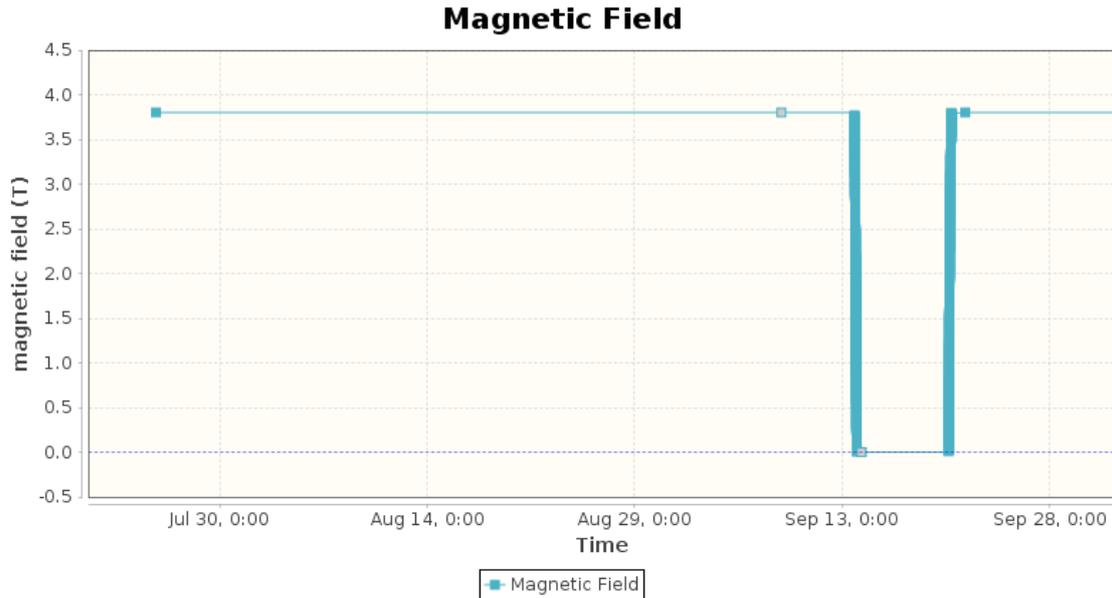


Figure 6.36 – CMS magnetic field from 2017.07.20.11.47 to 2017.10.02.16.05. The margins on either side of the data represent empty bins, not zero magnetic field.

Figure 6.36 shows the magnetic field present from 2017.07.20.11.47 to 2017.10.02.16.05. Once again, we see that during this time period, the solenoid magnet was ramped down and back up. If, then, the initial recovery of this channel was caused by the dislodging of a physical piece of debris which was blocking the VFAT input, but it was not dislodged by a significant distance or otherwise made inert, a further ramping of the magnetic field and subsequent flexing of the endcap disk could cause the debris to shift a second time and block the channel again.

Although the long period between the last bad S-curves and the first good S-curves for these channels (or, in the case of channel death after recovery, the last good S-curves and the first bad S-curves) makes it difficult to definitively say that it was the change of magnetic field which caused the recovery, nevertheless it provides some insight into how the magnetic field can affect the health of the GEM detector electronics within CMS.

6.5.7 Quality Control (and a Lack Thereof)

For approximately two months after the very start of operations (Lasting from April 3rd, 2017 through May 31st), both GEMINIm29L1 and GEMINIm29L2 experienced steady channel loss. However, it was only during the last week of this period, approximately May 23rd through May 31st, that there was any beam actually present within the LHC. The luminosity of this period is shown in Figure 6.37 for reference, taking note that this period encompasses the channel loss in GEMINIm01L2.

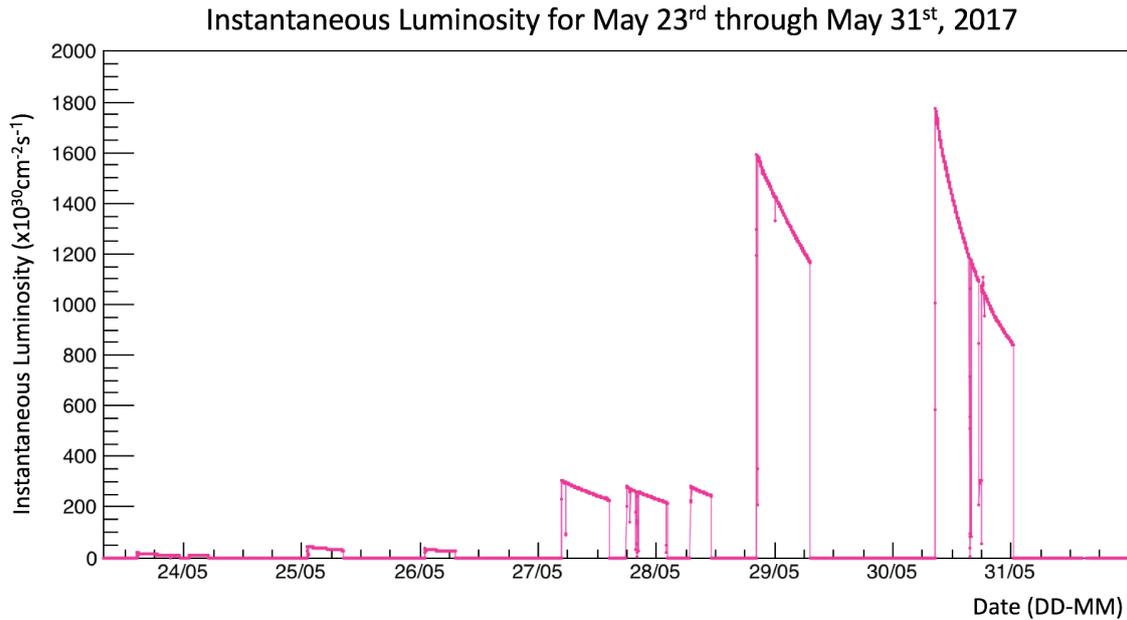


Figure 6.37 – LHC beam luminosity measured by the CMS detector from 2017.05.23 through 2017.05.31. The margins on either side of the data represent empty bins, not zero luminosity.

Since the channel loss began a significant amount of time before the introduction of beam, and since the introduction of beam did not correspond with a significant increase in the rate of channel loss, we can rule out the presence of beam as the cause of the channel loss.

Prior to this, but still during this two-month period, the CMS solenoid magnet was ramped up for the first time that operational year, as shown in Figure 6.38. However, the date of ramping does not correspond directly with a period of channel loss, or, taking into consideration that daily S-curves were not yet the norm at this point, a notable increase in the rate of channel loss.

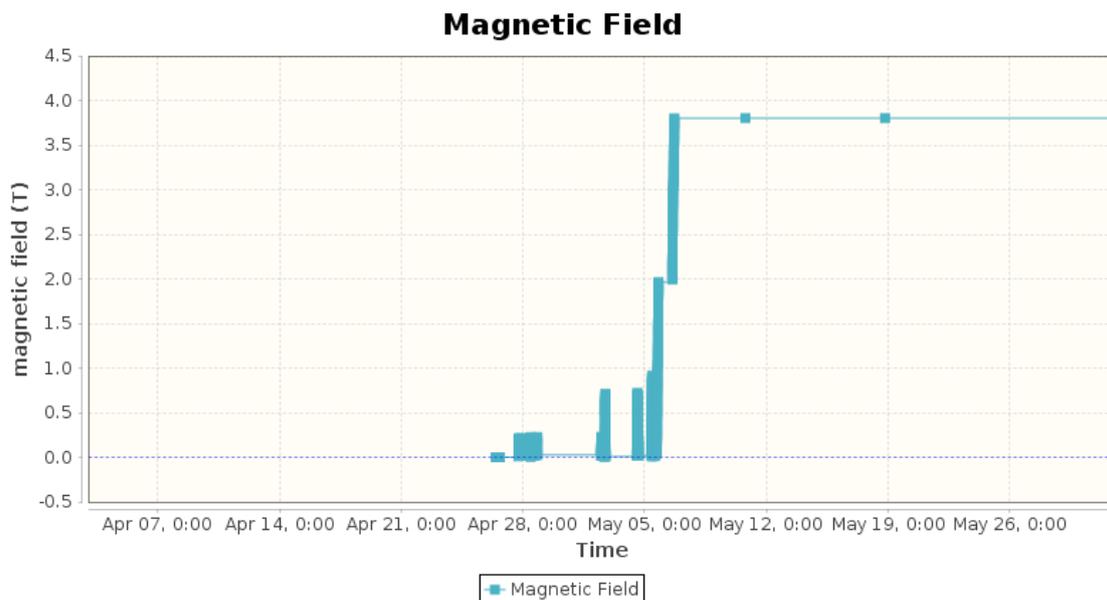


Figure 6.38 – CMS magnetic field from 2017.04.03 through 2017.05.31. The margins on either side of the data represent empty bins, not zero magnetic field.

Thus we must look to other reasons for this initial channel loss. Given that the channel loss began right at the very start of operations, we can consider the initial physical condition of the detectors.

For GE1/1, the sixth quality control step, or QC6, was implemented for each detector as part of the validation procedure. The aim of QC6 is to evaluate the high voltage stability of each detector and clean the GEM foils of any impurities that could promote discharges in the operation of the chamber, prior to the installation of any of the front-end readout electronics. It accomplishes this task in several steps. Initially, a Megger MIT430/2 insulation and continuity tester is used to apply 500 V to the tops of each foil, $G1_{top}$, $G2_{top}$, and $G3_{top}$, and the number of discharges within a minute is measured. Each discharge represents a piece of dust or other debris being burnt off of the foil, thereby cleaning it. In particularly “dirty” cases, up to 1 kV can be applied for a very short time in order to burn off the impurities.

Once the foil is clean, a series of five voltage scans are performed on each foil, in five-second 10 V steps from 0 V or from 1,000 V or until the detector trips. This process improves the voltage limit at which a trip occurs on a given foil, by approximately 50 V [96]. Afterwards, increasingly long stability tests are performed which have been shown to greatly reduce the discharges that occur in the chamber in the long-term.

From QC6, it has been shown that nearly every chamber which passes through the qualification procedure has a significantly higher rate of discharges at the very beginning of the tests than during the rest of the test duration, as one would expect from a working procedure. As QC6 had not yet been made a part of the quality control process for the slice test chambers, this initial state of the detectors can be compared to that of the slice test detectors within the CMS endcap during these initial months. Rather than being discharges of the traditional type that were described in Chapter 5, these are discharges caused by debris on the foils.

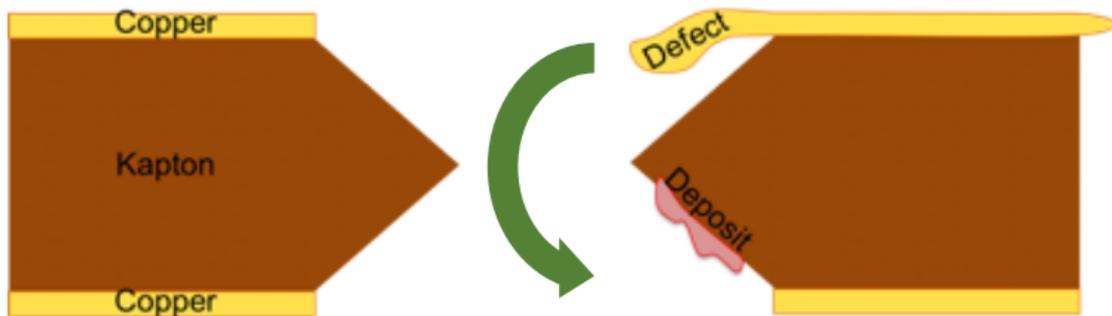


Figure 6.39 – Diagram showing example problems, including a foil manufacturing defect and a foreign deposit on the insulating polyimide inner layer of a GEM foil, which may cause discharges within a GEM detector [96]

Figure 6.39 shows a diagram which presents examples of two problems which may cause discharges of this type, including a manufacturing defect in the outer copper layer of the GEM foil and a deposit of foreign material (such as dust or other debris) on the inner polyimide layer. These impurities can result in an electrical bridge forming between the top and bottom copper layers of the foil, resulting in a discharge.

Foil defects notwithstanding, such debris is usually obliterated in the discharge process, essentially burning away to nothing. However, it has been shown that this process creates impurities in and of itself, although generally lesser than the original debris. For example,

when cleaning a foil in a laboratory setting, it is necessary to then clean the other two foils in the stack, to ensure that no further erroneous electrical connections can be made.

Thus, this initial increase in channel loss in the slice test detectors may be the result of both the burning off of debris present in the detector since production, and also a “runaway” effect of further discharges as a result of the impurities caused by that burning, particularly if the impurities are located on G3, closest to the readout board. This type of discharge is limited in scope, however, in the sense that no further discharges will occur once there is no longer any foreign material on the GEM foils. This may explain why the discharges stop entirely for GEMINIm29L2 in late May, and level off significantly for GEMINIm29L1 in early June.

6.5.8 Chamber Operation

In addition to the environmental conditions present, and initial startup issues, we must also consider the detector operational conditions. Since it is very clear that the worst of the channel loss occurred in the second half of the slice test, in 2018, and yet the environmental conditions were not so different between the 2017 and 2018, we must investigate if any changes were made to the detector operations.

When the slice test detectors began operation in March 2017, they were operated with the conditions shown in Table 6.10.

Detector ID	Slice Test Designation	HV Set Point (μA)	Minimum Gain	Average Gain	Maximum Gain
GE1/1-VII-S-CERN-0004	GEMINIm27 L1	686.8	$2.97 \cdot 10^3$	$1 \cdot 10^4$	$1.47 \cdot 10^4$
GE1/1-VII-S-CERN-0003	GEMINIm27 L2	711.7	$5.89 \cdot 10^3$	$1 \cdot 10^4$	$1.46 \cdot 10^4$
GE1/1-VII-L-CERN-0002	GEMINIm28 L1	667.5	$4.73 \cdot 10^3$	$1 \cdot 10^4$	$1.47 \cdot 10^4$
GE1/1-VII-L-CERN-0004	GEMINIm28 L2	677.3	$5.84 \cdot 10^3$	$1 \cdot 10^4$	$1.23 \cdot 10^4$
GE1/1-VII-S-CERN-0006	GEMINIm29 L1	696.5	$6.36 \cdot 10^3$	$1 \cdot 10^4$	$1.47 \cdot 10^4$
GE1/1-VII-S-CERN-0005	GEMINIm29 L2	706.0	$7.37 \cdot 10^3$	$1 \cdot 10^4$	$1.28 \cdot 10^4$
GE1/1-VII-L-CERN-0003	GEMINIm30 L1	686.9	$5.15 \cdot 10^3$	$1 \cdot 10^4$	$1.46 \cdot 10^4$
GE1/1-VII-L-CERN-0001	GEMINIm30 L2	680.2	$6.49 \cdot 10^3$	$1 \cdot 10^4$	$1.26 \cdot 10^4$
GE1/1-VII-S-CERN-0001	GEMINIm01 L1	697.3	$4.47 \cdot 10^3$	$1 \cdot 10^4$	$1.47 \cdot 10^4$
GE1/1-VII-S-CERN-0002	GEMINIm01 L2	683.2	$6.46 \cdot 10^3$	$1 \cdot 10^4$	$1.38 \cdot 10^4$

Table 6.10 – Initial operational values for the slice test, as they were known in March 2017 [97]. The HV set point is given as a divider current.

These values were determined from gain uniformity measurements as taken in QC5 in the GEM production laboratory prior to installation, such that the average gain value of each chamber would be equal to $1 \cdot 10^4$, and were calculated using the following snippet of code:

```
#-----PARAMS_GAIN-----
#Class for storing gain parameters
#Gain is defined as:
#
# G(x) = exp([0]*x+[1]) where x is an HV setpoint (either V_Drift in Volts or Divider
#                               Current in uA)
#
class PARAMS_GAIN:
    __slots__ = ['GAIN_CURVE_P0', 'GAIN_CURVE_P0_ERR', 'GAIN_CURVE_P1', 'GAIN_CURVE_P1_ERR']

    def __init__(self, gain_p0=1, gain_p0_err=1, gain_p1=1, gain_p1_err=1):
        self.GAIN_CURVE_P0=gain_p0
        self.GAIN_CURVE_P0_ERR=gain_p0_err
        self.GAIN_CURVE_P1=gain_p1
        self.GAIN_CURVE_P1_ERR=gain_p1_err

        return

    #G(x) = exp([0]*x+[1]) where x is hvPt
    def calcGain(self, hvPt):
        return np.exp(self.GAIN_CURVE_P0*hvPt+self.GAIN_CURVE_P1)

    #G(x) = exp([0]*x+[1]) where x is hvPt
    def calcGainErr(self, hvPt):
        return self.calcGain(hvPt)*np.sqrt(np.square(self.GAIN_CURVE_P0_ERR*hvPt)+np.square
            (self.GAIN_CURVE_P1_ERR))
```

However, the keen observer might notice that the detector IDs and slice test designations do not equal those from Table 4.2. The values were recorded based on the detector ID, but it is the slice test designation that would be used when applying the values in the DCS at P5. These designations were matched to their detector ID not with a physical label, but rather were matched based on Brian Dorney’s assumption from a poster hanging up in the laboratory. This assumption was, in fact, incorrect, and in fact each detector that is labeled “L1” in Table 6.10 is in fact “L2”, and vice versa. And this error was only discovered upon the extraction of the GEMINIm01 superchamber, when it was noticed that the detector ID did not match the layer orientation.

Thus, despite the careful calculation of the HV operational values to each detector’s needs, the real values as they were applied for the entirety of 2017 are listed in Table 6.11, as calculated by the above code. Gain values which are lower than the desired $1 \cdot 10^4$ are listed in blue, and gain values which are higher are listed in red for ease of understanding.

Detector ID	Slice Test Designation	Actual Average Gain
GE1/1-VII-S-CERN-0003	GEMINIm27L1	$3.7 \cdot 10^3$
GE1/1-VII-S-CERN-0004	GEMINIm27L2	$1.9 \cdot 10^4$
GE1/1-VII-L-CERN-0004	GEMINIm28L1	$7.5 \cdot 10^3$
GE1/1-VII-L-CERN-0002	GEMINIm28L2	$1.5 \cdot 10^4$
GE1/1-VII-S-CERN-0005	GEMINIm29L1	$5.9 \cdot 10^3$
GE1/1-VII-S-CERN-0006	GEMINIm29L2	$1.2 \cdot 10^4$
GE1/1-VII-L-CERN-0001	GEMINIm30L1	$1.4 \cdot 10^4$
GE1/1-VII-L-CERN-0003	GEMINIm30L2	$7.2 \cdot 10^3$
GE1/1-VII-S-CERN-0002	GEMINIm01L1	$1.5 \cdot 10^4$
GE1/1-VII-S-CERN-0001	GEMINIm01L2	$6.1 \cdot 10^3$

Table 6.11 – Actual operational values for the slice test during 2017

Initially, the HV values were kept essentially the same, with the values listed in Table 6.12 being first applied manually to the detectors on February 22nd², without the presence of LHC beam.

Slice Test Designation	Applied Divider Current (μA)	Applied Divider Voltage (kV)
GEMINIm27L1	686.7	3.678
GEMINIm27L2	718.5	3.575
GEMINIm28L1	672.6	3.578
GEMINIm28L2	677.9	3.627
GEMINIm29L1	696.8	3.702
GEMINIm29L2	706.1	3.777
GEMINIm30L1	687.7	3.649
GEMINIm30L2	682.7	3.627

Table 6.12 – Initial operational values for the slice test in 2018

It was not until April 17th, 2018³, when the DCS HV recipes were updated to include the correct values, that the mistake was finally rectified. These recipes are shown in Table 6.13, and were first applied with detectors on on April 20th, 2018. It is unclear why these values are not a simple swap of L1 and L2, as this was the stated intent of updating the recipes.

²<http://cmsonline.cern.ch/cms-elog/1029252>

³<http://cmsonline.cern.ch/cms-elog/1036672>

Slice Test Designation	Cosmic Muons	Physics Data Taking	Standby
GEMINIm27L1	711 μA , 3.754 kV	706 μA , 3.754 kV	472 μA , 2.500 kV
GEMINIm27L2	710 μA , 3.675 kV	705 μA , 3.675 kV	481 μA , 2.500 kV
GEMINIm28L1	683 μA , 3.575 kV	679 μA , 3.575 kV	475 μA , 2.500 kV
GEMINIm28L2	688 μA , 3.625 kV	683 μA , 3.625 kV	472 μA , 2.500 kV
GEMINIm29L1	720 μA , 3.775 kV	715 μA , 3.775 kV	475 μA , 2.500 kV
GEMINIm29L2	702 μA , 3.700 kV	697 μA , 3.700 kV	472 μA , 2.500 kV
GEMINIm30L1	698 μA , 3.650 kV	693 μA , 3.650 kV	476 μA , 2.500 kV
GEMINIm30L2	693 μA , 3.625 kV	688 μA , 3.625 kV	475 μA , 2.500 kV

Table 6.13 – Recipe-set operational values for the slice test in 2018

Regardless of why these new values were chosen, Table 6.14 presents the difference in applied divider current from 2017 to 2018, with increased values in red and decreased values in blue. Somewhat unsurprisingly, the three chambers with large, sudden losses of channels are also the three detectors with the largest difference in applied divider current from 2017 to 2018.

Slice Test Designation	2017 Applied Divider Current (μA)	2018 Divider Current (μA)	Difference (μA)
GEMINIm27L1	686.8	706.0	+19.2
GEMINIm27L2	711.7	705.0	-6.7
GEMINIm28L1	667.5	679.0	+11.5
GEMINIm28L2	677.3	683.0	+5.7
GEMINIm29L1	696.5	715.0	+18.5
GEMINIm29L2	706.0	697.0	-9
GEMINIm30L1	686.9	693.0	+6.1
GEMINIm30L2	680.2	688.0	+7.8

Table 6.14 – Change in operational values for the slice test from 2017 to 2018

As per the LHC schedule, beam recommissioning began on March 30th, 2018, the intensity began to ramp up on April 17th, 2018, and stable proton physics beams began on May 5th, 2018. And thus, we finally have a reasonable explanation for the intense period of channel loss on GEMINIm29L1.

Previously, we could not make a concrete connection between the channel loss and the luminosity of the LHC beam. However, this assumed that the operational values of the detectors were consistent throughout the duration of the slice test. With the knowledge that this was not the case, we can see a much clearer pattern emerge.

GEMINIm29L1 began its most intense period of channel loss on April 22nd, with the sudden loss of 15 channels over two VFATs, less than a week after having changed to the new operational HV values. If we look at the period of time from when the new operational values were first implemented to the S-curve that showed the first of these damaged channels, we see a very clear picture emerge from the luminosity.

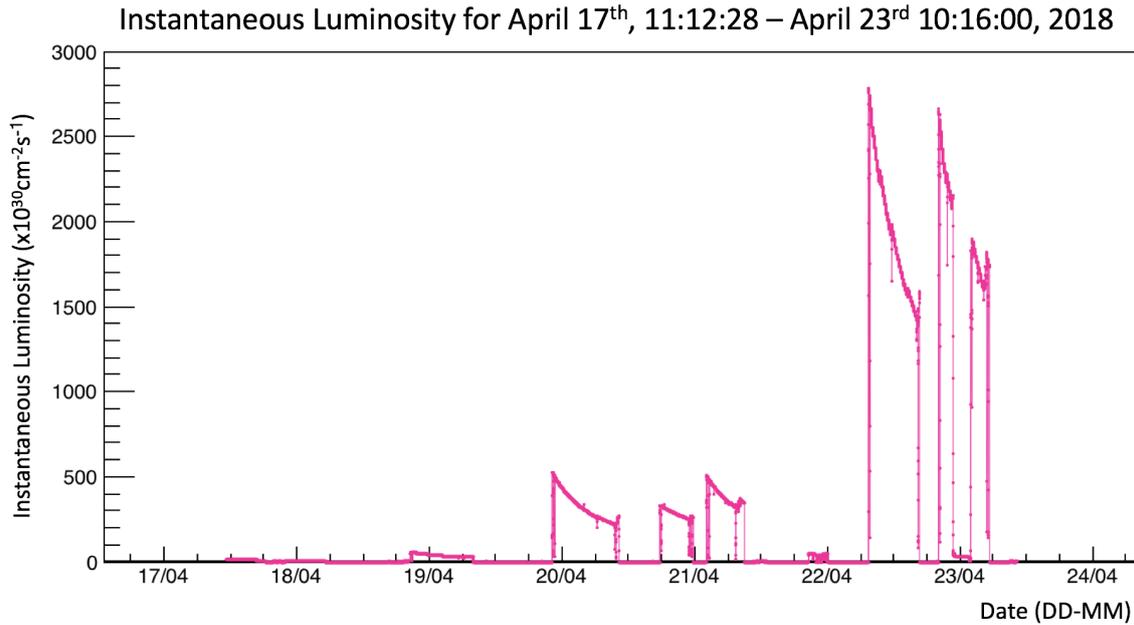


Figure 6.40 – LHC beam luminosity measured by the CMS detector from 2017.04.17.11.22 to 2017.04.23.10.16. The margins on either side of the data represent empty bins, not zero luminosity.

Figure 6.40 shows the luminosity during this time period, 2017.04.17.11.22 (HV recipe change) to 2017.04.23.10.16 (first large loss of channels). Note that, unlike in Figure 6.13, the zero luminosity between 22h00 on April 22nd and 02h00 on April 3rd is not an issue on the part of CMS as the result of a subdetector exclusion of VdM scan, but is an actual levelling of the beam luminosity of Fill 6595⁴ as CMS requested a low pile-up run that day⁵. Here, we see that the first high-luminosity beam after the recipe change occurs precisely between the window of channel damage to GEMINIm29L1: 2018.04.22.11.58 to 2018.04.23.10.16.

In summary, GEMINIm29L1 was operated under lower-than-ideal conditions for the first half of the slice test. When this was rectified, the operational values it was set to were larger than the ideal values that had been previously calculated, and represented the second highest change in applied divider current among the slice test chambers. These new values meant that the chamber was set to operate at a higher average gain than the generally-safe $1 \cdot 10^4$, and then, from the first interaction with intense beams, the detector began discharging rapidly, losing channels as a result. These discharges continued nearly daily until October 22nd, with the last discharge in any of the slice test chambers recorded on October 24th.

Indeed, this sudden cessation of channel loss further bolsters the case for our explanation above. October 23rd was the final day of proton beams in the LHC prior to LS2, with October 24th being the first day of a machine development period. So we see that the chambers which experienced the largest change in operational parameters consistently lost channels while under beam, and the channel loss immediately stopped upon the cessation of beam, exactly as one would expect if the damage was caused by discharges as a result of highly-ionizing particles (HIPs) interacting with a detector whose amplification is particularly strong.

⁴<https://cmswbm.cern.ch/cmsdb/servlet/FillReport?FILL=6595>

⁵<http://schwick.web.cern.ch/schwick/rctools/dailyReport/detailReport/display/2018-04-23>

What is left missing from our explanation, then, is mid-2018. Between June 27th and July 14th, 2018, we saw a period of high channel loss activity among the slice test detectors. GEMINIm27L1 and GEMINIm28L1 saw their first surge of channel loss, GEMINIm29L2 saw a resurgence of channel loss after its initial bout in 2017, and conversely, GEMINIm29L1 saw a uniform decrease in its channel loss rate. So, the question is, what happened?

From the official LHC schedule, we see that June 18th began the first technical stop since proton beams were resumed for the year. The experimental cavern was then open for access for five days, followed by a week and a half of “special physics runs”, including more VdM scans and a $\beta^* = 90$ m run. The luminosity during this period is shown in Figure 6.41.

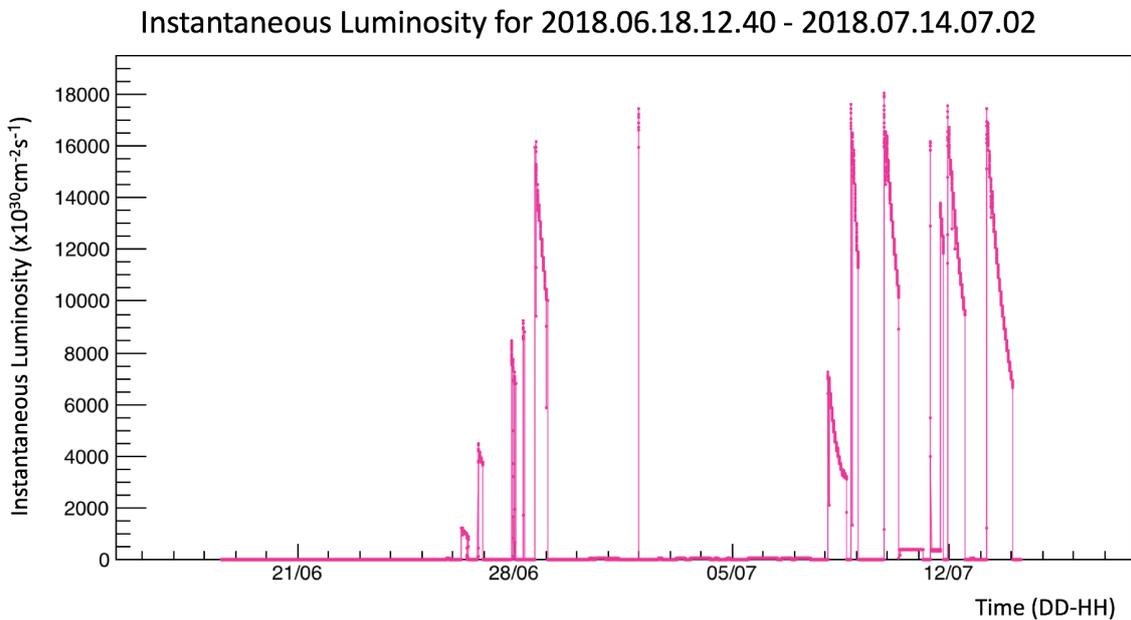


Figure 6.41 – LHC beam luminosity measured by the CMS detector from 2018.06.18.12.40 to 2018.07.14.07.02. The margins on either side of the data represent empty bins, not zero luminosity.

And, indeed, we see some a correlation. During the technical stop, when the chambers were vulnerable to human access, GEMINIm29L2 experienced the first of several lost channels, and the loss rate change for GEMINIm29L1 took effect, indicating that something may have occurred during the interventions that physically affected the slice test detectors. This is followed by the first reintroduction of beam for the special physics runs, which coincides directly with the first channel loss for GEMINIm28L1, and the following period of beam, which coincides directly with the start of the channel loss for GEMINIm27L1. While it is not entirely clear what intervention could have caused these problems, we can nevertheless surmise that the change in operational parameters further exacerbated the problem.

6.6 Plan of Action Moving Forward

Following these events, it has become abundantly clear that we must ensure three things:

- That the GEM detectors are operated at their ideal gain and HV settings to prevent

the Raether limit being exceeded, resulting in a discharge within the detector

- That the foils are cleaned well before the detectors are exposed to beam to prevent discharges as a result of impurities on the foils
- That the readout electronics must be protected from any discharges which occur despite the previous two precautions

The latter precaution will be discussed in great detail in Section 7.2, but we can briefly discuss the former two precautions here.

In regards to the first precaution, this is particularly important for GE1/1, as the two detectors of a given superchamber will share the same power supply. That is, a given HV channel has a branch that goes to L1 and a branch that goes to L2, such that a given electrode (say, $G1_{top}$) will receive the same power values for both chambers.

In order to ensure this type of operation is safe, when the superchambers are being coupled in the lab, prior to being brought to P5, they are coupled together based on sharing similar gain profiles. Thus, a set of operational values that is good for one layer is good for the other layer, and a mix-up like the one that happened in the slice test is, by definition, impossible.

For the second precaution, we turn to the following section.

6.6.1 HV Training

As discussed in Section 6.5.7, for all GE1/1 detectors, a sixth quality control step was introduced which cleaned the foils prior to the addition of any readout electronics to the detector. However, this is not the only foil cleaning / stability procedure which was put into place. The detectors now undergo a process called “HV training”, both as part of the final QC8 validation step, and then later once they are installed within the CMS endcap. This procedure, described below, reveals any shorts that may be present in the detector, further cleans the foils, and increases the stability of the detectors when under HV.

- Check for the presence of shorts
 - Apply 50 V for 1 minute to $G1_{top}$, $G2_{top}$, and $G3_{top}$.
 - After ramping up, check that the monitored current returns to 0 μA .
 - If it does not, power down the chamber and prepare for investigation. A current of $\geq 1\mu\text{A}$ indicates a short.
- Train G1. All others remain off.
 - Set G1Top to 100 V and monitor for 10min
 - Set G1Top to 200 V and monitor for 20 min
 - Set G1Top to 300 V and monitor for 30 min
 - Set G1Top to 350 V and monitor for 1 hour
 - Set G1Top to 380 V and monitor for 1 hour
 - Set G1Top to 400 V and monitor for 1 hour
 - Set G1Top to 410 V and monitor for 1 hour
 - Set G1Top to 420 V and monitor for 1 hour

- Set G1Top to 430 V (the nominal value) and monitor for 12 hours
- Train G2 by repeating the above process for $G2_{top}$. All others remain off.
- Train G3
 - Set G3Top to 100 V and monitor for 10min
 - Set G3Top to 200 V and monitor for 20 min
 - Set G3Top to 300 V and monitor for 30 min
 - Set G3Top to 340 V and monitor for 1 hour
 - Set G3Top to 360 V and monitor for 1 hour
 - Set G3Top to 370 V and monitor for 1 hour
 - Set G3Top to 380 V (the nominal value) and monitor for 12 hours
- Power $G1_{top}$, $G2_{top}$, and $G3_{top}$ at their nominal values and monitor for 12 hours
- Train half of the GEM stack (Drift, $G1_{top}$, $G1_{bottom}$, $G2_{top}$). All others remain off.
 - Set equivalent divider current⁶ to 660 μA and monitor for 1 hour
 - Set equivalent divider current to 670 μA and monitor for 1 hour
 - Set equivalent divider current to 680 μA and monitor for 1 hour
 - Set equivalent divider current to 690 μA and monitor for 6 hours
- Train the full gem stack, minus the induction gap, with $G2_{bottom}$ powered at 50%, by repeating the above process.
- Train the full gem stack, minus the induction gap, with $G2_{bottom}$ powered at 70%, by repeating the above process.
- Train the full gem stack, minus the induction gap, with $G2_{bottom}$ fully powered, by repeating the above process.
- Check the front-end electronics for any damage, take a full set of S-curves and analyze them.
- Train the full GEM stack, including the induction gap, at nominal values.
 - Set equivalent divider current to 660 μA and monitor for 1 hour
 - Set equivalent divider current to 670 μA and monitor for 1 hour
 - Set equivalent divider current to 680 μA and monitor for 1 hour
 - Set equivalent divider current to 690 μA and monitor for 12 hours

Initially, this procedure is performed in pure CO_2 , provide better protection against electrical damage to the detectors, as CO_2 is a quenching gas. This allows for an identification of the most problematic chambers before moving to the operational gas mixture.

⁶Note that while we use "equivalent divider current", there is no physical HV divider present on the GE1/1 detectors. Rather, this is a calculation made by the DCS which takes in the current that would be applied to an HV divider, and calculates the values required for each of the GEM electrodes of the multi-channel power supply.

Barring significant problems, the procedure can then be repeated in Ar:CO₂ 70:30. Although a long process, this thoroughly removes any residual contaminants in the detectors, reducing their discharge potential and greatly improving their operational stability.

With these improvements made to the validation and commissioning of the GE1/1 detectors, we can turn to the next chapter to discuss the third and final precaution for GE1/1.

Chapter 7

Discharge Damage Mitigation Studies

7.1 Channel Loss Probability

The rate of channel loss in a given GEM detector is given by the following pseudo-formula: channel loss rate = (background rate) · (discharge probability) · (propagation probability) · (damage probability). As the background rate is a function of the position of the detectors within CMS, and therefore is an unchangeable quantity, in order to reduce the rate of channel loss it is necessary to focus on reducing the discharge probability, the propagation probability, and the damage probability.

The original design of the GE1/1 detectors made all attempts, based on past experience, to minimize both the rate of discharges and the potential damage caused by them, so as to preserve the intended lifetime of the GE1/1 system. In order to reduce the discharge probability, the detectors were designed to be operated at a gain of less than $2 \cdot 10^4$, and the voltage across G3, where the charge density is higher, would be lowered. And in order to reduce the probability of damage as a result of discharges, a protection resistor was installed on the tops of the GEM foils, in order to decouple the energy stored in the cables and power supply. The tops of the GEM foils were also segmented into HV partitions of $\leq 100 \text{ cm}^2$ each - 40 sectors for short detectors and 47 sectors for long detectors [98] - in order to limit the energy available to discharges.

As seen in the slice test, however, these attempts were not sufficient to prevent discharged-induced channel loss, and further actions had to be taken in order to mitigate any possible channel loss on the final GE1/1 detectors. However, by the time the channel loss in the slice test was discovered and the cause of the damage was deduced, the GE1/1 detectors were already well into production. As a result, the possible channel loss mitigation strategies for the GE1/1 system were limited to lowering the damage probability alone, as the discharge probability is largely intrinsic to a given gas detector at its chosen operating settings, and lowering the propagation probability would require changing the physical structure of the inner detector hardware such as the GEM foils. This limitation, coupled with the strict time frame in which the changes had to be made in order to comply with the CMS global schedule for LS2, meant that only two mitigation strategies were ultimately pursued: external protection circuits in front of the inputs to the VFAT3 hybrid, and additional resistance added to the on-detector HV filter.

7.2 VFAT3 Protection Circuits

In order to investigate the protection capabilities of any proposed protection circuits, tests were performed at INFN Bari using the set-up shown in Figure 7.1. Here, the HV source charges the capacitor, C, by means of the resistor. This resistor has the high resistance value of $1 \text{ G}\Omega$ in order to limit the current and have a slow, controllable rate of discharges, as well as to decouple the high voltage during the discharge process. The gas discharge tube component has an intrinsic property known as the “DC sparkover voltage”, which is the voltage required between the two pins of the gas discharge tube in order to create a discharge within the tube. When the capacitor becomes sufficiently charged and this

sparkover voltage is achieved, the gas discharge tube creates a spark, and as the tube is in series with the input protection circuit, this spark is injected into the circuit as the measured electrical discharge.

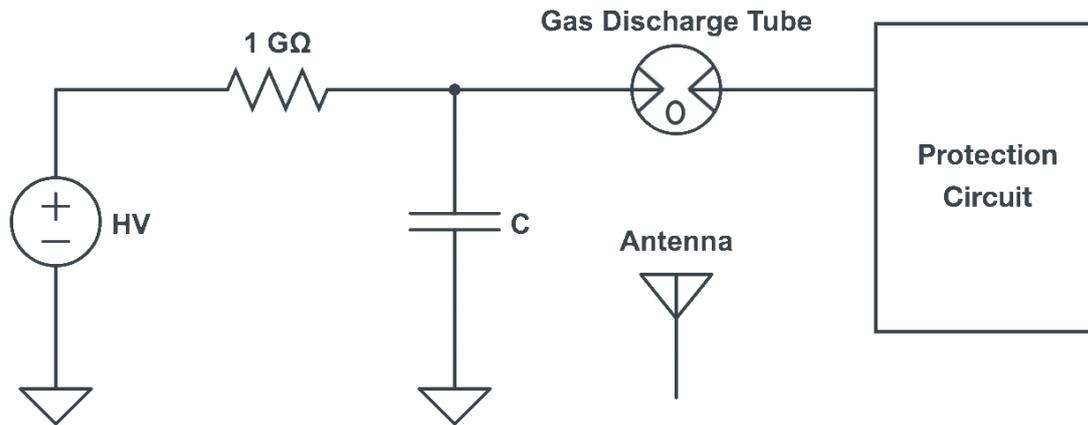


Figure 7.1 – Discharge testing circuit for VFAT3 hybrids

With this set-up, the energy of the incoming discharges could be tuned by changing either the capacitor or the gas discharge tube. This tuning allowed for the precise determination of the maximum energy at which the protection circuit would fail, according to Equation 7.1, and, if applicable, the number of discharges at that energy required for failure.

$$E_{discharge} = \left(\frac{1}{2}\right)CV_{th}^2 \quad (7.1)$$

For example, for a $C = 15$ nF discharge capacitor and a tube threshold voltage of $V_{th} = 460$ V, the discharge energy applied to the protection circuit would be $E = \left(\frac{1}{2}\right)(15 \cdot 10^{-9} \text{ F})(460 \text{ V})^2 = 1.587 \text{ mJ}$ [99].

Crucial to this set-up is the presence of the antenna, which detects the electromagnetic pulse generated during a discharge. This antenna is connected to an oscilloscope, with an example output seen in Figure 7.2, which allows the user to monitor the discharges as they occur, noting their time and amplitudes. In this way, the user can count the precise number of discharges necessary to destroy a given protection circuit.

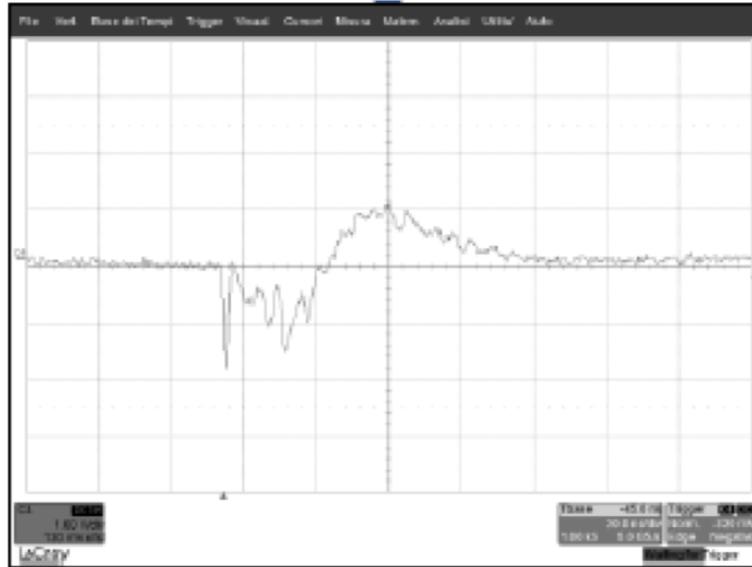
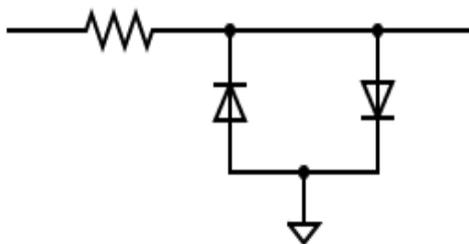


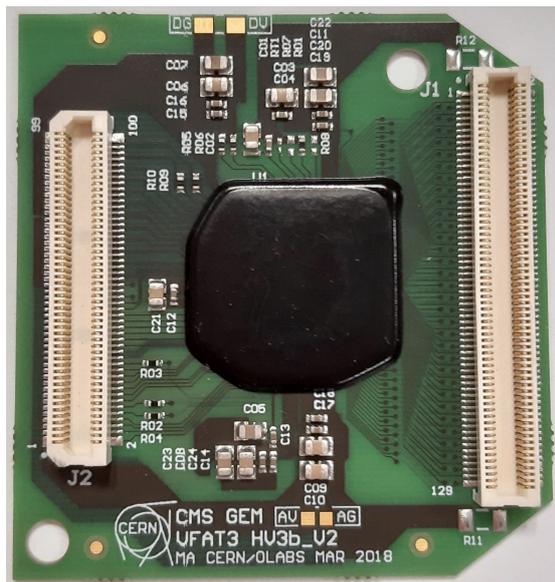
Figure 7.2 – Output of the antenna for the protection circuit testing set-up

7.2.1 HV3b_v2

HV3b_v2 is the baseline VFAT3 hybrid, whose chip contains the default internal protection circuit - a grounded series of diodes and an 8.5Ω current-limiting, heat-dissipating resistor. The circuit diagram of this protection circuit, and its corresponding application on a VFAT3 hybrid, can be seen in Figure 7.3. The schematic for the internal protection circuit can be seen in Figure 7.4, where both the 8.5Ω resistor can be seen as the second structure from the left, followed by the 4-stage protection diodes on the right.



(a) Circuit Diagram for HV3b_v2



(b) Example of a HV3b_v2 hybrid

Figure 7.3 – HV3b_v2 protection circuit diagram and physical implementation

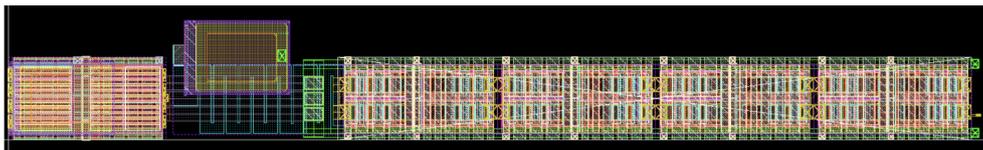


Figure 7.4 – The schematic for the VFAT3 internal protection circuit [100]



Figure 7.5 – Destroyed internal protection resistor on a VFAT2 chip (left) [100]

When a discharge occurs, almost all of the discharge energy is dissipated through the resistor in this protection circuit. When a given discharge is above a certain threshold value, the resistor is destroyed, as seen on the left of Figure 7.5, in comparison to the undamaged resistor on the right.

In order to test the HV3b_v2 internal protection circuit using Bari's set-up, the following expected operational conditions were assumed: $V_{drift} = 430$ V, $C_{GEM,short} \approx 3$ nF, $C_{GEM,long} \approx 4.4$ nF, and maximum discharge energy in GE1/1 = $410 \mu\text{J}$. From real measurements with the chambers, the assumed condition of a maximum discharge energy of $410 \mu\text{J}$ was shown to be incorrect - while this value is true for a single, simple discharge, a propagating discharge can result in more than one foil releasing its energy into the discharge, thus increasing the energy cap. However, as this value was assumed in all tests described in the following sections, it is listed here for clarity.

Multiple discharge energies were then utilized in order to determine the precise energy and number of discharges needed to render the protection circuit inoperative. The results from these tests can be found in Table 7.1.

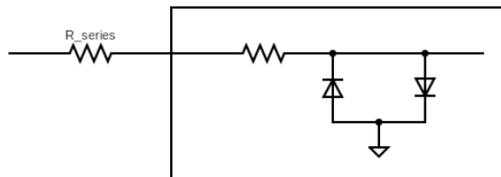
C (pF)	V_{th} (V)	E (μJ)	$N_{d,ok}$	$N_{d,fail}$
220	460	23.2760	26	0
242	460	26.6036	25	0
267	460	28.2486	1*	3
330	460	34.9140	1	2
440	460	46.5520	0	1

Table 7.1 – Results of the discharge testing of the HV3b_v2 protection circuit

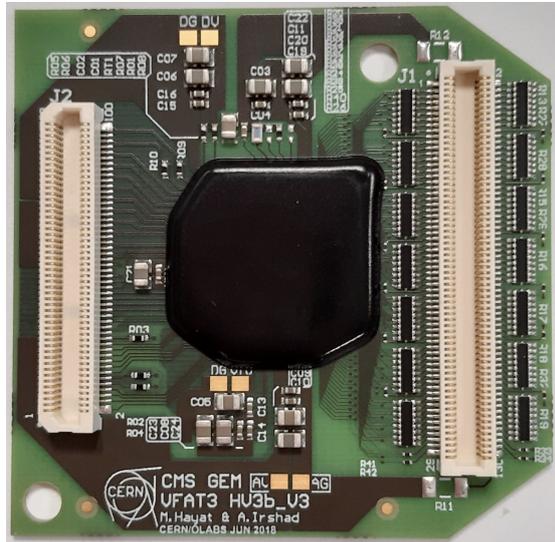
For the test at $E = 28.2486 \mu\text{J}$, one tested channel survived a single discharge, but broke after the second discharge. Three proceeding channels broke after the first discharge. Therefore it can still be said that $\sim 28 \mu\text{J}$ is the limit for the HV3b_v2 internal protection circuit. As the average discharge energy seen in the sustained operations studies was $\sim 110 \mu\text{J}$, the HV3b_v2 circuit is clearly insufficient for operation within the CMS detector, and a more robust solution would be needed. To that end, two additional external protection

circuits were proposed, called HV3b_v3 and HV3b_v4.

7.2.2 HV3b_v3



(a) Circuit Diagram for HV3b_v3



(b) Example of a HV3b_v3 hybrid

Figure 7.6 – HV3b_v3 protection circuit diagram and physical implementation

HV3b_v3 is the first of the two proposed external protection circuits, which adds a series of resistors in front of the default internal protection circuit and input channels of the VFAT3 hybrid. These resistors would dissipate a portion of the energy of any incoming discharge, protecting the channels. Its circuit diagram, and its corresponding application on a VFAT3 hybrid, can be seen in Figure 7.6. In comparison to the HV3b_v2 hybrid in Figure 7.3b, here we can clearly see the additional resistors on either side of the Panasonic connector on the right of the image, external to the ASIC which is protected underneath the black globtop in the center of the hybrid.

A number of different resistance values were tested for their resilience against discharges, ranging from 0Ω up to 470Ω , with the results seen in Figure 7.7. For each resistance value, discharges of increasingly-greater energy were applied until the resistor failed, with survived discharges marked as green data points and damaging discharges marked as red data points, and the blue line dividing the two. From these results, we see that a series resistance of 470Ω is durable against discharges of energies up to 1.5 mJ , an order of magnitude larger than the discharges seen in the sustained operations studies.

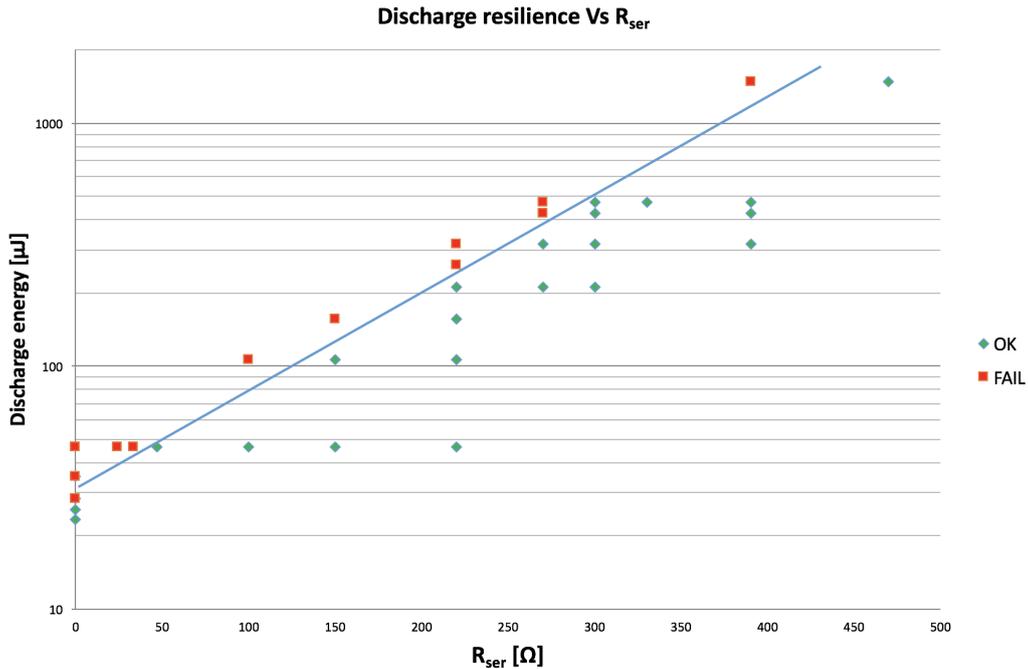


Figure 7.7 – The discharge resilience of different series resistors [93]

Producer	Part Number	Type
Panasonic	EXB2HV471JV	Resistor Array
CTS	741X163471JP	Resistor Array
Bourns	CAY16471J8LF	Resistor Array
Koa Speer	RK73B1HTTC471J	Single Resistor
Panasonic	ERJ1GNJ471C	Single Resistor
Multicomp	MCRE000132	Single Resistor
Yageo	RC0201FR-07470RL	Single Resistor
Walsin	WR02X4700FAL	Single Resistor

Table 7.2 – Resistors tested for use on HV3b_v3 [100]

Many resistors of the same value were tested to find the most resilient and suitable, as listed in Table 7.2.

In order to test the resiliency of each resistor or resistor array, each subject was tested in the same way as the protection circuits, as in Figure 7.1. The resistance was measured first, prior to any discharges. A discharge was then applied to the resistor, after which the resistance was measured again. If the resistance remained constant, the discharge-and-measurement process was repeated until the resistance changed, at which point the resistor would be changed out for another resistor of the same type in order to test the uniformity of the response. Of the tested resistors and resistor arrays, Panasonic’s EXB2HV471JV resistor array was chosen as being the most suitable, with its results seen in Figure 7.8, compared to a less-stable resistor array seen in Figure 7.9.

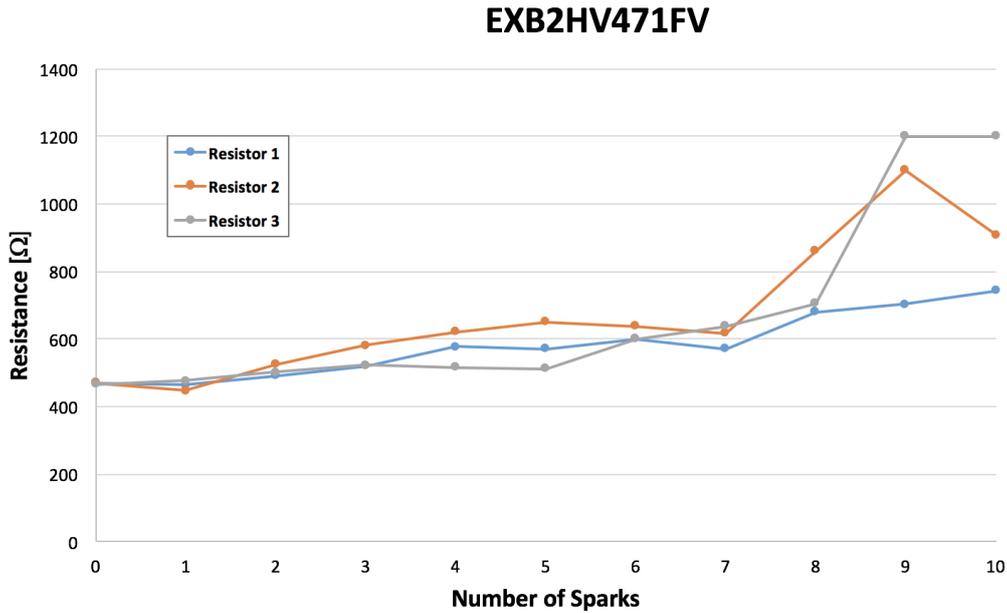


Figure 7.8 – Stability of EXB2HV471JV resistor array under discharges [100]

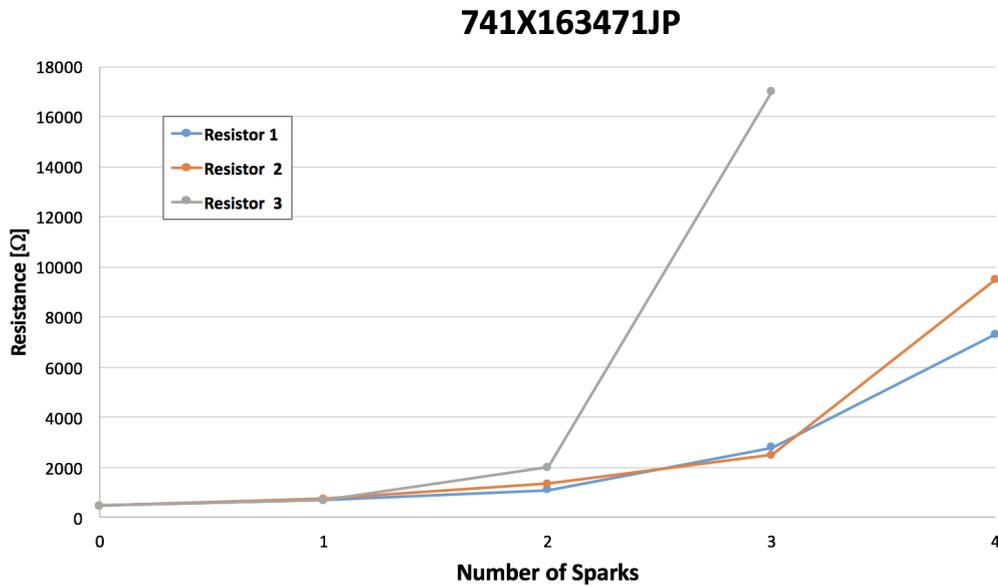


Figure 7.9 – Stability of 741X163471JP resistor array under discharges [100]

The EXB2HV471JV resistor array has been successfully used in previous LHC experiments and have been shown to have a viable radiation tolerance for this application, thus negating the need for any additional radiation testing. However, as it has not been previously used in a GEM application, it needed to be studied in terms of any additional noise or cross-talk they might inject into the system. These results, which were obtained by simulation, can be found in Table 7.3, utilizing the internal injection circuit seen in Figure 7.10a and the external injection circuit seen in Figure 7.10b

Discharge Capacitance (pF)	Noise, 0 Ω (fC)	Noise, 330 Ω , Internal Injection (fC)	Noise, 330 Ω , External Injection (fC)	Noise, 470 Ω , Internal Injection (fC)	Noise, 470 Ω , External Injection (fC)
1	0.193	0.194	0.194	0.194	0.194
12	0.199	0.240	0.242	0.254	0.258
23	0.210	0.323	0.333	0.351	0.372
34	0.225	0.413	0.439	0.448	0.499
45	0.244	0.499	0.548	0.534	0.627
56	0.265	0.578	0.659	0.608	0.755
67	0.288	0.649	0.769	0.673	0.881
78	0.314	0.714	0.878	0.728	1.004
89	0.340	0.772	0.985	0.776	1.126
100	0.367	0.824	1.091	0.818	1.244

Table 7.3 – Noise injected by the HV3b_v3 protection circuit, set to maximum peaking time

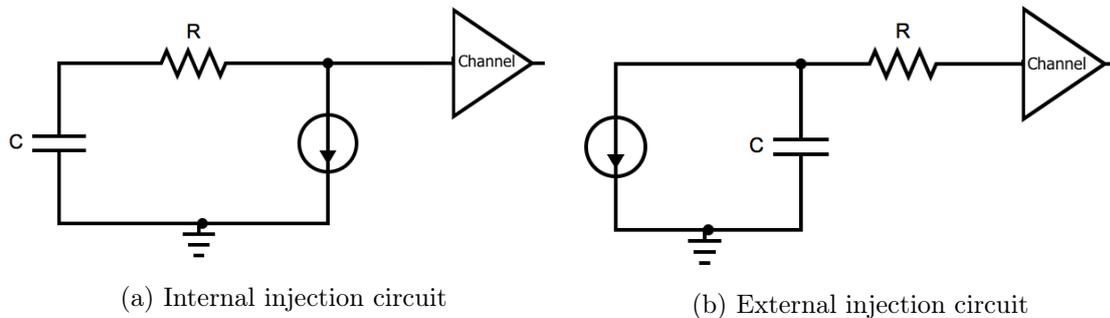


Figure 7.10 – Injection circuits used for the noise study of the HV3b_v3 protection circuit resistors [100]

The two injection points each produce a different differential amplifier output, as seen in Figure 7.11. As the injected charge for the external injection point follows the same path as the GEM signal, this is taken to be the correct injection type. This intuitively makes sense, given that the internal injection circuit is simply that of the VFAT3’s internal calibration circuit, as described on page 24 of the Reference [46], rather than a circuit which more closely represents a true signal.

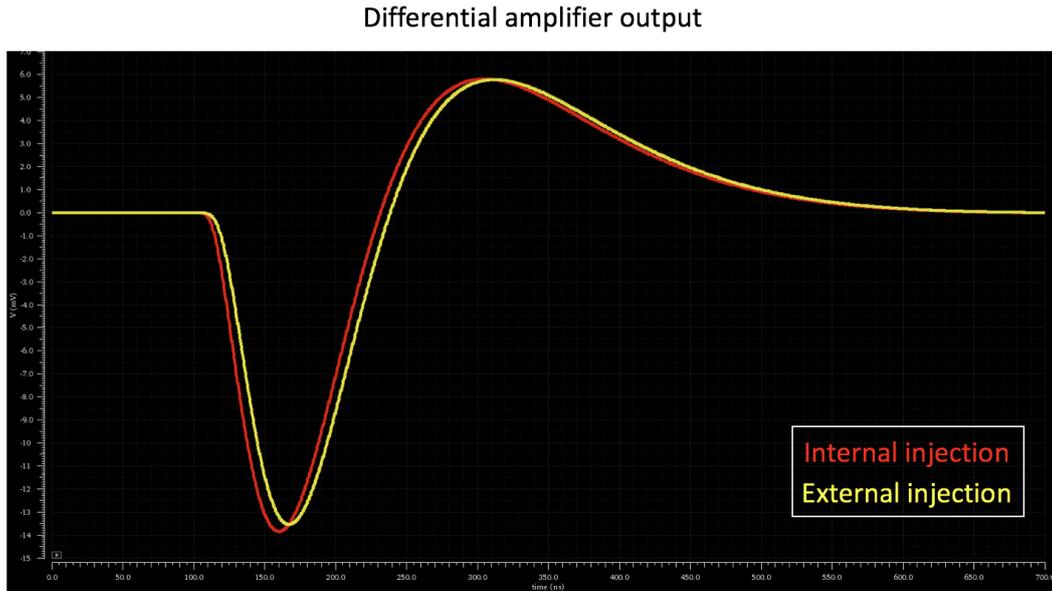


Figure 7.11 – Simulation of differential amplifier output for internal and external injection circuits [100]

The choice of the external injection point is further backed by the beginnings of noise saturation seen for the internal injection circuits when the values from Table 7.3 are plotted on a graph, as in Figure 7.12. Here, we see that the simulated noise measurements for the internal injection circuit begins to saturate at higher values of the discharge capacitance, whereas this saturation is not seen for the noise from the external injection circuit, thus making the internal injection unsuitable for a significant portion of the simulated measurements.

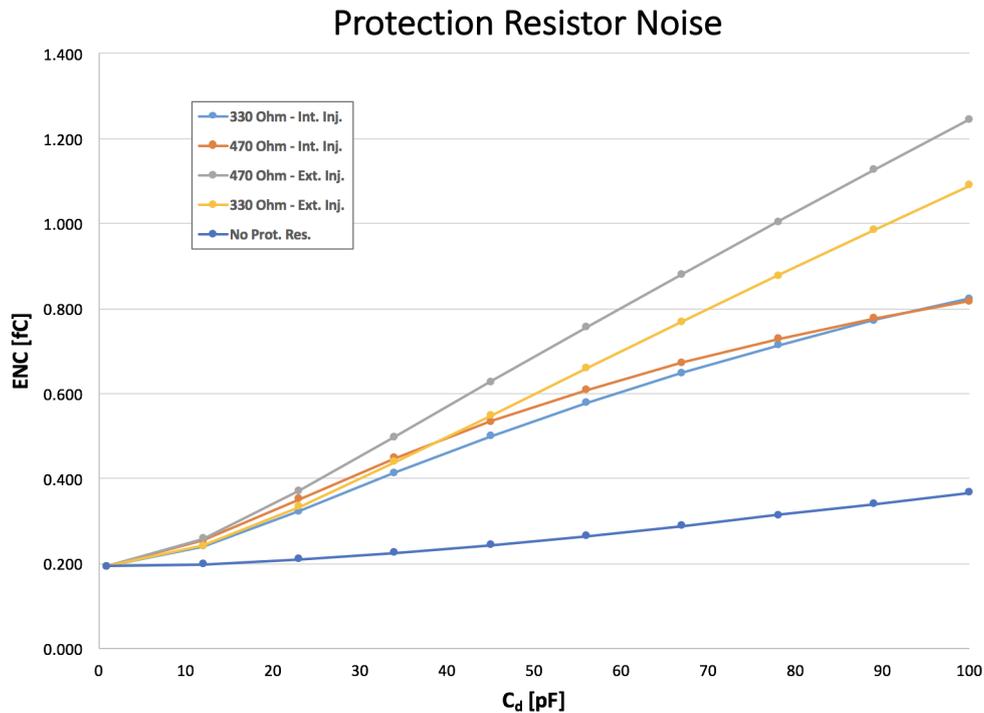


Figure 7.12 – Noise injected by the HV3b_v3 protection circuit [100]

Thus, from Table 7.3, we can compare column 2 against column 6 to see the amount

of noise injected by the chosen $470\ \Omega$ EXB2HV471JV resistor array, as in Table 7.4.

Discharge Capacitance (pF)	Noise, $0\ \Omega$ (fC)	Noise, $470\ \Omega$, External Injection (fC)	Absolute Increase in Noise (fC)	Percent Increase in Noise (%)
1	0.193	0.194	0.001	0.518
12	0.199	0.258	0.059	29.648
23	0.210	0.372	0.162	77.143
34	0.225	0.499	0.274	121.777
45	0.244	0.627	0.383	156.967
56	0.265	0.755	0.490	184.905
67	0.288	0.881	0.593	205.902
78	0.314	1.004	0.690	219.745
89	0.340	1.126	0.786	231.176
100	0.367	1.244	0.877	238.964

Table 7.4 – Noise injected by the HV3b_v3 protection circuit, set to maximum peaking time

In addition to noise, also of interest is any additional cross-talk that may be added to the system by the HV3b_v3 protection circuit. In order to study that, a simulation was done using the set-up seen in Figure 7.13. Here, the simulation settings were $t_p = 100\ \text{ns}$, $C_d = 4\ \text{pF}$, $C_c = 8\ \text{pF}$. Using Equation 7.2, the cross-talk was determined for $0\ \Omega \leq R_{prot} \leq 400\ \Omega$. Unfortunately, no simulation data exists for the ultimately-chosen value of $R_{prot} = 470\ \Omega$, although a value of approximately 20% can be extrapolated from the existing data. These cross-talk values are plotted in Figure 7.14.

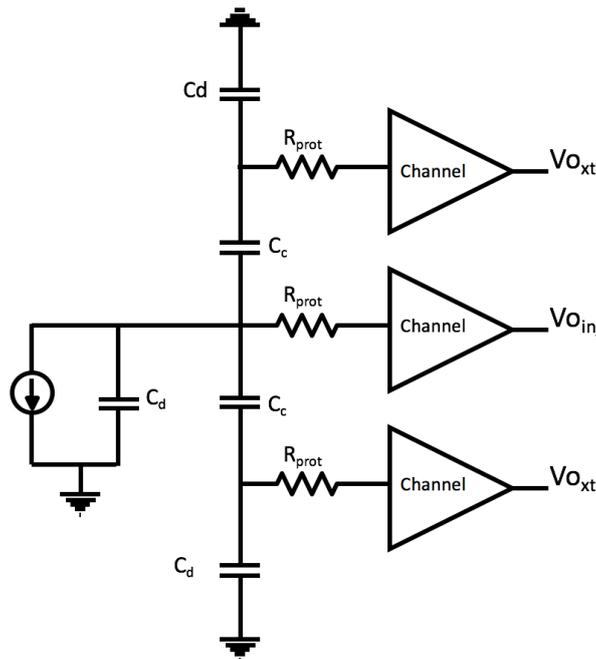


Figure 7.13 – Simulation set-up for the cross-talk study of the HV3b_v3 protection circuit [100]

$$XT = \left(\frac{V_{Oxt}}{V_{Oinj}} \right) \cdot 100 \quad (7.2)$$

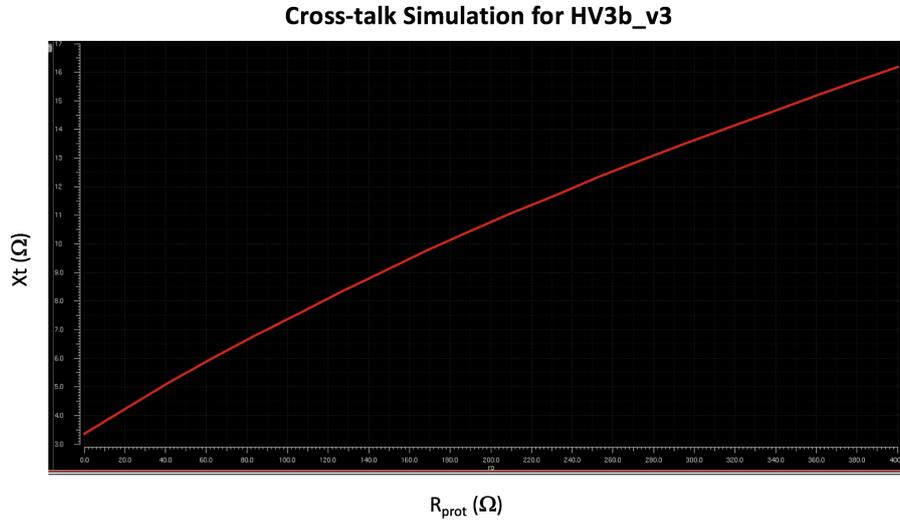
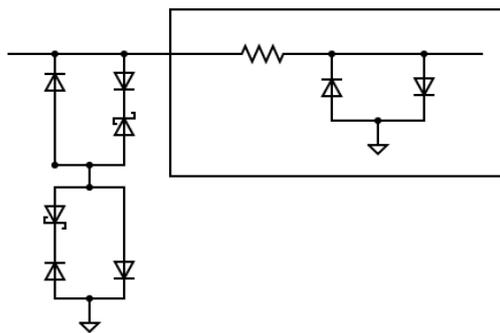


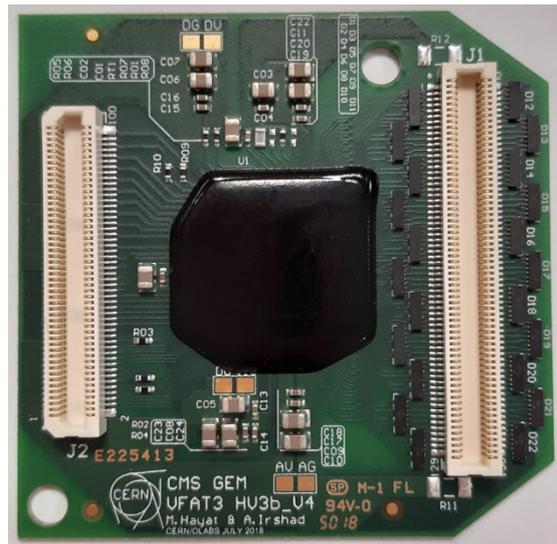
Figure 7.14 – Cross-talk simulation results for the HV3b_v3 protection circuit [100]

Despite the drawbacks of the increased noise and cross-talk, an initial run of 75 HV3b_v3 VFAT3 hybrids was ordered from Rui De Oliveira at CERN and sent to HybridSA for assembly. Testing of the HV3b_v3 protection circuit using Bari's set-up showed that the circuit was capable of withstanding 500 electrostatic discharges at $470 \mu\text{J}$ per discharge without sustaining any damage to the input channels [101].

7.2.3 HV3b_v4



(a) Circuit Diagram for HV3b_v4



(b) Example of a HV3b_v4 hybrid

Figure 7.15 – HV3b_v4 protection circuit diagram and physical implementation

HV3b_v4 is the second of the two proposed external protection circuits, which adds a set of diodes, ON Semiconductors' ESD-7016 low capacitance protection diodes, in parallel

in front of the default internal protection circuit and input channels of the VFAT3 chip. Its circuit diagram, and its corresponding application on a VFAT3, can be seen in Figure 7.15. In comparison to the HV3b_v2 hybrid in Figure 7.3b, here we can clearly see the additional diodes on either side of the Panasonic connector on the right of the image, external to the ASIC which is protected underneath the black globtop in the center of the hybrid. These diodes injected no additional cross-talk or noise into the system, making them an ideal candidate for the protection solution.

An initial run of one thousand HV3b_v4 VFAT3 chips was ordered from the PCB manufacturer Fineline Global. Testing of the HV3b_v4 protection circuit using Bari's set-up showed that the circuit was capable of withstanding 540 electrostatic discharges at 470 μJ per discharge without sustaining any damage to the input channels [101]. The noise levels did not increase following the discharge testing, remaining at ~ 800 electrons [93].

7.2.4 On-Detector Testing

With the radiation tolerance and robustness of both HV3b_v3 and HV3b_v4 established, the final step towards choosing the ultimate solution was on-detector testing. This was done in the GEM production lab in Building 904 at the CERN Preveessin site.

First the different versions were placed on a $10 \times 10 \text{ cm}^2$ detector with a 110 k Ω HV filter resistance and 0 Ω drain resistance, and filled with the standard gas mixture of 70% Argon, 30% CO_2 . An ^{241}Am source was placed on the detector to provide 5.5 MeV α particles in order to trigger the discharges. The probability of damage under discharges of different energies was recorded, with the discharge energies calculated from Equation 7.3, where C_{foil} is the capacitance of the foil plus any additional capacitors and V_{applied} is the applied voltage.

$$E_{\text{discharge}} = \left(\frac{1}{2}\right)C_{\text{foil}}V_{\text{applied}}^2 \quad (7.3)$$

In order to obtain different discharge energies, two methods could be used. The easier of the two methods simply requires changing the voltages applied to each of the GEM foils. However, this method would of course change the electric fields present within the detector. In order to preserve these fields, the second method of changing the capacitance could be used. As the capacitance of the foil is an intrinsic property based on the size of the foil, changing the capacitance requires adding additional external capacitors between $G3_{\text{bottom}}$ and ground. For the result plot seen in Figure 7.16, this second method was used. Here, although both HV3b_v3 and HV3b_v4 provide greater damage resilience compared to the default HV3b_v2, HV3b_v4 resulted in a much less steep curve and the ability to withstand discharges of much greater magnitude than HV3b_v3.

After these promising results, the different versions were then moved to a standard GE1/1 detector for final testing, once again using an ^{241}Am source to trigger discharges. These results, seen in Figure 7.17, however, were unexpected and far less promising.

For each type of protection circuit, the damage probabilities were higher on the GE1/1 detector than on the smaller $10 \times 10 \text{ cm}^2$ detector. This was expected, as the larger detector allows for more complex discharge propagation which allows the discharges to accumulate more energy. What was not expected, however, was that upon a single discharge of 1 mJ, all of the HV3b_v4 input channels suffered significant damage, resulting in the total destruction of the chip. This was observed multiple times and was not a statistical outlier.

Although it is not yet fully understood why this occurred, the current hypothesis is that

the large current induced by the discharge evacuated through the diode to the ground line in the GEB. This ground line, however, isn't thick enough to support such a large current, and has a small impedance. When the current sees the impedance, it creates a local voltage drop on the ground line. As the VFAT3 chip can only support $\sim 1\text{-}2\text{ V}$ between the low voltage and ground lines, this larger voltage drop results in the total destruction of the chip.

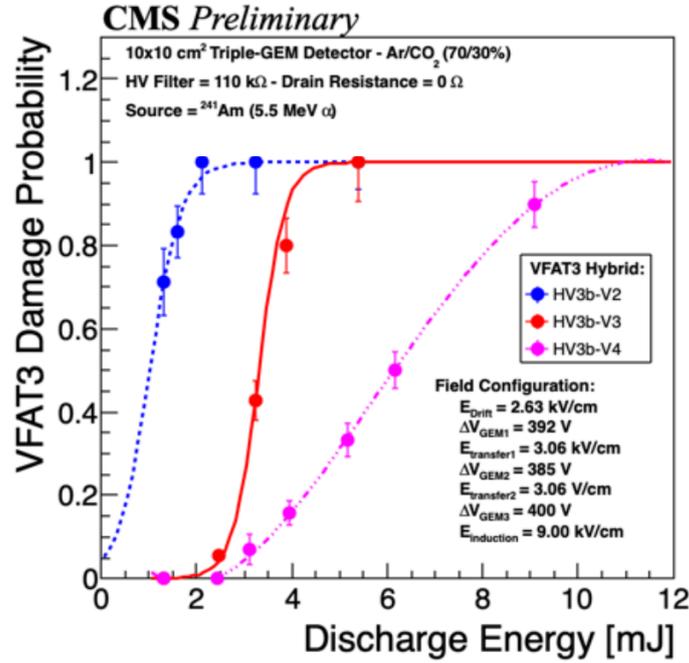


Figure 7.16 – VFAT3 damage probability vs discharge energy for different protection circuits on a $10 \times 10\text{ cm}^2$ GEM detector [85]

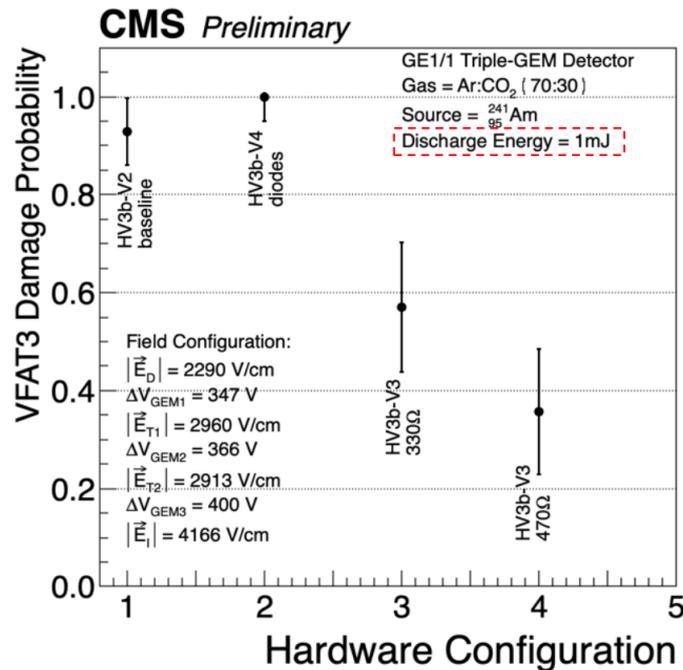


Figure 7.17 – VFAT3 damage probability vs hardware configuration for different protection circuits on a GE1/1 detector with 1 mJ discharge energy [85]

Thus, although the HV3b_v4 seemed better than HV3b_v3 in nearly all aspects, it was unsuitable for use in the CMS GEM systems, and HV3b_v3 was chosen as the ultimate on-chip protection solution.

7.3 HV Filter Protection

As described in Section 5.2, in larger detectors the complex propagation of discharges can result in the reignition of the propagation after it first propagates to the readout board. This requires a certain amount of stored energy, which is largely held within the HV filter of the detector, whose circuit diagram is seen in Figure 7.18a. Thus, we can limit the potential for propagation reignitions, and hence limit the damage that can be caused to the readout electronics and detector by any given discharge event, by tuning the capacitance of the HV filter. We do this by placing an additional resistor between the GEM electrode and the default 10 k Ω resistor on the HV filter, as shown in Figure 7.18b.

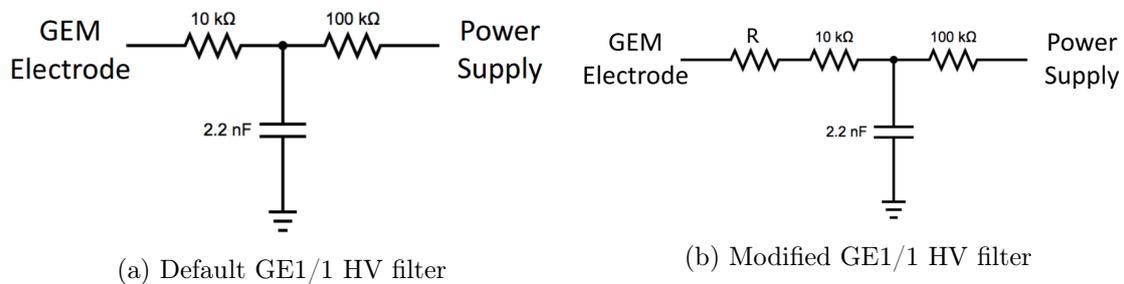


Figure 7.18 – GE1/1 HV filters

Multiple values for this additional resistor were tested, ranging from 100 k Ω to 1 M Ω . Figure 7.19 shows a representative discharge event given a 1 M Ω additional filter resistor, with values plotted as a moving average of 2, 5, and 8 data points for the antenna line, G3_{top}, and G3_{bottom} respectively. This plot, being a representative of all other plots taken at this filter resistance, shows no reignitions after a propagating discharge event and represents the drastic reduction of the probability of reignition.

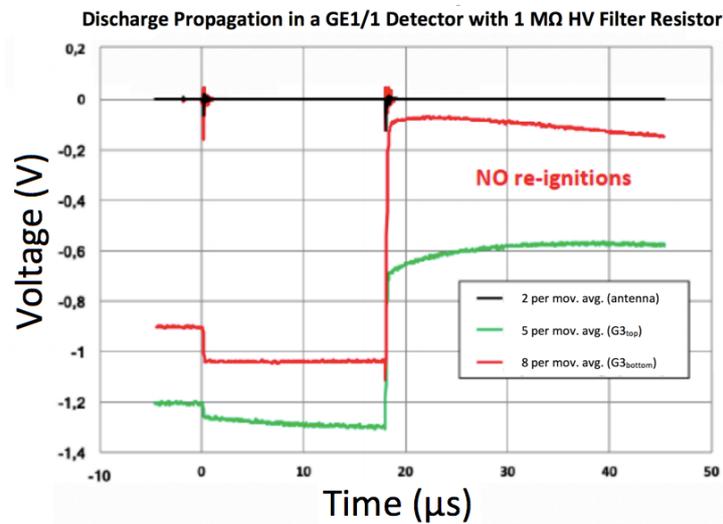


Figure 7.19 – Discharges in a GE1/1 detector with an additional 1 M Ω HV filter resistor [85]

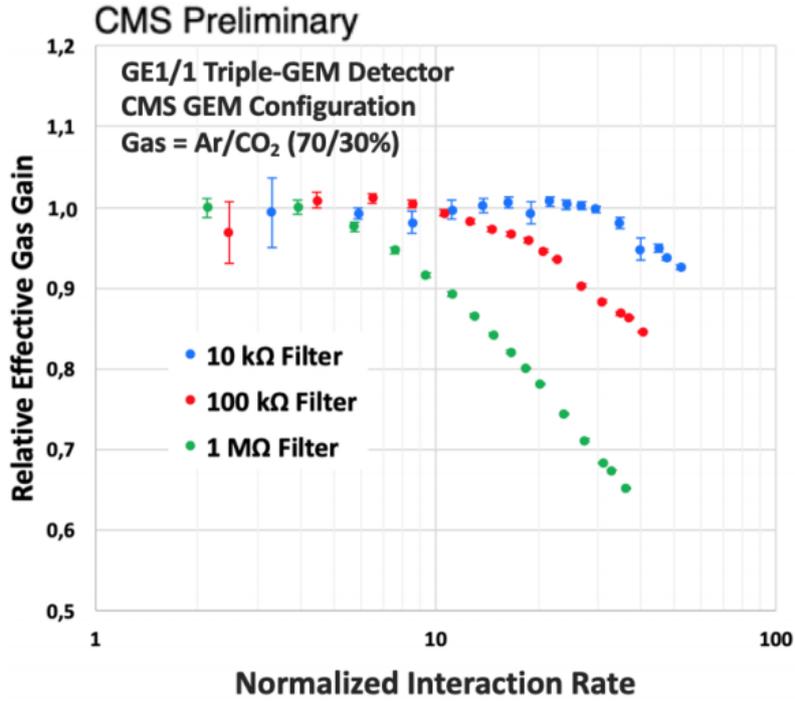


Figure 7.20 – Normalized interaction rates for different HV filter resistances [85]

However, this reduction comes at the cost of also negatively affecting the rate capabilities of the detector, as seen in Figure 7.20, where the relative effective gas gain drops more sharply with the normalized interaction rate as the filter resistance increases. Thus, a balance had to be reached between sufficiently reducing the probability of discharge propagation reignition without drastically reducing the rate capability of the detector. This was reached by ultimately choosing to add a 200 kΩ resistor on the GEM electrode side.

7.4 GE1/1 Final Protection Configuration

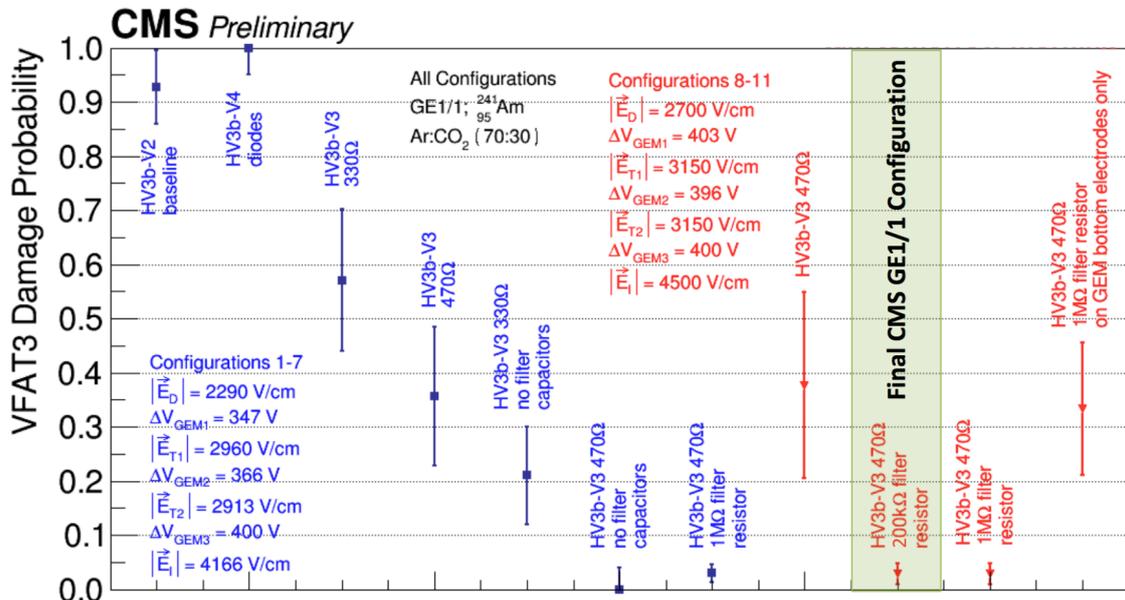


Figure 7.21 – VFAT3 damage probability for various protection methods and configurations

Figure 7.21 shows a summary of the tests described in Sections 7.2 and 7.3. Ultimately the final protection configuration for the GE1/1 detectors was chosen to be the HV3b_v3 VFAT with a series of resistors totalling to 470Ω and an additional $200 \text{ k}\Omega$ resistor added to the HV filter on the GEM electrode side. With this protection configuration, the probability of damage to the VFAT3 was reduced from $92.8\% \pm 6.8\%$ in the slice test detectors to $2.86\% \pm 1.99\%$.

Given this information, we can compare the damage that occurred during the slice test to the projected damage expected in the final GE1/1 system.

Taking the cause of the discharges within the detector to be HIPs, let's consider only those HIPs with a deposited energy of $E_{deposited} > 30 \text{ keV}$. The maximum rate of these particles in the area of the slice test and GE1/1 detectors is $R_{HIP,max} = 3.15 \cdot 10^3 \frac{\text{Hz}}{RO_{sector}}$. From the average of the tests performed by J. Merlin et al, which were measured using $E_{deposited} = 300 \text{ keV}$, we can take the discharge probability of a chamber to be $P_{dis} = 1.24 \cdot 10^{-9} \pm 1.16 \cdot 10^{-9}$. Then, the maximum discharge rate is $3.90 \cdot 10^{-6} \frac{\text{Hz}}{RO_{sector}}$. Assuming an LHC duty cycle of 0.33, this results in 39 discharges per year per sector.

Now, from the design of the GEM foils utilized in these detectors, we take the transfer propagation probability to be 0.5, and the induction propagation probability to be 1.

Finally, factoring in the probability of damage to the VFAT3 as given by the protection configurations in Figure 7.21, where $P_d = 0.928 \pm 0.068$ for the slice test detectors and $P_d = 0.0286 \pm 0.0199$ for GE1/1, we get the average expected values for channel loss over the 10-year lifetime of the system, as listed in Table 7.5.

System	Number of damaged readout channels per readout sector per 10 years	Percentage of damaged readout channels per readout sector per 10 years
Slice Test	185	100
GE1/1	6	4.79

Table 7.5 – Projected loss of readout channels per readout sector for GE1/1-style detectors within CMS [102]

Thus, we see the success and great importance of the changes made. Without the channel loss information gained from the slice test, and without the mitigation campaign that was sparked as a result, the GE1/1 system would have failed to meet the basic requirement of maintaining healthy operation within the CMS detector over its projected 10-year lifespan, reaching complete channel loss prior to those 10 years. With the new implemented changes, however, over those ten years we expect less than 5% of total channel loss under the harshest of conditions.

However, so far we have only considered the protection measures made to the GE1/1 detectors. And while this is a great accomplishment, there is a great deal that was not able to be considered given the time constraints that were present. So, we turn now to the future CMS GEM systems GE2/1 and ME0, and the mitigation measures we can consider with a much larger amount of time and design flexibility available to us.

Chapter 8

Future CMS GEM Systems: GE2/1 and ME0

Although GE1/1 is a monumental step forward towards a new era in CMS data taking, it is merely the first such step in a series of CMS GEM upgrades. As GE1/1 is in its final stages of production, installation, and commissioning, two additional GEM systems are in the process of being born: GE2/1 and ME0. The location of these new systems within the CMS muon endcap can be seen in Figure 8.1, with GE2/1 in red at $z \approx 8$ and ME0 in orange at $z \approx 5.5$.

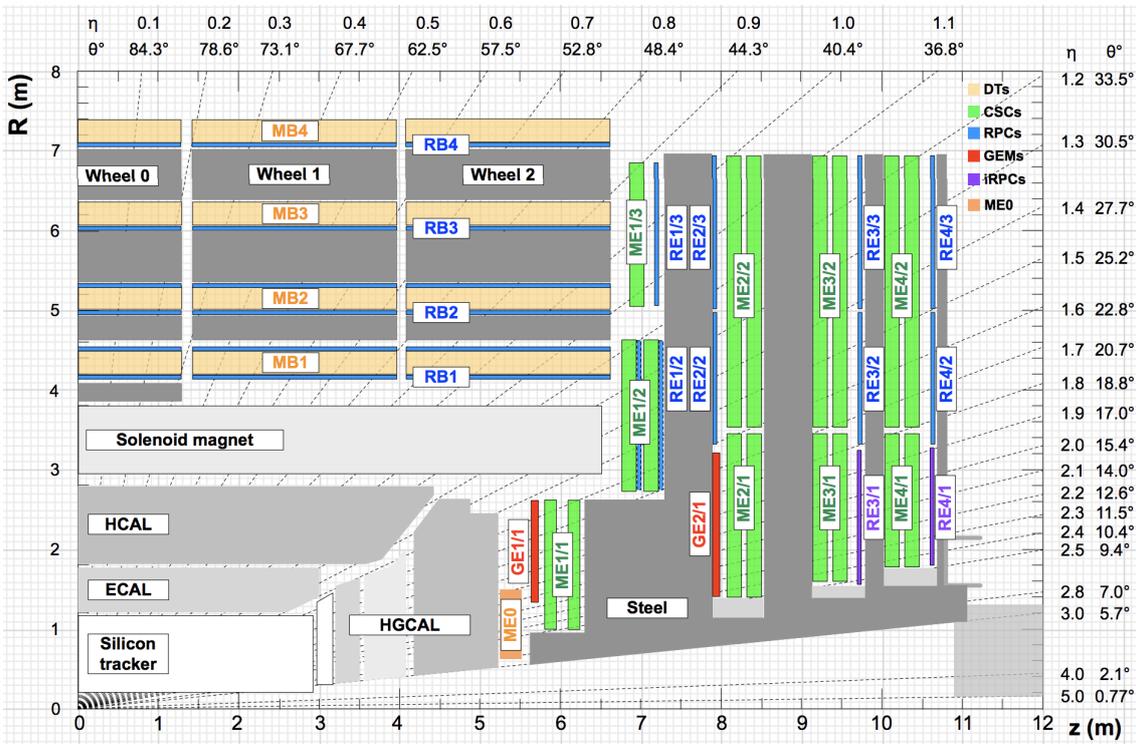


Figure 8.1 – Cross-sectional view of the CMS detector detailing the location of the three GEM-based systems within the CMS muon endcap

8.1 GE2/1

8.1.1 Motivation

As with GE1/1, the GE2/1 subsystem was designed with a complementary ideal in mind, positioned within the CMS muon endcap so as to complement the ME2/1 system of CSCs. They also complement the GE1/1 system, both overlapping its pseudorapidity range and extending the GEM detector acceptance to include the range of $2.15 < |\eta| < 2.4$, a range which was previously only covered by CSCs. With the inclusion of both the GE1/1 and GE2/1 systems, the number of hits for a traversing muon will be increased from six to

eight. This increase is essential for obtaining a robust track reconstruction at the Level 1 trigger, in addition to enhancing the overall redundancy of the muon endcap system.

At present, the measurement of momentum at the Level 1 trigger is based on both the known vertex position within the $r - \phi$ plane and the measured positions of the track segments in the CSCs. The addition of these new GEM chambers will increase the local lever arm for the segment reconstruction by greater than a factor of two. The additional local level arm is important for displaced vertices for which we cannot exploit the vertex constraint. By having two independent measurements of direction in stations 1 and 2 of the endcap, we will be able to make a standalone momentum measurement at the Level 1 trigger.

Since there are no iron yoke elements between the proposed GE2/1 and the existing ME2/1 chambers, the multiple scattering is small, allowing for the precise measurement of the muon direction within a single station. This is illustrated in Figure 8.2. Here, ϕ_1 and ϕ_2 are the muon local bending angles, which would be measurable at Level 1 with the addition of the GE1/1 and GE2/1 detectors. These measurements are crucial, as the measured direction of a given muon is strongly correlated with its momentum. By measuring this more precisely than in the current CSC-only system, we can greatly reduce both the fraction of soft muons with an overestimated momentum and the trigger rate for a given p_T threshold while still maintaining a high trigger efficiency.

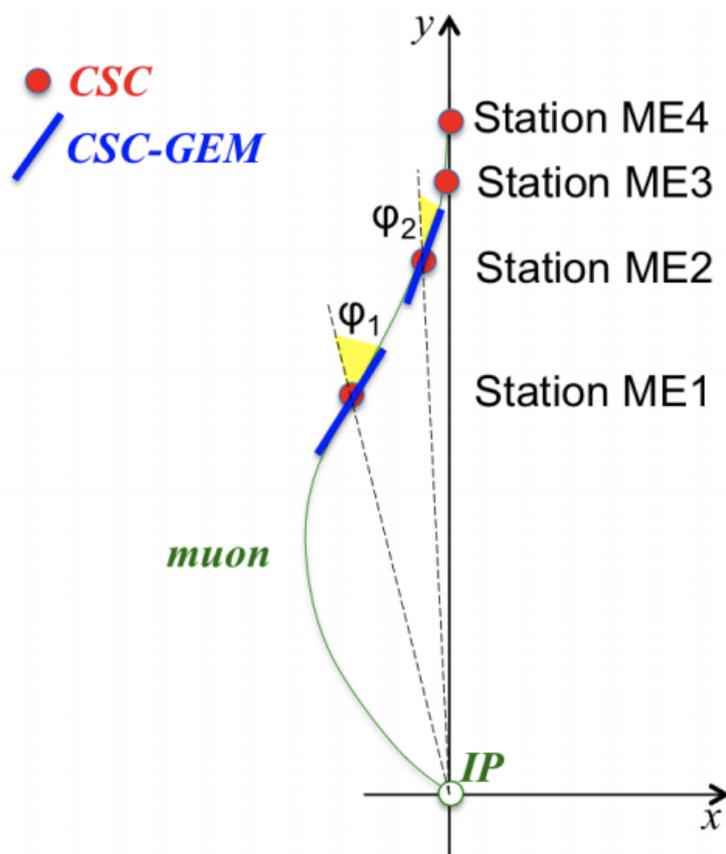


Figure 8.2 – A schematic of a muon trajectory through the GE2/1 and ME2/1 detectors from an axial point of view [8]

8.1.2 Detector Design Specifications

Despite the similarities in motivation between GE2/1 and GE1/1, the design of the GE2/1 detectors differs greatly from their predecessors. The basic outline for the design specifications can be found in Table 8.1, as taken from the technical design report for the Phase 2 upgrade of the CMS muon detectors.

Specification / Parameter	GE2/1
Detector technology	Micropattern gas detector (MPGD)
Charge amplification element	GEM foil (triple, cascaded, tensioned at $\sim 5 \frac{N}{cm}$)
Number of chambers in overall system	72 [36 (2×18) in each endcap]
Chamber shape (active readout area)	Trapezoidal, assembled from 4 modules M1 - M4; opening angle 20°
Active area overlap in adjacent chambers	15.2 mrad (32 readout strip pitches)
Chamber dimensions (active volume)	L: 183.3 cm (centerline), W: (53.3–117.4) cm, H: 0.72 cm
Total chamber thickness	H: 36.6 mm
Active readout area	1.45 m ²
Active chamber volume	10.4 liters
Radial distance from beam line	1365.5 mm (at inner edge of active readout area)
Geometric acceptance in $ \eta $	1.62 – 2.43
Signal readout structure	Truly radial readout strips
Readout strip dimensions	474 μ rad angular pitch
Number of η -segments in readout	8
Number of readout strips per η -segment	768
Number of readout strips per chamber	6144
Counting gas mixture	Ar:CO ₂ 70:30
Nominal operational gas flow	1 chamber volume per hour
Number of gas inlets	4 (one per module)
Number of gas outlets	4 (one per module)
Nominal HV applied to drift electrode	3200 V (Ar:CO ₂)
Nominal operational gas gain	$1 - 2 \times 10^4$
Demonstrated rate capability	few $\frac{MHz}{cm^2}$

Table 8.1 – Main specifications and parameters for the design and operation of GE2/1 [8]

The first thing of note in this table is that the number of chambers in the GE2/1 system is 72 - 18 two-detector superchambers in each endcap - or half the number of chambers in GE1/1. The reason for this is a simple one: the GE2/1 chambers are larger than GE1/1, spanning an opening angle of 20° rather than 10° . This difference in size is illustrated in Figure 8.3, which compares the rings of the two detector systems. Because the GE2/1 chambers are so much larger than their predecessors, they must, by necessity, be broken down into multiple modules, named M1 through M4 for the back layer of chambers (the chamber in each superchamber which is further from the interaction point) and M5 through

M8 for the front layer of chambers (the chamber in each superchamber which is closer to the interaction point).

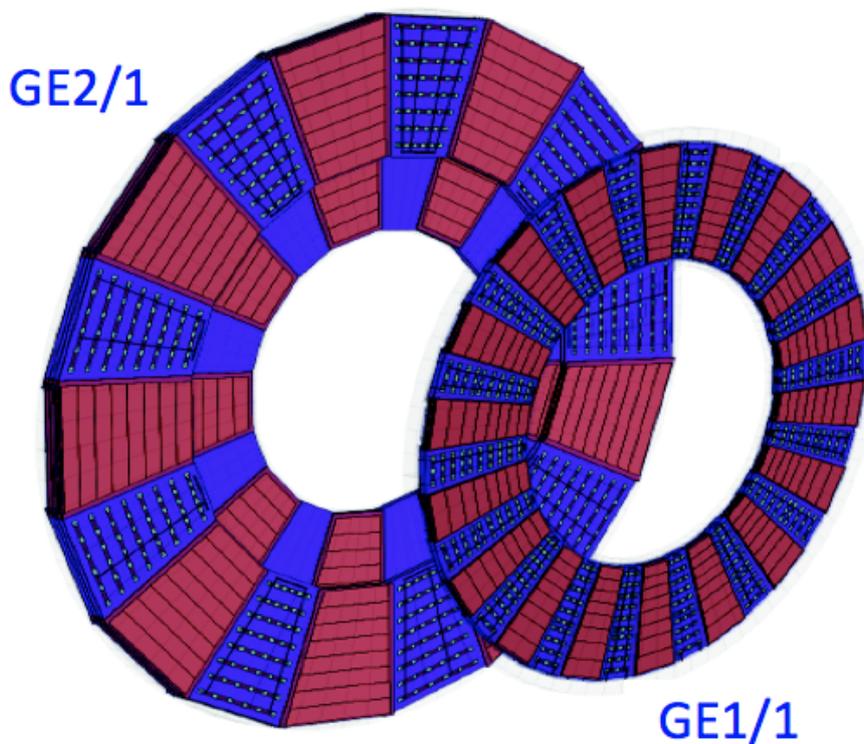


Figure 8.3 – The GE1/1 ring of detectors superimposed on the GE2/1 ring of detectors, illustrating their size differences

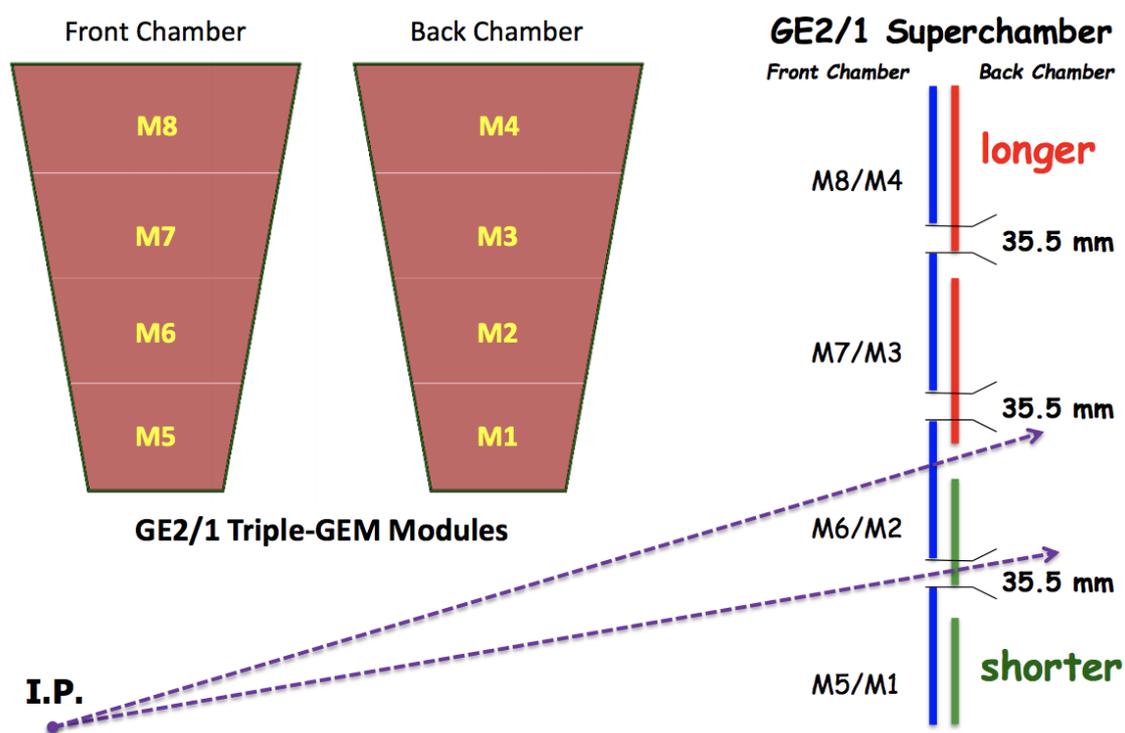


Figure 8.4 – The relative positioning and partitioning of the front and back GE2/1 detector layers [8]

While the breaking of a GE2/1 detector into four individual modules is a necessity for the ease and cost of production of the detectors, it comes with the unfortunate downside of introducing a dead area of 35.5 mm in between each module, for a total of 100.5 mm dead space in total per detector. In order to minimize the effect this has for CMS data taking, the individual modules for the front and back detector layers are of slightly different sizes, as shown in Figure 8.4. In this way, when projecting back to the interaction point, the dead space gaps for the two detector layers do not overlap, and thus do not create a gap in the GE2/1 acceptance.

In order to provide mechanical stability to the full detector, the modules are coupled together with an aluminum frame which runs around the edges of the full detector. Three aluminum stiffeners are placed between the modules, and an additional two radial aluminum bars are fixed onto both the frame and stiffeners and run below a protective aluminum chimney which encloses the detector. Figure 8.5 shows the first prototype of this structure, in which an M4 module is fixed into position.

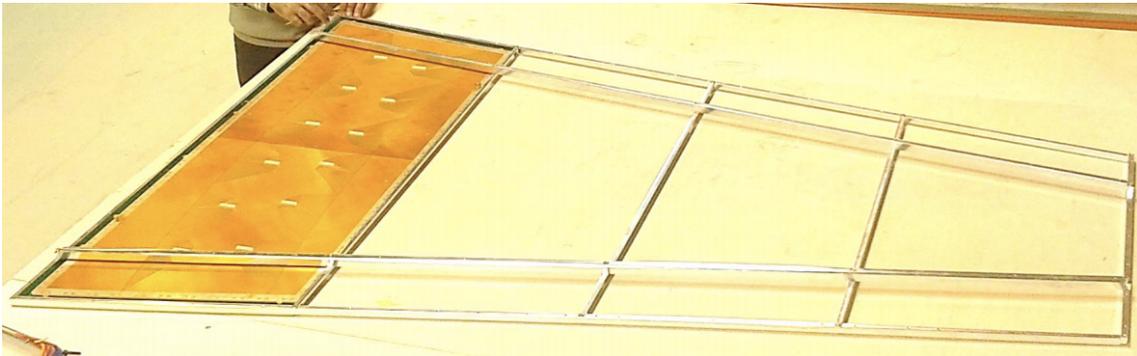


Figure 8.5 – First prototype of the mechanical support structure of a GE2/1 detector with M4 module in place [8]

8.1.3 Readout Sectors and VFAT3s

Each module is, in its own right, a complete GEM detector, consisting of a drift and readout PCB, external and internal frames, and a set of GEM foils specific to each module. Each module is segmented into two partitions along the η direction and 768 strips along the ϕ direction. Strips belonging to the same η partition are routed to the readout connectors in groups of 128 strips, as with GE1/1, in order to match the granularity of the front-end electronics, with 64 strips routed to either side of the readout connectors. This routing is shown in Figure 8.6 for an M1 module. Given this segmentation, each module is partitioned into 12 readout sectors ($2 \eta \times 6 \phi$) for a total of 48 readout sectors per GE2/1 detector.

Each module's front-end electronics consists of 12 VFAT3 chips, which are connected to the detector via a hybrid plug-in card. This plug-in card, shown in Figure 8.7, consists of a rigid portion, where the VFAT3 chip is located, and a flexible portion which serves to absorb any residual misalignment between the GEB and the readout board. The rigid portion is a key upgrade from the VFAT hybrids used in GE1/1, as its larger area allows for the use of larger, more robust protection circuits than those that were able to be considered in Section 7.2 due to the size constraints. The potential protection circuits currently under study will be discussed in much greater detail in Section 8.1.6.

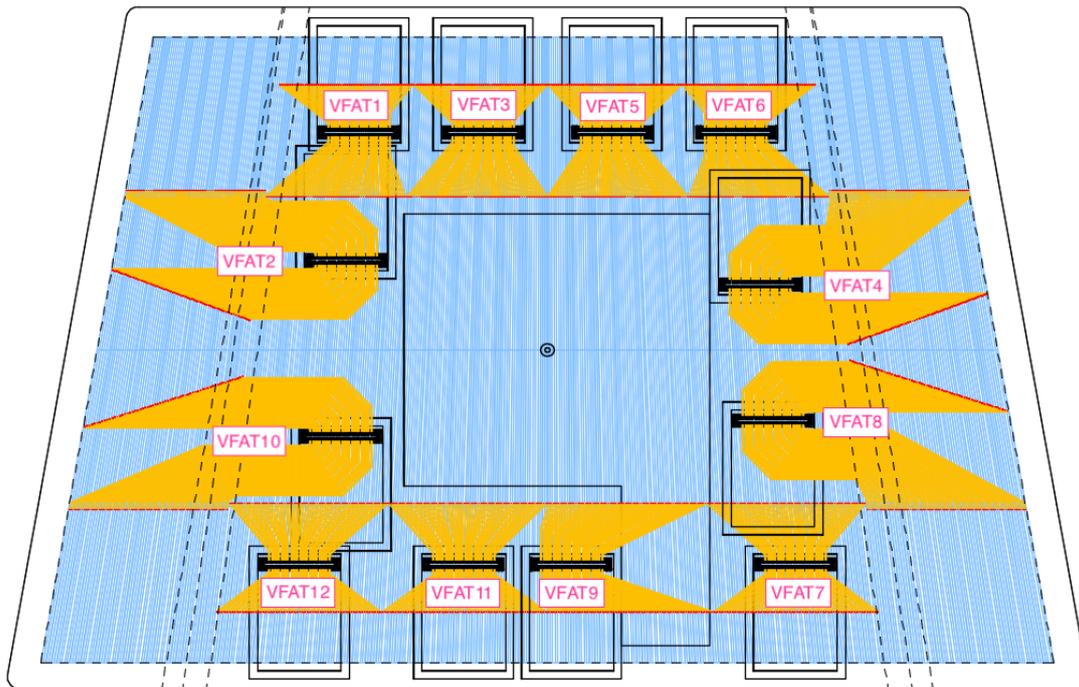
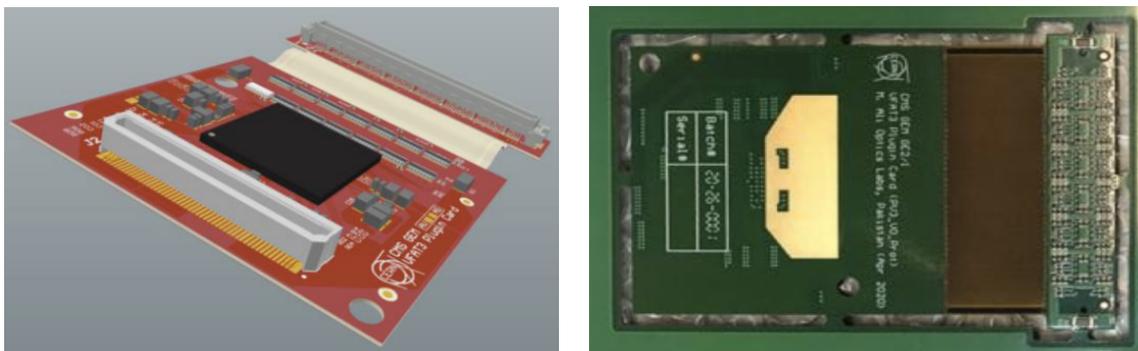


Figure 8.6 – Schematic of the GE2/1 M1 readout board, with VFAT positions labeled [103]



(a) Visualization of the GE2/1 plug-in card (b) First prototype of the GE2/1 plug-in card

Figure 8.7 – GE2/1 VFAT3 plug-in card versions [52]

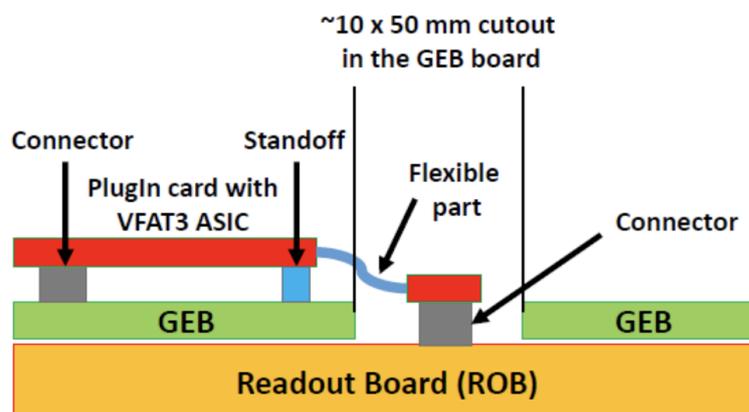


Figure 8.8 – Side view schematic of the GE2/1 plug-in card as installed on a detector [104]

8.1.4 GE2/1 Optohybrid

In addition to the 12 VFAT3 hybrids, each GE2/1 module contains an optohybrid, as shown in Figure 8.9, which provides readout and trigger interfaces for the VFAT3s. By giving each module its own optohybrid, the potential for mechanical stress that existed in GE1/1 by virtue of the optohybrid straddling two GEBs is eliminated, and independent powering and separate grounds can be maintained. By virtue of operating one GBTx in wide bus mode (See Section 8.1.4.1 for more detail), the trigger data can be sent to the μ TCA back-end through the GBT links rather than utilizing the VTTx optical links. As a result, each module only requires one VTTx in order to transmit the trigger data to the CSC OTMB, with the exception of the smallest, innermost module - due to the larger hit rate present here, this module alone has two VTTx links.

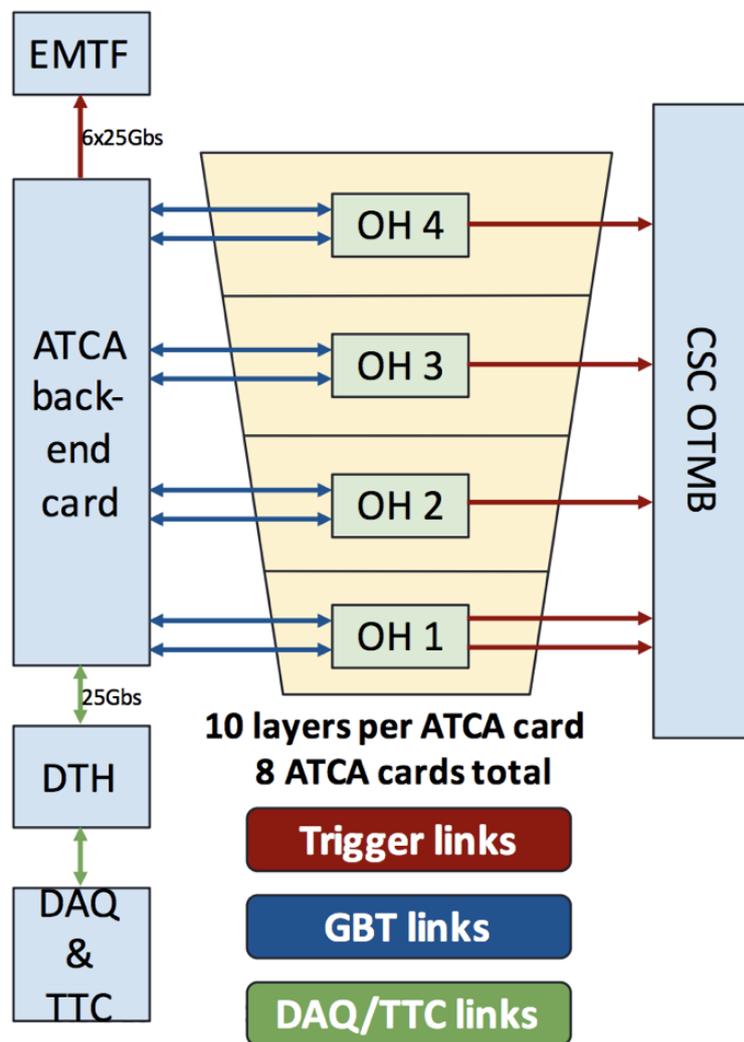


Figure 8.9 – Schematic of the GE2/1 electronics scheme

Prior to the decision to use the wide bus mode in this manner, the number of trigger optical links was reduced via the use of a master-slave configuration, with OH 1 and 2 and OH 3 and 4 connected via a simplex link. This simplex link allowed for direct FPGA-to-FPGA connection without intermediate buffers and had 12 differential data lines which ran at 320Mbps. In this configuration, while all four optohybrids had two VTRx optical transmitters, only the master optohybrids had two VTTx optical transmitters as well

[105]. While this master-slave configuration model ultimately went unused, it remains a possibility should a future upgrade be required to reach a higher trigger bandwidth.

Unlike the GE1/1 optohybrids, which feature a Virtex-6 FPGA, the GE2/1 optohybrids have been redesigned to utilize an Artix-7 200T FPGA. This selection was made for several reasons, including the presence of a sufficient amount of IO pins (285 total) and logic resources, a low power consumption, a flexible IODELAY block, sufficient radiation tolerance for long-term use in CMS, and a low cost and high availability [105].

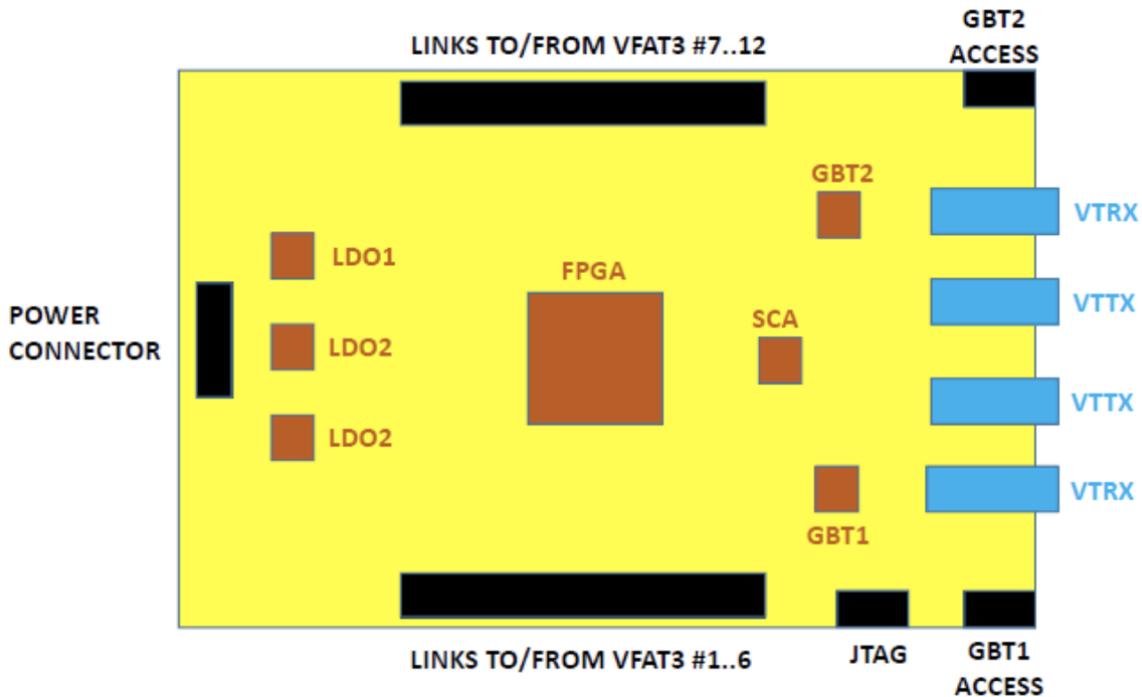


Figure 8.10 – Block diagram of the GE2/1 optohybrid, master configuration [105]

Also unlike the GE1/1 optohybrids, which utilized three GBTx ASICs, the GE2/1 optohybrid uses only two. The first of these GBTx serves VFATs 1 through 6, and the second serves VFATs 7 through 12, which are connected to the optohybrid through links on opposite sides of the board, as shown in Figure 8.10. Each of these links comprises three differential signal pairs: the outputs from the GBTx to VFAT3, the inputs from VFAT3 to GBTx, and the clock outputs from GBTx to VFAT3. The programming of each GBTx is possible through a dedicated connector (8-pin FCI75867-132F) from an I2C dongle. Each GBTx ASIC is connected to its own VTRx optical transceiver, and provides two general purpose e-links to the FPGA. GBTx1 also provides e-links to download the FPGA, supports the SCA interface, and provides four global clocks to the FPGA - two dedicated MGT clocks and two general purpose global clocks [105].

As with GE1/1, in order to remove any dependence on radiation-soft components, promless programming of the FPGA was required, in which the back-end system programs the FPGA through the GBTx links described above. On the optohybrid, the reset of the FPGA is controlled by the SCA chip, and the FPGA's parallel programming port is connected to eight GBTx e-links, each of which runs at 80 Mbps. Once a TTC hard reset signal is received by the back-end board, it must reset all of the optohybrid FPGAs and send a bitstream to all of the optohybrid FPGAs in parallel. In order to save block RAM resources, the back-end board must have storage outside of the FPGA which is larger than

the bitstream. For the Artix-7 200T FPGA, this bitstream is 3.6 MB, but in order to allow for the potential future upgrade of the FPGA in the final system, a storage requirement of 9.2 MB has been taken into account [106].

This storage must also take into account that the sending of the bitstream must be done without interruption: for every 40 MHz clock cycle, the back-end must deliver two bytes of the bitstream until its end is reached. Thus, the back-end bitstream storage must sustain a constant data rate of at least 640 Mbps, as shown in Table 8.2.

8.1.4.1 VTRx and VTTx Optical Links

As mentioned above, each optohybrid contains two duplex optical links based on the CERN-designed radiation-hard VTRx transceiver, shown as the blue double-arrows in Figure 8.9. These links are used for communication with the 12 VFAT3 chips, SCA chip, and optohybrid FPGA. For the VFAT3 communication, this includes delivery of the VFAT3 control data, which is comprised of the clock, TTC, and slow control signals, to the VFAT3 chips, as well as reception of DAQ and slow control data from the chips. For the SCA communication, this includes delivery of the SCA control data and reception of the SCA reply data in order to control the GPIO, ADC, I2C, and JTAG modules. And for the FPGA communication, this includes the remote programming of the FPGA using 8 bit / 80 MHz slave SelectMAP mode, the delivery of clock, TTC, and slow control signals, the reception of slow control signals, and the reception of S-Bit cluster trigger data [106].

Each master-configured optohybrid, but not the slave optohybrids, contains two other optical links, which are based on the CERN-designed two-channel VTTx transmitter. These are shown as the red arrows in Figure 8.9. Here, two channels are required in order to provide trigger bits to the OTMB board for local triggering, and the other two channels provide that same data to the CMS EMTF. Both transmission types occur at 3.2 Gbps or higher [105].

A standard feature of the GBTx is “wide bus mode”, which converts four output e-links to input e-links, each of which run at 320 Mbps. These output e-links were previously unused, making this conversion virtually cost-free. Operating the GBTx in wide bus mode provides 3.52 Gbps of unused bandwidth from the optohybrid to the back-end electronics, which is more than enough to fit the S-Bit cluster of trigger data, given the need to transfer up to five clusters per bunch crossing for a total of 2.6 Gbps. The breakdown of the bandwidth needs for both the back-end Tx and Rx is provided in Table 8.2

	12 VFAT3s	SCA	FPGA Program- ming	FPGA TTC and Slow Control	Trigger Data	Total
Back-end Tx	3.84 Gbps	80 Mbps	640 Mbps	320 Mbps	N/A	4.88 Gbps
Back-end Rx	3.84 Gbps	80 Mbps	N/A	320 Mbps	3.52 Gbps	7.76 Gbps

Table 8.2 – Bandwidth breakdown of a single GE2/1 optohybrid board’s interface [106]

In the optohybrid-to-back-end direction, the forward error correction (FEC) code will also be removed and replaced with user data. The extra bandwidth can then be used for

error detection codes (CRC) and provide ample margin for possible future needs. CRC was previously used in GE1/1 for error detection in the DAQ data, though not in the trigger data as it will be for GE2/1, and the CSCs of CMS have been successfully operating with CRCs and no FEC for many years. The CSCs will also be adding a new board which uses GBTx in wide bus mode for future operations [107].

The embedding of the trigger data into the GBTx links in this manner marks a large deviation from GE1/1, and doing so allows the back-end electronics to be fit into just eight ATCA cards as opposed to the twelve CTP7s needed for GE1/1 [107].

8.1.5 Back-End Electronics

Where the back-end of GE1/1 was based on the μ TCA standard, GE2/1 instead uses the advanced telecommunications computing architecture (ATCA) standard. As of the writing of this dissertation, the specific ATCA card has not been chosen. However, a baseline option has been identified in the APx Consortium's Advanced Processor - Trigger (APT) card. The architecture diagram and a photo of a card can be seen on the left and right of Figure 8.11, respectively.

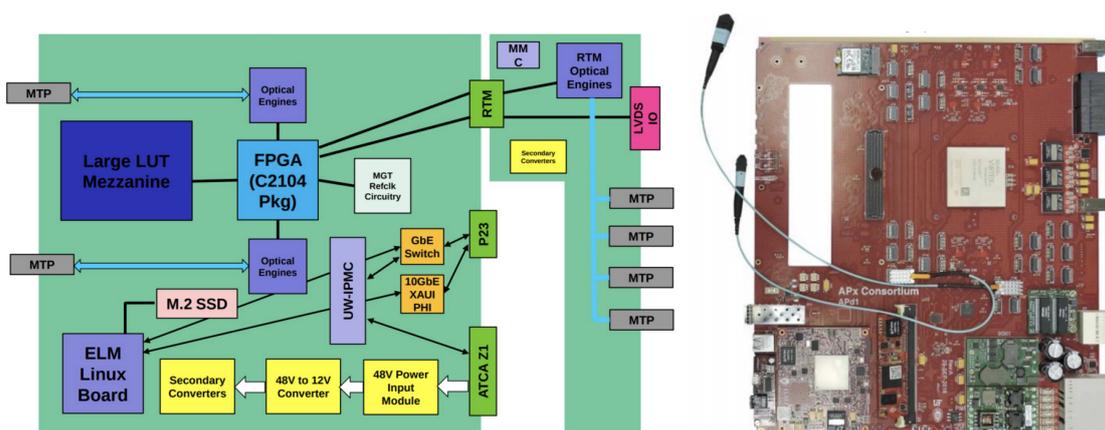


Figure 8.11 – Architecture of the APx APT card (left) and photo of a prototype of the physical card (right) [107]

This baseline option, which already satisfies all of the needs of GE2/1, ME0, and the CSC upgrades, contains onboard Linux like the CTP7, but has a faster CPU and faster link to the FPGA than the CTP7. It utilizes a large Virtex Ultrascale+ VU9P FPGA and supports up to one hundred 28 Gbps optical Tx and Rx links [107].

8.1.5.1 EMTF Interface

After the trigger data is received from the optohybrid, the back-end electronics multiplex that data and retransmit it to the EMTF on 25 Gbps unidirectional optical links, where one 25 Gbps link is needed for a single GE2/1 superchamber's worth of trigger data. This is seen at the top left of Figure 8.9.

The EMTF itself is segmented into six trigger sectors per endcap (12 total), each of which covers 60° in ϕ , plus an additional $10^\circ - 20^\circ$ of overlap from its neighboring sector. As each GE2/1 superchamber covers 20° in ϕ , each EMTF trigger sector is sent trigger data from four superchambers. This means that interfacing with each trigger sector requires four links, with one link dedicated to overlap, for a total of 48 links [106].

8.1.5.2 DTH Interface

The ATCA crate also hosts the DAQ and timing hub (DTH) board, which replaces the functionality that the AMC13 held for GE1/1. A 2019 prototype of the DTH board can be seen in Figure 8.12.

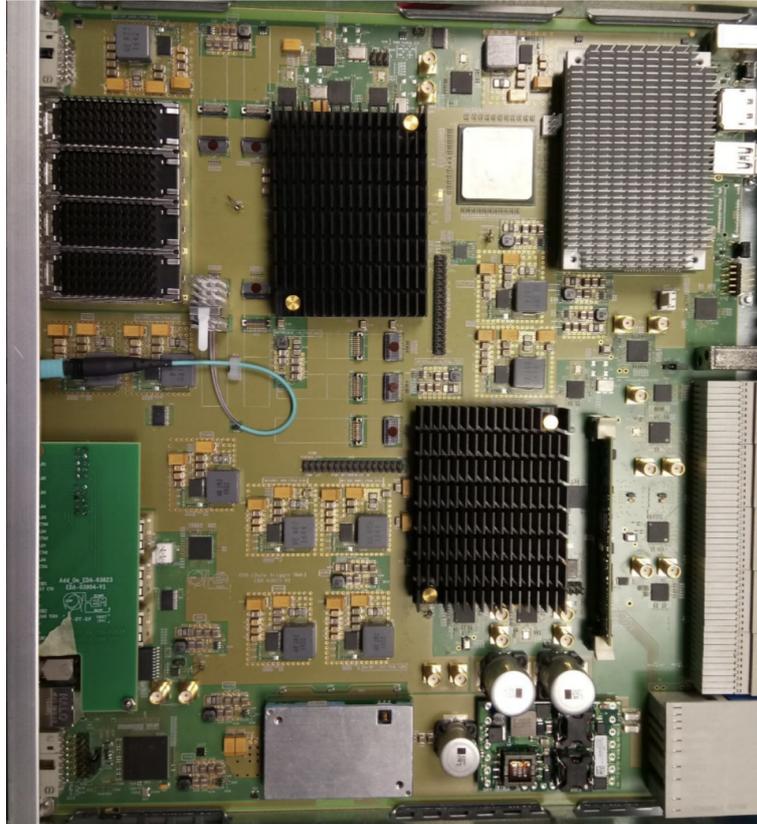


Figure 8.12 – 2019 prototype of the DTH to be used in GE2/1 and ME0 [107]

The DTH acts as an interface between the GE2/1 back-end system and both the TCDS and DAQ system, and is installed into the hub slot of the ATCA crate. From here, it provides both the LHC clock and TCDS commands to all of the GE2/1 back-end boards via the ATCA backplane. The DTH board can also be fitted with up to three DAQ modules which add an additional 200 Gbps of bandwidth to the DAQ system. However, despite this capacity, the total average DAQ data rate of the GE2/1 system is expected to be 6.64 Gbps in HL-LHC conditions, which is well below the capabilities of the DTH board, and as such, only a single DAQ module will be needed for a single DTH board [106].

Communication between the DTH and the ATCA back-end card, and the DTH and the DAQ module, is accomplished via a single 25 Gbps bidirectional optical link each, as shown in green in Figure 8.6.

8.1.6 Protection Circuits for the GE2/1 VFAT3 Plug-In Cards

As stated in Section 8.1.3, the new VFAT3 plug-in card has a greater area than the VFAT3 hybrid used in GE1/1, and this greater area allows for larger, more robust protection circuits to be used. As the studies following the slice test focused solely on those protection circuits which could physically fit on the VFAT3 hybrid, further studies were needed for GE2/1 to explore other options which could not be considered for GE1/1 for size reasons

alone.

In order to make the most complete study possible, three types of protection circuits were considered: resistor-only (R), resistor and capacitor (R+C), and resistor and diode (R+D). These three options are shown in the circuit diagram in Figure 8.13, with each option highlighted in a separate color. A R&D plug-in card was produced for the study which had the physical space for all three circuit types - for each study performed, the unused circuits were replaced with jumpers in order to be able to use the same hardware.

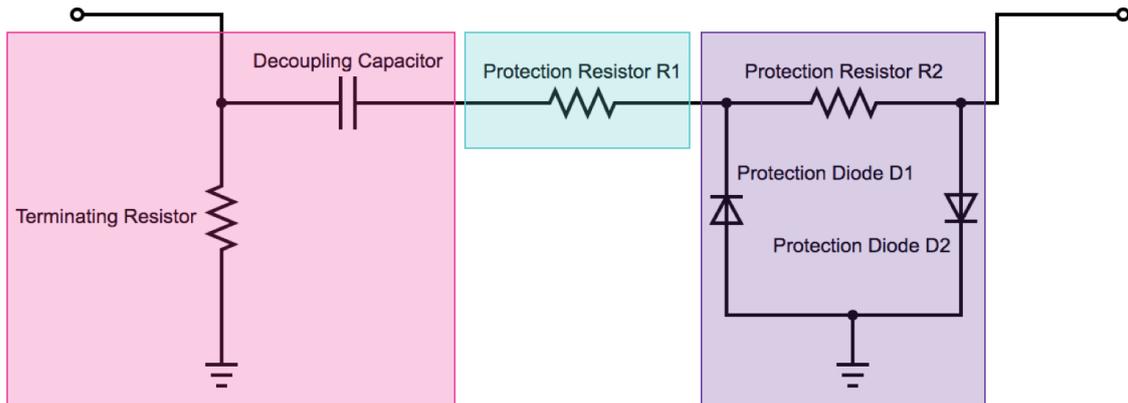


Figure 8.13 – Circuit diagram of the protection circuit types considered for GE2/1

Each circuit type was chosen with a specific goal in mind. The R circuit type, shown highlighted in turquoise, was chosen to mimic the resistor-only protection circuit of the HV3b_v3 VFAT3 hybrid of GE1/1. This circuit would allow for the testing of larger, more robust resistors without changing the type of protection used between systems. It would also allow for the opportunity to test the resistor-only protection under much harsher conditions than were able to be tested for prior to the finalization of the GE1/1 electronics.

The R+D circuit, shown highlighted in purple, was chosen so as to revisit the ideas behind the HV3b_v4 protection circuit that was ruled out for GE1/1. This circuit would prevent damage to the inputs of the VFAT3 chip in the event of a discharge propagation.

Finally, the R+C circuit, shown highlighted in pink, was chosen so as to prevent a discharge from propagating in the first place. Unlike the resistor and diode circuit, however, if this prevention failed and a discharge did propagate to the readout electronics, this circuit would not prevent damage from occurring.

As this last option is the only circuit type which was not previously evaluated for use in GE1/1, special care would need to be taken to ensure its suitability. In particular, it would need to be proven that the addition of an R+C circuit to the VFAT3 plug-in card would not negatively affect the efficiency of the detector.

In order to do so, the narrow end of a GE1/1 detector was populated with HV3b_v2 VFAT3 hybrids connected to the R&D plug-in card, as shown in Figure 8.14. This allowed a production-grade VFAT3 hybrid to be tested without the interference of any previously-installed protection circuit, as HV3b_v2 is the baseline GE1/1 VFAT3 hybrid. A pair of scintillators was installed over the center of the VFAT cluster, in order to provide a trigger signal via a coincidence from cosmic muons. Due to the large jitter of the scintillator coincidence scheme, a $5 \mu\text{s}$ coincidence window was chosen, along with a low threshold setting. The GEM acquisition window was chosen to be 200 ns so as to capture a maximal number of events.

Using knowledge gained from prior tests with GE2/1 detectors, this test detector was outfitted with an improved grounding scheme. A connection between the AGND pins on the VFAT3 hybrids and the readout board's copper layer was added by soldering $0\ \Omega$ resistors at positions R10 and R11 on the VFAT3 hybrid boards and then connecting the corresponding pads on the plug-in card with a small wire.

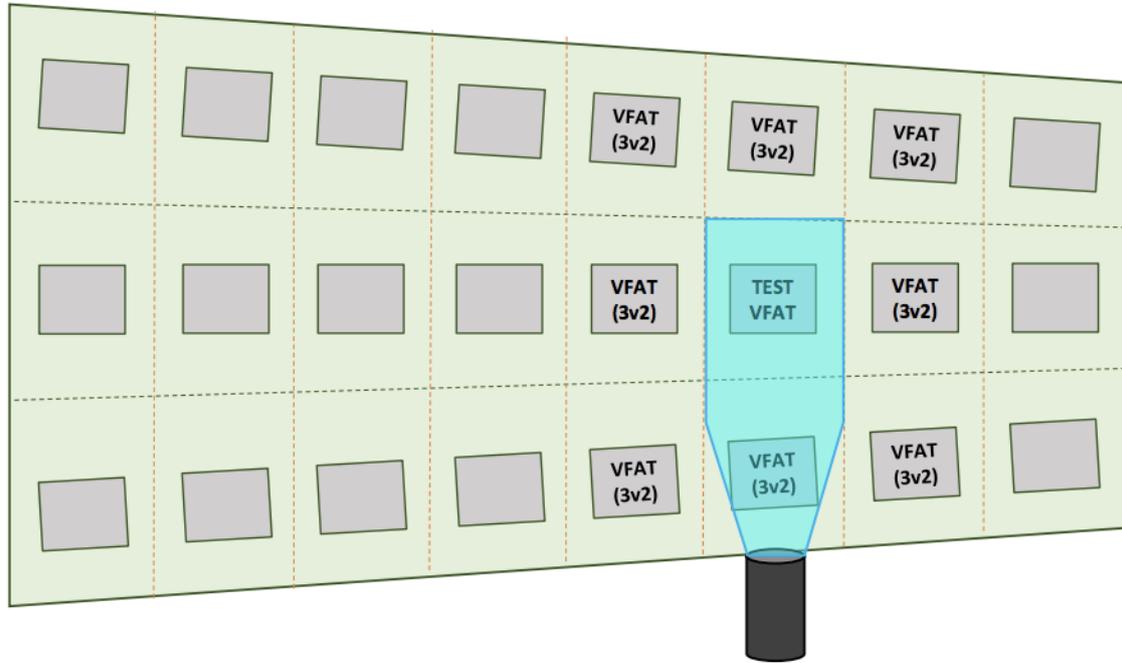


Figure 8.14 – Schematic of the GE2/1 protection circuit efficiency study setup

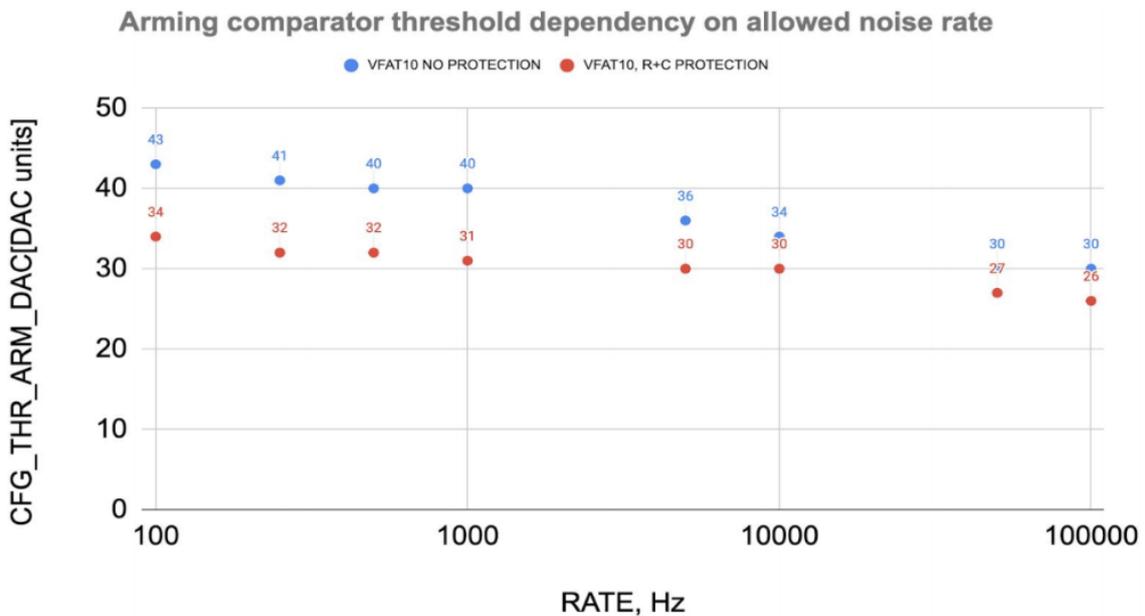


Figure 8.15 – CFG_THR_ARM_DAC values for VFAT10 of the GE2/1 protection circuit test set-up at different noise rates

Prior to data taking, an electronic noise scan was performed in order to obtain the set of comparator threshold (CFG_THR_ARM_DAC) values which allow different levels of noise. This would then allow for the efficiency of the detector to be determined at varying

allowed noise levels. These threshold values both with and without the presence of the resistor and capacitor protection circuit are shown in Figure 8.15, with the protectionless data in blue and the protection data in red. Note that these values come only from VFAT10, the VFAT directly under the area covered by the scintillator paddles, labeled in Figure 8.14 as “TEST VFAT”.

One other crucial piece of information to note is that here, the improved grounding scheme described above was implemented prior to taking the data points with the protection circuit, but after taking the data points without the protection circuit. The protectionless data was taken with the detector still using the traditional “star point” grounding scheme from GE1/1. This discrepancy may contribute to the higher values of CFG_THR_ARM_DAC for the protectionless VFAT3s.

The efficiency of the detector with the protection circuits applied was then calculated as the ratio between the total number of triggered events and the number of events with at least one verified hit from a cosmic muon. This calculation was performed for a set of data obtained at different noise rates, and is presented in Figure 8.16.

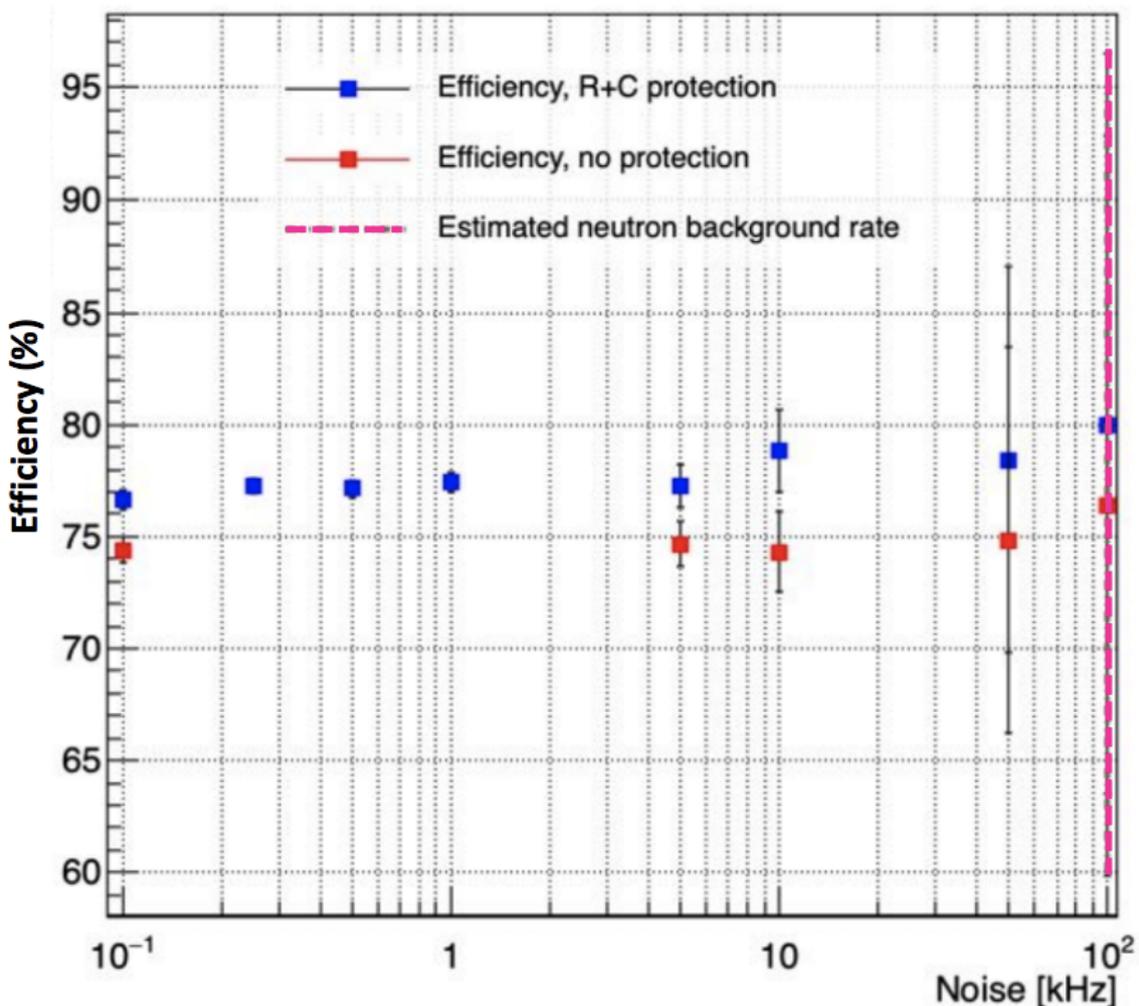


Figure 8.16 – Efficiency of the test detector at different noise rates, with and without a protection circuit applied on the plug-in card

Here, it is crucial to note is that the detector used in this test was an R&D detector which had not formally undergone all of the usual quality control tests required of a production-grade detector. As such, the baseline efficiency of the detector is far lower than

would normally be acceptable, at 75% - 80% rather than the required 97% for a detector to be validated for inclusion into CMS. This is why the inclusion of the non-protected data set is key, because it allows the protected data set to be evaluated in comparison to that baseline rather than the usual 97% efficiency baseline requirement.

For both data sets, the error is calculated as the estimated systematic error added in quadrature to the statistical error. The systematic error is derived from the electronic noise present in the system, where the probability of a noise signal occurring during the acquisition window is assumed to come entirely from the nine VFATs under the scintillator paddle and estimated from Poisson statistics via Equations 8.1 and 8.2. Here, $t_{acq} = 200$ ns, $N(VFAT) = 9$, and R_{VFAT} varied from 100 Hz to 100 kHz in different experimental trials. The probability function, P , is plotted as a percentage in Figure 8.17 as a function of the noise rate. The range of the x-axis was chosen as this source systematic error doesn't play a significant role until the allowed noise rate surpasses 50 kHz, as evidenced by the error bars shown in Figure 8.16.

$$P(\text{coincidence}) = 1 - e^{-R(t_{acq})} \quad (8.1)$$

$$R = N(VFAT) \times R_{VFAT} \quad (8.2)$$

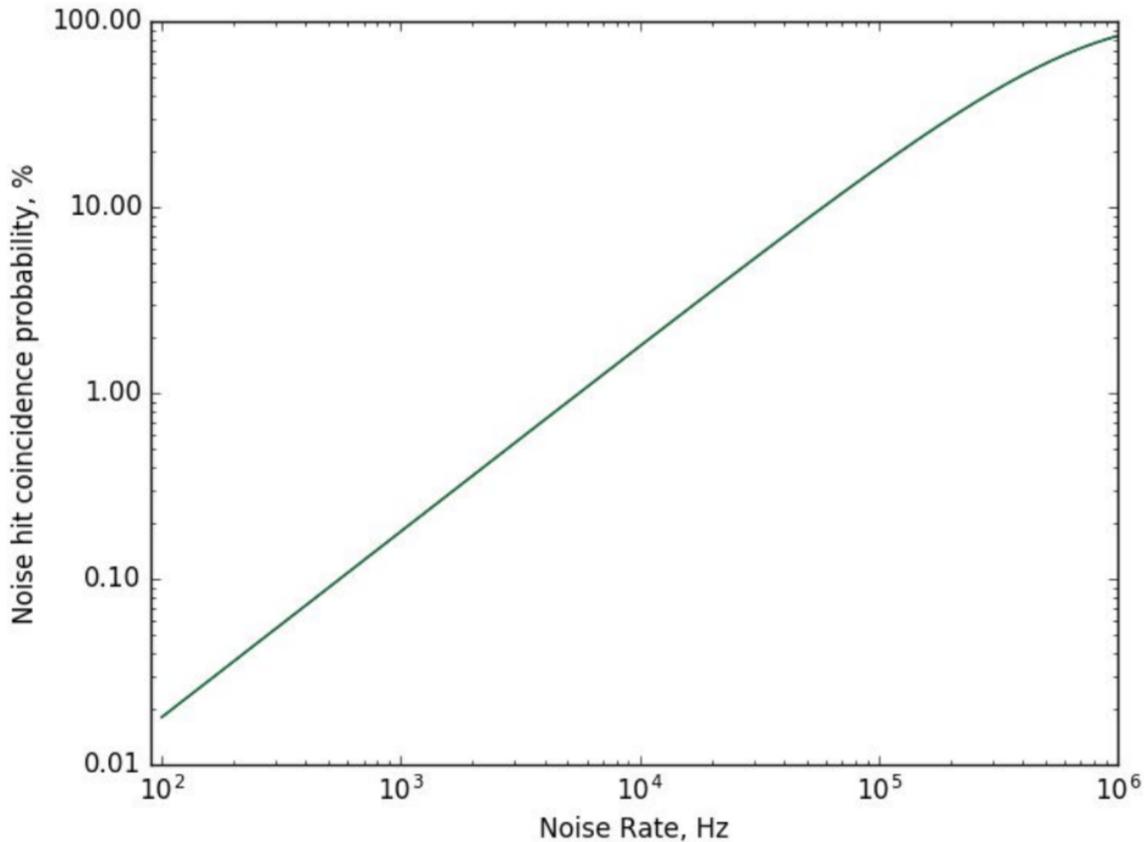


Figure 8.17 – Probability of a noise hit during the acquisition window as a function of the noise rate

Looking back to Figure 8.16, we see that until this noise becomes a significant factor in the efficiency measurement, the efficiency of the detector with the protection circuit applied remains nearly constant, indicating a stable mode of operation.

Furthermore, when Figure 8.16 is replotted as a ratio between the detector efficiencies with and without the protection circuit applied, as in Figure 8.18, we see that in all cases, the ratio remains above 1, indicating that this R+C protection circuit is, indeed, a viable option for GE2/1's protection scheme.

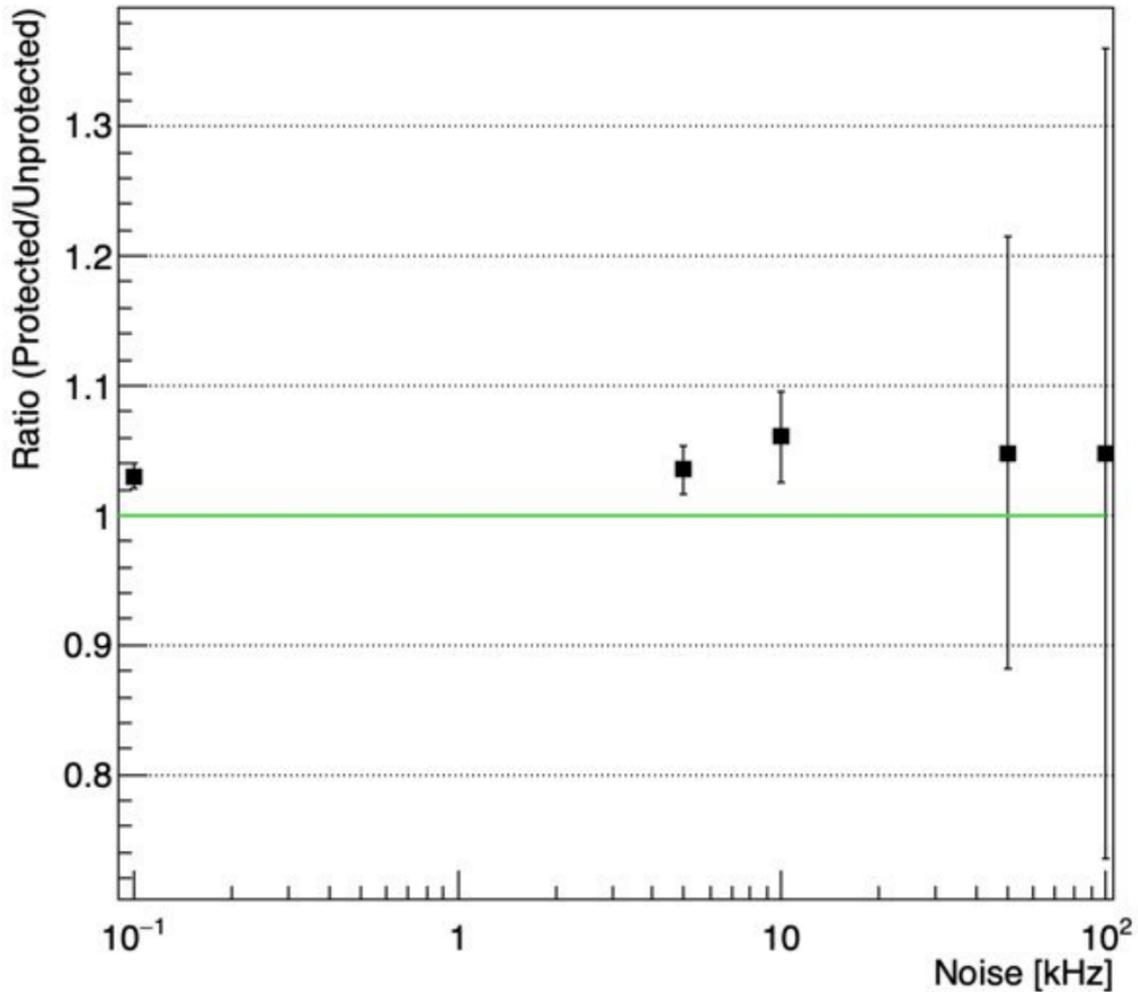


Figure 8.18 – Ratio of the efficiency of the test detector at different noise rates between the cases of being with and without a protection circuit applied on the plug-in card

Of course, further investigation is needed in order to verify this viability. In particular, further studies would be needed in the following areas:

- Testing of the protection circuit vs. the bare detector when the detector is in the same grounding configuration for both cases, for a truly comparable ratio between the two.
- Use of a production-grade GE2/1 detector for the tests to ensure that the efficiency with the protection circuit in place remains above 97%.
- Testing of various types of resistors and capacitors for this protection circuit type
- Testing and/or simulation of the detector with an R+C circuit at the expected interaction rate within CMS.
- Testing of the other protection circuit types described above, including R and R+D.

These investigations will be carried out after the completion of the GE1/1 installation and the lifting of the more stringent COVID-19 restrictions that are in place as of the writing of this dissertation. The fourth test is of particular importance, as the addition of the R+C circuit makes the detector AC-coupled rather than DC-coupled. And although this is not a problem for cosmic muons, and is generally expected to not be a problem for the rate expected in GE2/1, the expected ME0 interaction rate of a few 100 kHz per strip could cause this switch to be an issue, invalidating this option for ME0's protection.

At present, 12 prototype plug-in cards have been manufactured with the R+C circuit of values $R = 100 \text{ k}\Omega$ and $C = 1 \text{ nF}$. The resistors have been placed on the smaller of the plug-in card, the so-called daughter board, whereas the capacitors have been placed on the larger side of the plug-in card where the VFAT lies, the so-called mother board. The reasons behind choosing the 1 nF value of the capacitor were two-fold: for one, the footprint of the 1 nF capacitor is the largest that can be physically accommodated on the plug-in card. And for another, the resistor of the R+C circuit acts as both a drain resistor and a “pull-down” resistor, in that during a discharge it collapses the electric field of the induction gap. As GEM3's operational voltage is (generally) 412 V, the protection circuit must also be functional at this voltage. Previously a 3.8 nF capacitor had been considered, but its max operational voltage was only 500 V compared to the 630 V of the 1 nF capacitor, and this was deemed to be too close to the 412 V of GEM3 to ensure safe operation. These prototype cards will be used in the validation of the GE2/1 double-segmentation re-designs, as will be discussed in Chapter 9.

8.2 ME0

8.2.1 Motivation

Unlike GE1/1 and GE2/1, which were designed with the ideal of complementing the previously-installed CSC systems at those locations, ME0 is intended to fill an entirely new niche within the CMS muon endcap. As the very forward muon tagger, ME0 will extend the acceptance of the CMS endcap muon system from $|\eta| = 2.4$ with its range of $2.0 < |\eta| < 2.8$, the maximum possible range allowed by mechanical constraints. This extension into $2.4 < |\eta| < 2.8$ is beyond the reach of any other CMS muon detector, making ME0 a key addition to the CMS muon system, particularly in the detection of multi-muon final states such as $H \rightarrow 4\mu$ and for forward particle production [8].

In the region of $2.15 < |\eta| < 2.4$, which will be covered by CSCs, RPCs, and GE2/1, ME0 will help to maintain the muon identification and trigger capabilities even at high background rates. The additional track hits that ME0 provides increases the lever arm for momentum determination, in a very similar way as to what we saw in Figure 8.2 for GE2/1, suppressing the L1 trigger rate due to mismeasured low p_T muons by an order of magnitude within this η region. Together with ME1/1, ME0 will provide a second muon direction measurement (the first being provided by GE2/1 and ME2/1), thus enabling the standalone triggering on muons which result from the decays of long-lived particles.

In the region of $2.4 < |\eta| < 2.8$, the so-called “very-forward region”, ME0 will be used to provide a new trigger signal. Unlike GE1/1 and GE2/1, whose superchambers consist of only two layers of detectors, an ME0 superchamber or “stack” is composed of six GEM detectors. This provides up to six track points, which provides enough redundancy to reject neutron-induced backgrounds and form muon track segments in the L1 trigger. Although a standalone ME0-based L1 trigger is not feasible, L1 cross-triggers and the HLT will

make use of the ME0 trigger data, particularly for low p_T very-forward muons. In the HLT, ME0 allows for the performance of track reconstruction regionally, which enables the use of very-forward-going muons in the trigger algorithms at nearly the same quality as offline analysis. Thanks to the low background contamination ME0 can provide, even for transverse momenta as low as $p_T = 3$ GeV, the identification efficiency has been simulated to be approximately 95 % [8].

That being said, ME0 does still play a complementary role, though not in the same way as the GEM-CSC couplings. ME0's extended range of $2.4 < |\eta| < 2.8$ is covered by the inner tracker which is being added for Phase II. This tracker cannot identify charged particles as muons, but the GEM detectors of ME0 can. Therefore ME0 is needed to make full use of the new tracker's capabilities.

8.2.2 Detector Design Specifications

The basic outline for the design specifications of ME0 can be found in Table 8.3, as taken from the technical design report for the Phase II upgrade of the CMS muon detectors.

Specification / Parameter	ME0
Detector technology	Micropattern gas detector (MPGD)
Charge amplification element	GEM foil (triple, cascaded, tensioned at $\sim 5 \frac{N}{cm}$)
Number of chambers in overall system	216 [36 (6 \times 18) in each endcap]
Chamber shape (active readout area)	Trapezoidal; opening angle 20°
Chamber dimensions (active volume)	L: 78.8 cm (centerline), W: (23.6–51.4) cm, H: 1.8 cm
Total chamber thickness	H: 33.4 mm
Active readout area	0.296 m ²
Active module volume	2.1 liters
Radial distance from beam line	630 mm (at inner edge of active readout area)
Geometric acceptance in $ \eta $	2.03 - 2.8
Signal readout structure	Truly radial readout strips
Readout strip dimensions	940 μ rad angular pitch
Number of η -segments in readout	8
Number of readout strips per η -segment	384
Number of readout strips per module	3072
Counting gas mixture	Ar:CO ₂ 70:30
Nominal operational gas flow	1 module volume per hour
Number of gas inlets	4 (one per module)
Number of gas outlets	4 (one per module)
Nominal HV applied to drift electrode	3200 V (Ar:CO ₂)
Nominal operational gas gain	$1 - 2 \times 10^4$
Demonstrated rate capability	100 $\frac{MHz}{cm^2}$

Table 8.3 – Main specifications and parameters for the design and operation of ME0 [8]

The first thing to note in comparison to GE2/1 is the greater number of detectors. This is because, as mentioned in Section 8.2.1, each ME0 stack is composed of six detectors rather than two, as seen in Figure 8.19 along with their corresponding cable tray for better cable management. These cable trays are a first for a CMS GEM system, being integrated into the stack itself. Figure 8.20 shows a mock-up of the stacks being installed into the endcap, with the installation jig (dark blue) taking the cable tray into account. The stacks themselves are each mounted on a 15mm-thick aluminum plate, which supports the stack and creates an independent, complete unit which can be tested and qualified as a whole prior to installation into CMS. The rails which are used for aligning and sliding the stacks into the endcap are seen in blue on the endcap itself and orange on the stacks.

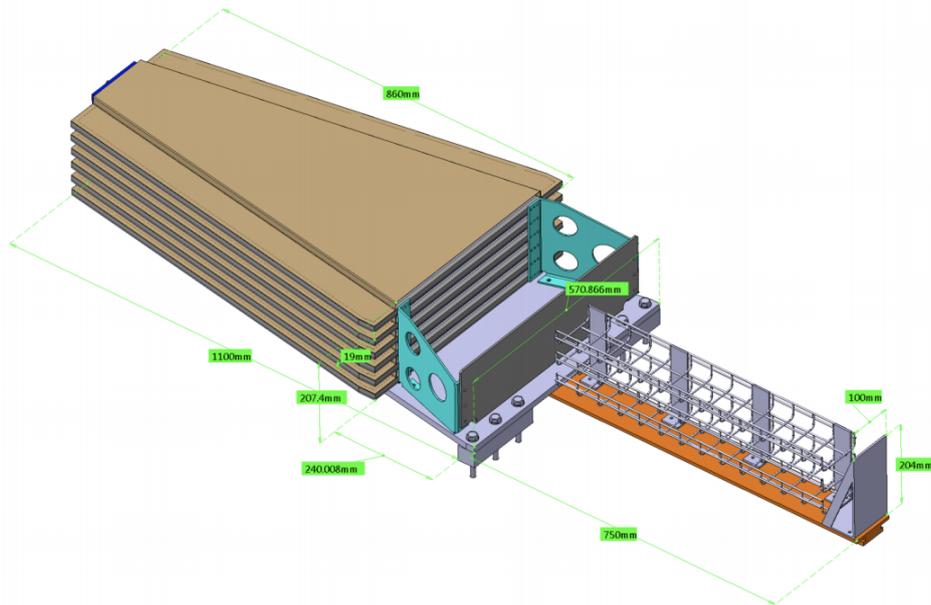


Figure 8.19 – Schematic of a stack of six ME0 detectors with corresponding cable trays, including all major dimensions [8]

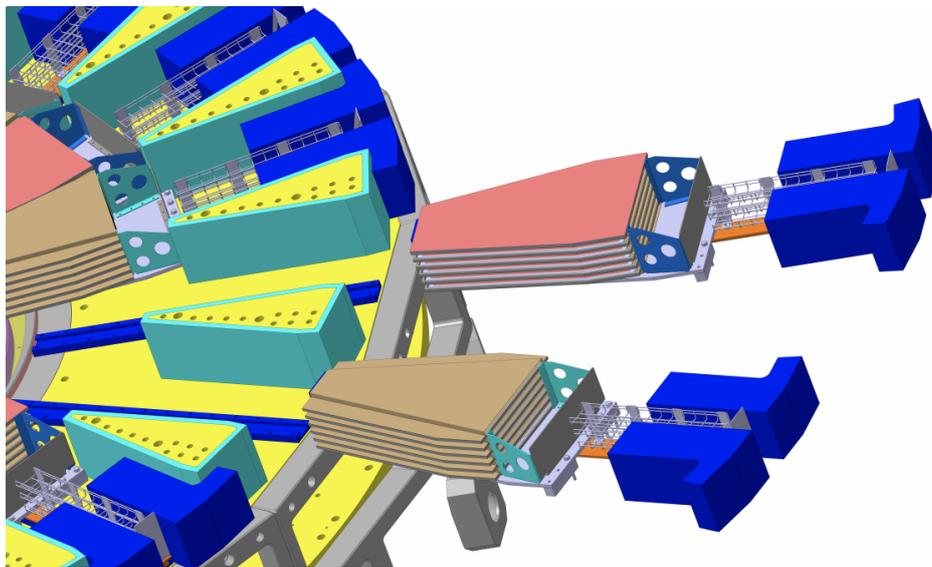


Figure 8.20 – CAD drawing of the installation method and placement of the ME0 stacks within the CMS endcap

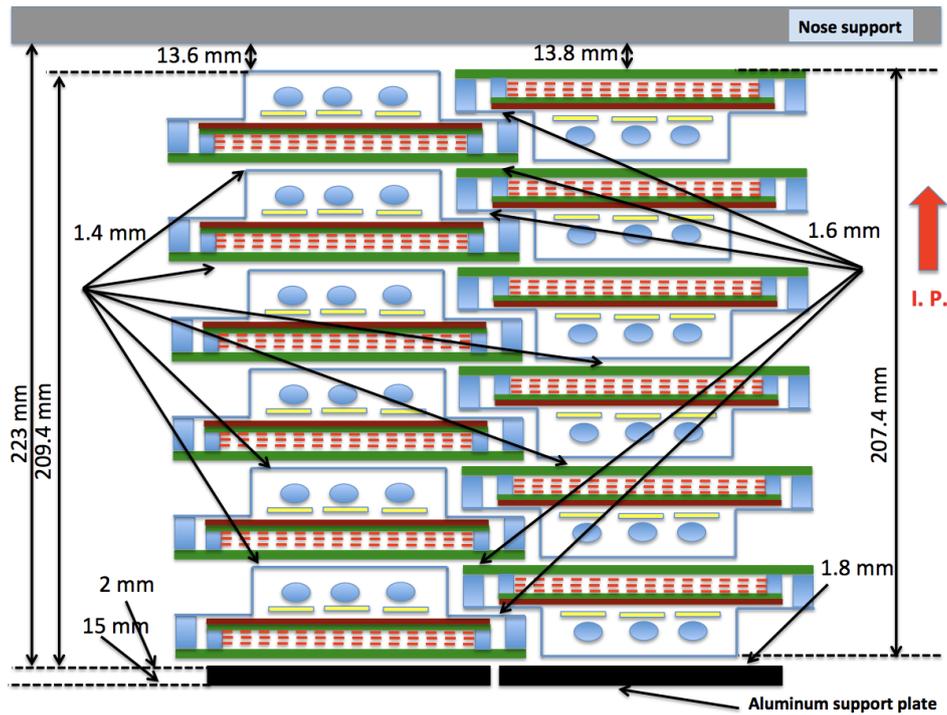


Figure 8.21 – Cross-sectional schematic of two stacks of overlapping ME0 detectors, including their aluminum support plates [8]

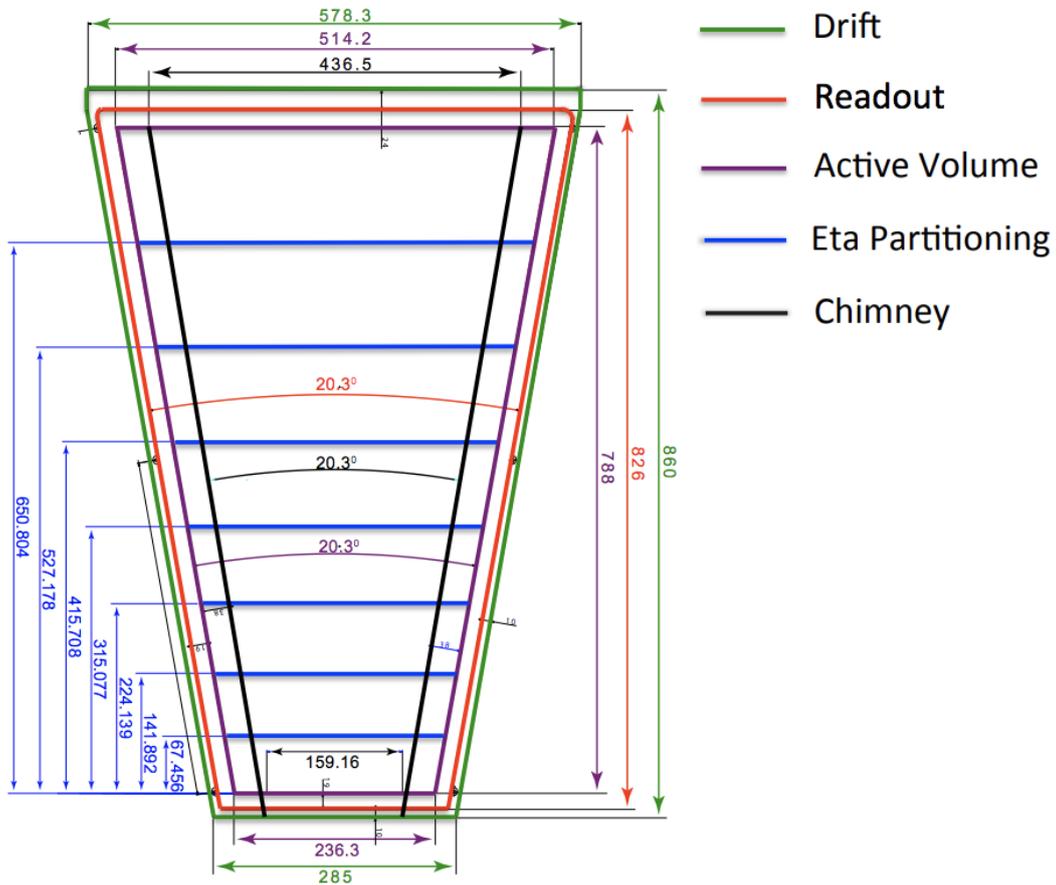


Figure 8.22 – Dimensions of an ME0 detector, including drift and readout boards, aluminum chimney, and active readout area [8]

This figure also shows the layout of the detectors within the endcap, with every other stack being flipped 180° about its longitudinal axis, so that they can be made to slightly overlap in ϕ and maximize the coverage of the system. The dimensions of this overlap can be seen in the cross-sectional view of Figure 8.21, including all required clearance distances, although the figure itself is not to-scale. Here, the interaction point of CMS is oriented towards the top of the figure.

The detectors, although they cover the same 20° opening angle, are also smaller than either the GE1/1 or GE2/1 detectors, with both a smaller active readout area and active module volume. The main dimensions of an ME0 detector, in addition to being listed in Table 8.3, can be seen in Figure 8.22. As with GE2/1 in Figure 8.3, the relative size of the ME0 system to the GE1/1 and GE2/1 systems can be seen in Figure 8.23.

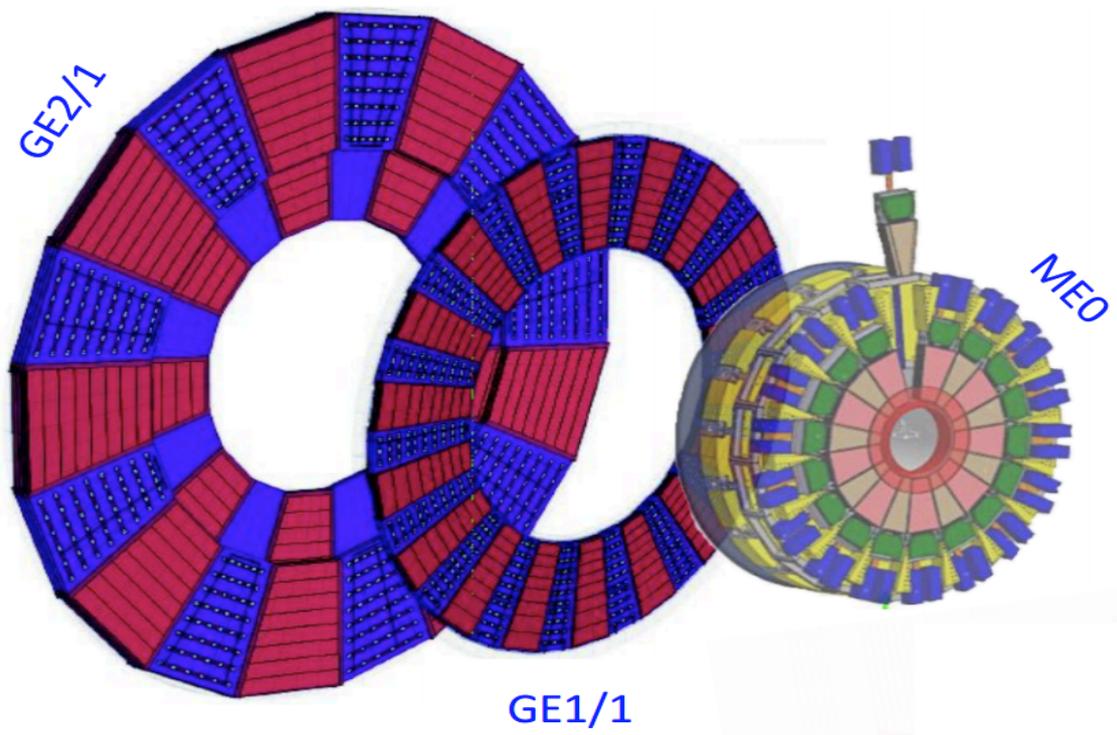


Figure 8.23 – The ME0 ring of detectors superimposed on the GE1/1 and GE2/1 rings of detectors, illustrating their size differences

8.2.3 ME0 DAQ Electronics System Overview

The flow diagram for the ME0 front-end and back-end DAQ electronics is shown in Figure 8.24. Beginning on the right of the figure we have an ME0 chamber, which is broken down into 24 readout sectors - three columns in ϕ and eight partitions in η , just like GE1/1. Each readout sector contains 128 radial strips which are read out by a single VFAT3 chip.

From here, for each VFAT, one 320 Mbps elink carries the DAQ information, and eight 320 Mbps elinks carry the trigger information to the optohybrid board, which is described in the next section. From simulations using 240 pile-up interactions per bunch-crossing and three times the expected neutron background (see: Reference [108] for more information on FLUKA simulations of background particle rates within the CMS GEM detectors), the average data rate from the VFAT3 closest to the beam pipe for ME0 has been determined to be 275 Mbps without zero suppression, well below the 320 Mbps capabilities of these

elinks.

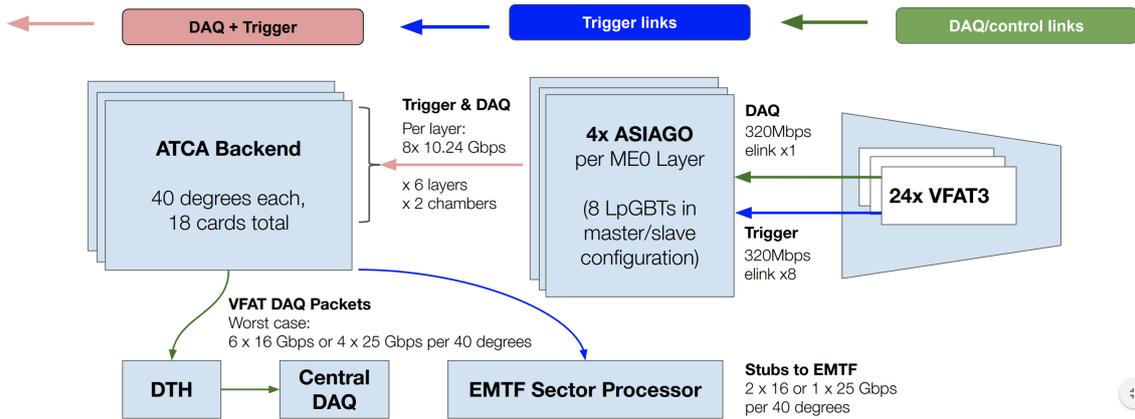


Figure 8.24 – Flow diagram of the ME0 DAQ electronics [109]

8.2.4 ME0 “ASIAGO” Optohybrid

The most notable change in the ME0 electronics is the use of a new class of optohybrid board, called the ASIC And Gigabit Optics board, or “ASIAGO”, a first prototype of which can be seen in Figure 8.26. What sets the ASIAGO apart, aside from its cheesy name and adorable mouse logo, is the complete lack of an FPGA on-board, making the optohybrid and all other on-detector electronics fully radiation-hard. All data compression is instead done on the back-end, which will be explained later in this subsection.

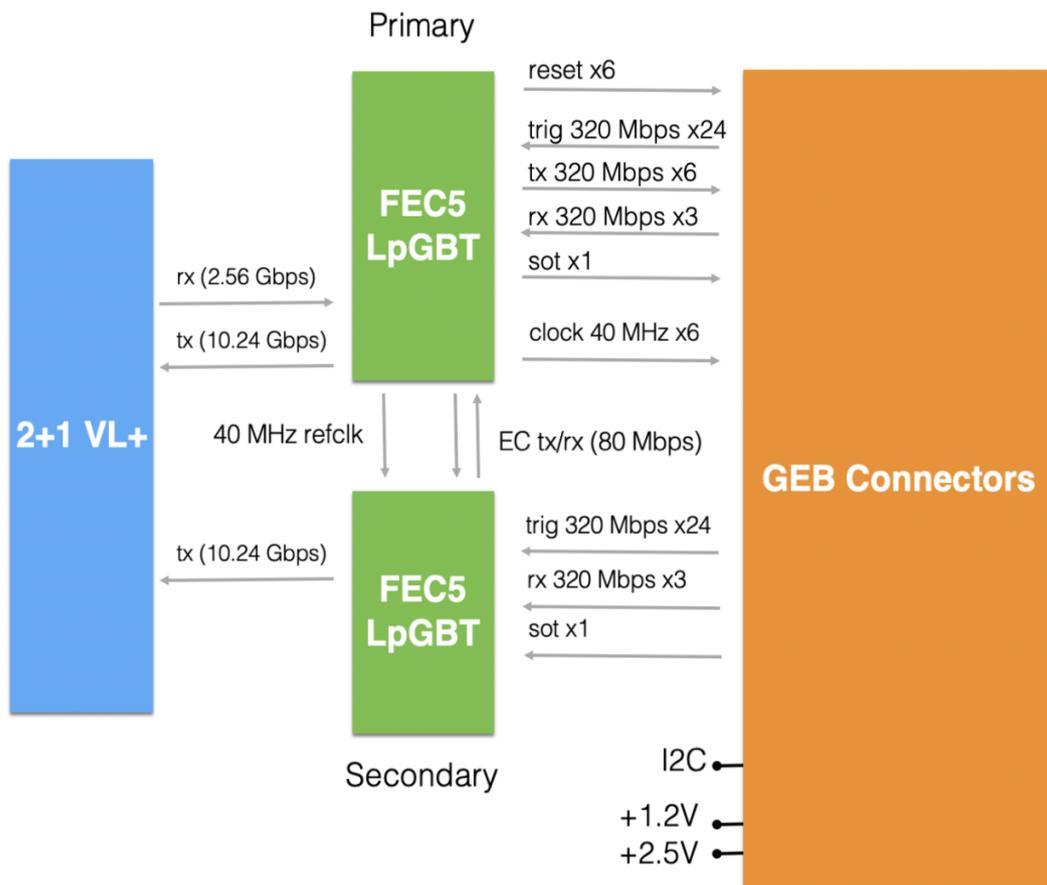


Figure 8.25 – Architecture diagram for the ASIAGO board [110]

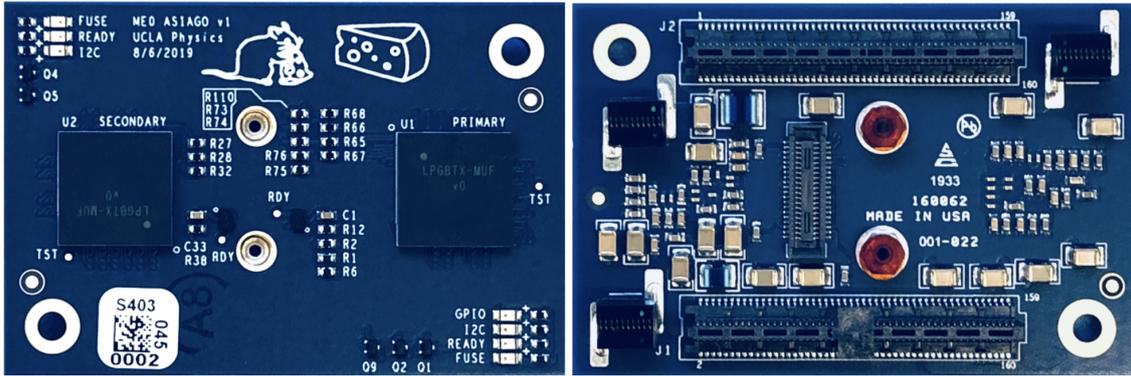


Figure 8.26 – First prototype of an ME0 ASIAGO board, as designed by the GEM electronics group at UCLA [110]

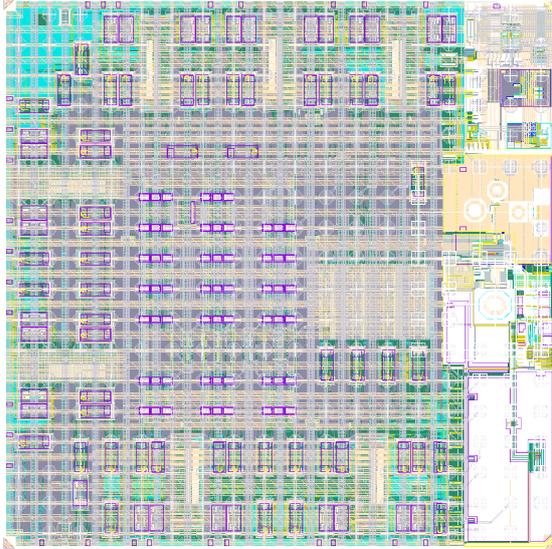
Physically, the ASIAGO is significantly smaller than its predecessors, at only $\sim 3 \text{ cm} \times 5 \text{ cm}$. Its main components include two LpGBT chips (see: Reference [111]), versatile link PLUS (VL+) optics (see: Reference [112]), baseboard connectors, and miscellaneous passive parts, making each board relatively inexpensive to produce. An architecture diagram for the ASIAGO can be seen in Figure 8.25.

8.2.4.1 LpGBTs

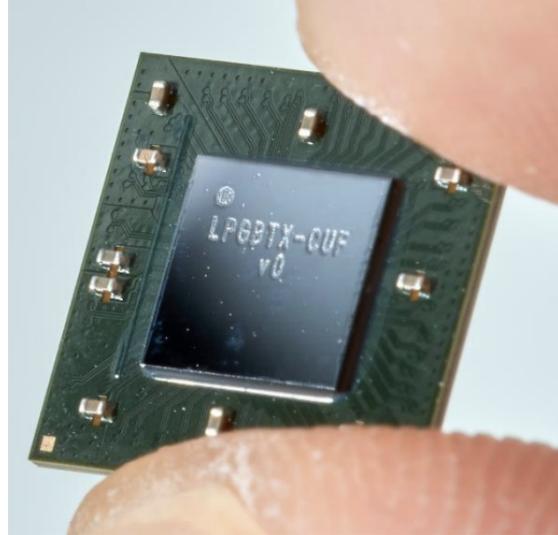
In place of the GBTx chips of GE1/1 and GE2/1, the ME0 ASIAGO makes use of two Low Power GigaBit Transceiver (LpGBT) chips, arranged in a master/slave configuration. The LpGBT chip is a radiation tolerant ASIC that is being designed at CERN for data, timing, trigger and control applications in the context of the HL-LHC. Its design specifications were specifically chosen for radiation hardness, with a total ionizing dose of 200 Mrad and extensive SEU protection. Its high-speed links have speeds of up to 10.24 Gbps for transmitting and 2.56 Gbps for receiving data. It also has many features for use in control and monitoring, including three I2C masters, a 16-bit general purpose I/O port, a 10-bit ADC with 8 multiplexed inputs, an 8-bit voltage DAC and 8-bit current DAC, as well as a temperature sensor. The chip itself is also physically very small, at just $9 \text{ mm} \times 9 \text{ mm} \times 1.25 \text{ mm}$. Its block diagram and a photo of one of the earliest prototypes can be seen in Figures 8.27a and 8.27b respectively [113].

By operating the LpGBTs in FEC5 mode, at the cost of having less error correction, each LpGBT can have up to 28 elinks operating at 160 Mbps. And as mentioned in Section 8.2.3, each VFAT has nine elinks. Thus, a single LpGBT chip can handle three VFATs, and consequently a two-LpGBT chip ASIAGO can handle six VFATs. As the ME0 detectors have a total of 24 VFATs each, four ASIAGO boards are needed for a given ME0 detector, similar to GE2/1. This has the added benefit of keeping all signal lines short, rather than having to route them long distances across the GEB.

In Figure 8.25, we see the master LpGBT labeled as “Primary” and the slave LpGBT labeled as “Secondary”. Although both LpGBTs deal with a start of transmission (sot) signal, three DAQ (Rx) links, and 24 trigger links to or from the GEB connectors, it is only the master LpGBT that also carries six Tx links, six reset signals, six 40 MHz clocks, the latter of which is transmitted to the slave LpGBT as a 40 MHz reference clock. As well, though both LpGBTs have a Tx VL+, only the master LpGBT has an Rx VL+ link.



(a) Layout of the LpGBT ASIC [111]



(b) Prototype LpGBT chip [114]

8.2.4.2 Versatile Link PLUS

The versatile link PLUS is the next generation of versatile link technology. The differences between VL+ and its predecessor can be found in Table 8.4.

	Versatile Link (VL)	Versatile Link PLUS (VL+)
Upgrade phase	Phase I	Phase II
Optical mode	Single- and multi-mode	Multi-mode
Flavors	1 Tx + 1 Rx, 2Tx	up to 4 Tx (+1 Rx)
Radiation resistance	Up to calorimeter-grade	Up to tracker-grade
Form factor	SFP+	Custom miniature
Data rate	Tx/Rx: 5 Gbps	Up: 5/10 Gbps, Down: 2.5 Gbps

Table 8.4 – Differences in specifications between VL and VL+ optical links [113]

Compared to its predecessor, the new VL+ is physically smaller, as shown in Figure 8.28, with prototypes measuring $10 \text{ cm} \times 20 \text{ cm} \times 2.5 \text{ mm} - 4 \text{ mm}$. And yet it is far more versatile, with up to 4 Tx and 1 Rx, as configurable by masking channels. It is designed to be pluggable like the VTRx, with an electrical connector on one side and an optical connector or short pigtail on the other side. It utilizes an 850 nm vertical cavity surface emitting laser (VCSEL), which is characterized by high energy efficiency and high data rates over short (few 100's of meters) data links. In particular, for the VTRx+, Tx can be 5 or 10 Gbps and Rx is 2.5 Gbps. It is also designed to withstand harsh environmental conditions: it can operate safely in a temperature range of $-35^\circ\text{C} - +60^\circ\text{C}$, has been qualified with a total radiation dose of 1 MGy (with investigations of up to 2 MGy), and can withstand a total fluence of 10^5 hadrons per cm^2 , the latter two of which meet the stringent requirements of the CMS tracker [113].

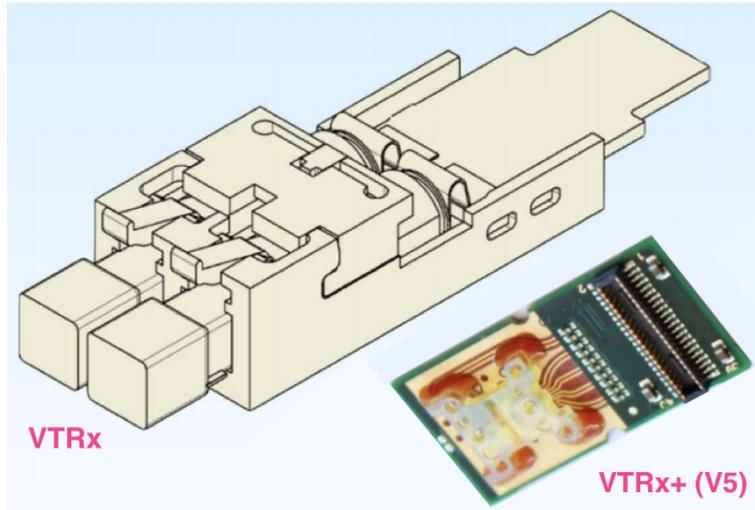


Figure 8.28 – To-scale size comparison of the VL VTRx module with the VL+ VTRx+ V5 prototype

8.2.5 ATCA Back-End

As with GE2/1, ME0 utilizes ATCA technology for its back-end electronics - specifically, the APx boards which are being designed for use in the CMS Phase II upgrade for the barrel calorimeter, muon, and trigger systems. Each ATCA card covers a 40° area, or two ME0 stacks. To connect to the front-end electronics, each ATCA card has 96 10.24 Gbps uplinks from the LpGBTs for trigger, DAQ, and control and 48 2.56 Gbps to the LpGBTs for TCDS and control. Then to connect further down the line of back-end electronics, it will have either one 25 Gbps or two 16 Gbps links to the EMTF, and either four 25 Gbps or six 16 Gbps links to the DTH. Unlike GE2/1, which will require one DTH per ATCA crate, ME0 will require two. From this point, the back-end can be considered to be nearly identical to that of GE2/1.

Chapter 9

Foil Double Segmentation: a Discharge Mitigation Strategy for GE2/1 and ME0

In GE1/1 detectors, the surface of the GEM foil which faces the readout board is a single continuous conductor, whereas the other surface, which faces the drift board, is segmented into sectors, each of which are separated from each other and connected to the HV filter via a $10\text{ M}\Omega$ protection resistor which is soldered directly onto the foil. Each of these sectors has an area of less than 100 cm^2 by design, imitating the area of a $10\times 10\text{ cm}^2$ detector. This segmentation and the corresponding HV filters are illustrated in Figure 9.1.

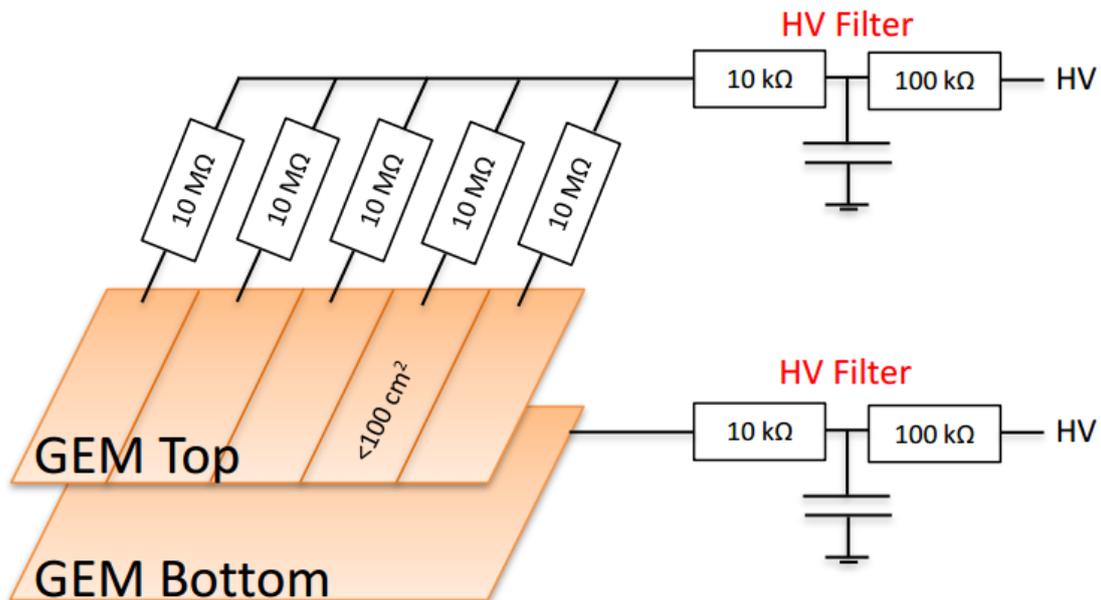


Figure 9.1 – Figure illustrating the single segmentation of a traditional GE1/1 GEM foil [85]

However, by leaving the bottom of the foil to act as a single continuous conductor with a high total capacitance, a higher probability for discharge propagation remains. To lessen this probability further, the bottom of the foil can be segmented as well, also into sectors of an area less than 100 cm^2 , as shown in Figure 9.2. The primary difference here lies in the resistors. Unlike the traditionally-segmented top of the foil, the resistors present between the bottom segments are not protection resistors, but rather serve as decoupling resistors. In this way, not only is the capacitance of the foil reduced, but the HV sector decoupling is improved. Although this is still under study, a generally-acceptable value for these decoupling resistors has been found to be $100\text{ k}\Omega$. In this way, a large ratio is maintained with the top protection resistors, and the bottom potential will not jump towards the top and potentially create a discharge propagation [115].

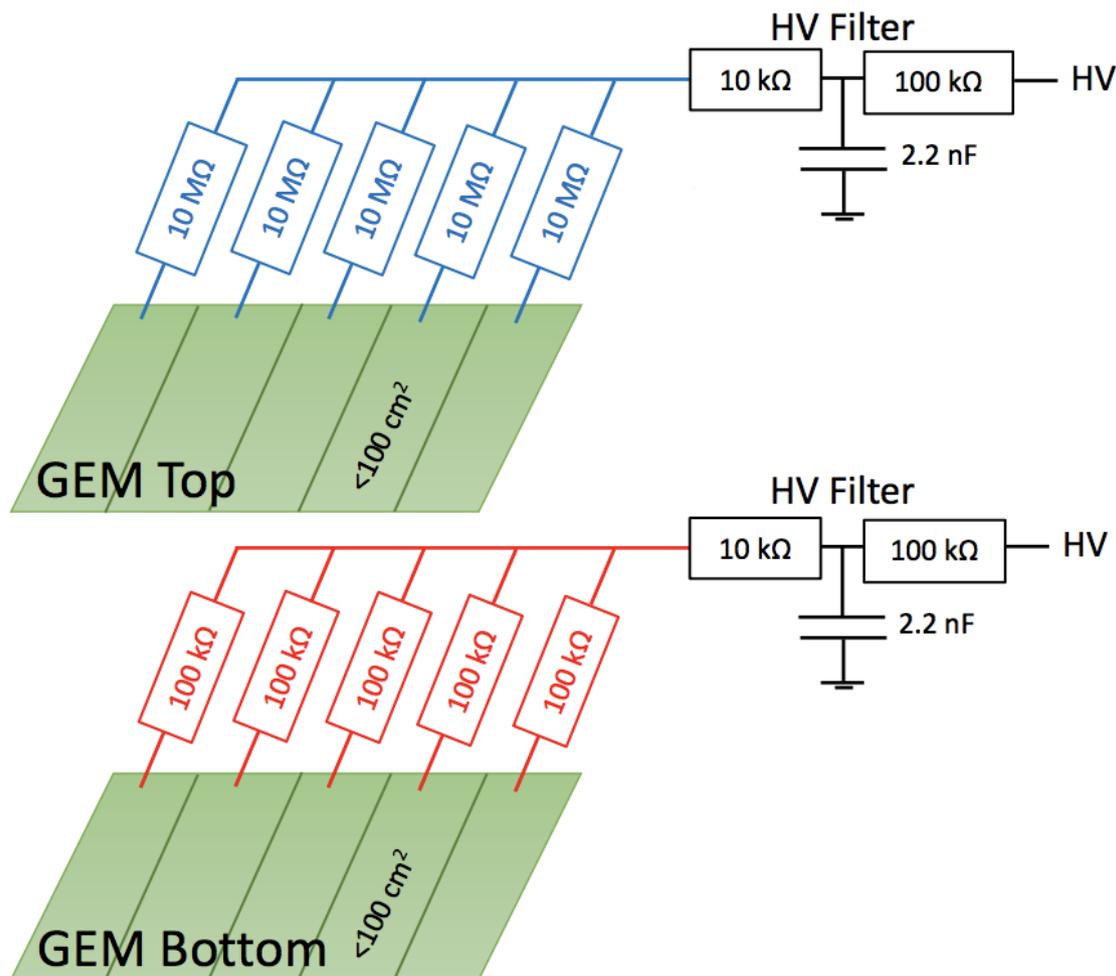


Figure 9.2 – Figure illustrating the double segmentation of a GE1/1 GEM foil [115]

This chapter discusses the studies done into the reduction in discharge propagation this double-segmentation affords us, and any issues that it may present.

9.1 Discharge Rate in a Double-Segmented GE1/1 Chamber

In order to study the propagating discharge rate in a large double-segmented GEM detector over time, and therefore confirm the effectiveness (or lack thereof) of double segmentation as a successful mitigation strategy, a modified GE1/1 detector, seen in Figure 9.3, was constructed with a double-segmented G3 foil. Several mm-diameter holes were drilled into the drift board of the detector and covered with food-grade plastic wrap and Kapton tape in order to maintain gas tightness. These holes would allow for the placement of an ^{241}Am source to allow α particles to penetrate the detector and trigger discharges at a faster rate than with cosmic rays alone. This source, which was measured to have an activity of 39.33 kBq on June 6th, 2018, was taped to a midline hole at the wide end of the detector using Kapton tape. All Panasonic connectors were terminated with 50 Ω Panasonic-to-LEMO terminators, and a 100 k Ω resistor was added to each HV sector. The voltage settings can be found in Table 9.1, with the note that this GE1/1 chamber still utilizes the current divider method. All electronics modules used in this study were NIM modules, as seen in Figure 9.4, which were connected as described in the flow chart of Figure 9.5. Also seen in

this flow chart and in Figure 9.3, an HV probe (yellow) was connected to the G3 HV pad in order to better monitor the presence of propagating discharges within the chamber.

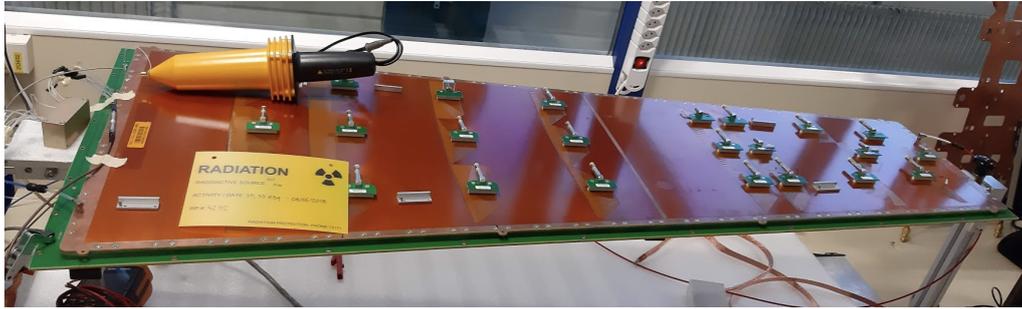


Figure 9.3 – Detector portion of the experimental set-up for the double-segmented GE1/1 propagating discharge rate study

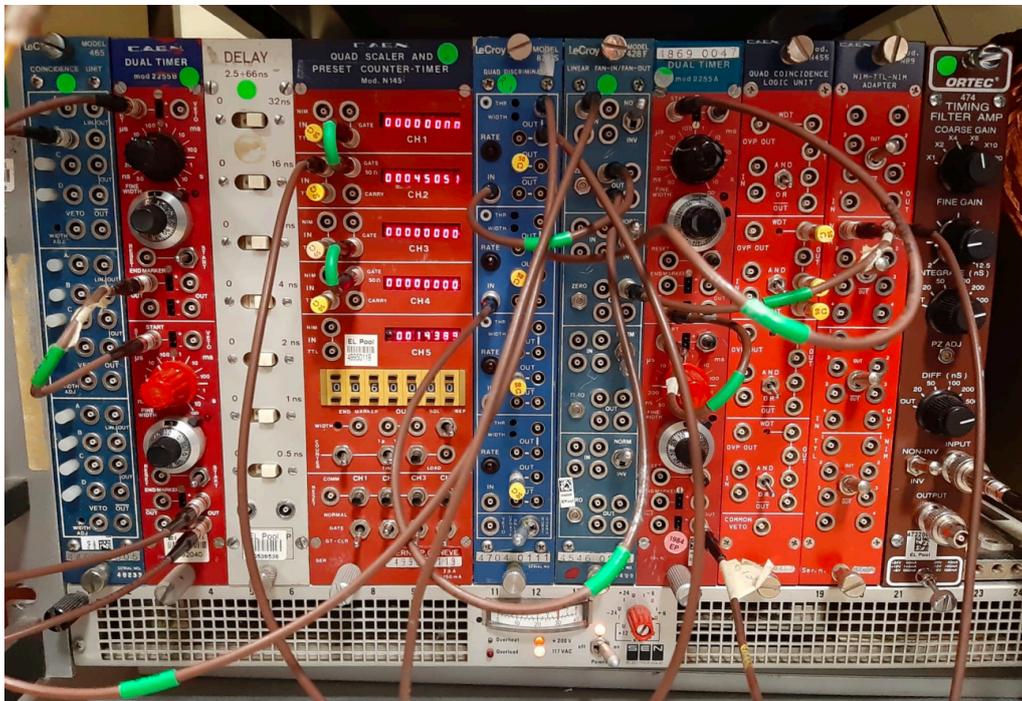


Figure 9.4 – NIM module portion of the experimental set-up for the double-segmented GE1/1 propagating discharge rate study

Detector Layer	Vmon (V)	V ₀ (V)	I _{mon} (μA)	I ₀ (μA)
Drift	787.50	788.23	100	0.239
G1 _{top}	397.00	396.95	100	0.032
G1 _{bottom}	306.60	306.67	100	0.043
G2 _{top}	385.00	385.16	100	0.280
G2 _{bottom}	512.50	612.52	100	0.398
G3 _{top}	400.00	400.05	100	0.418
G3 _{bottom}	437.50	437.54	100	2.516

Table 9.1 – HV settings and readings for the double-segmented GE1/1 propagating discharge rate study

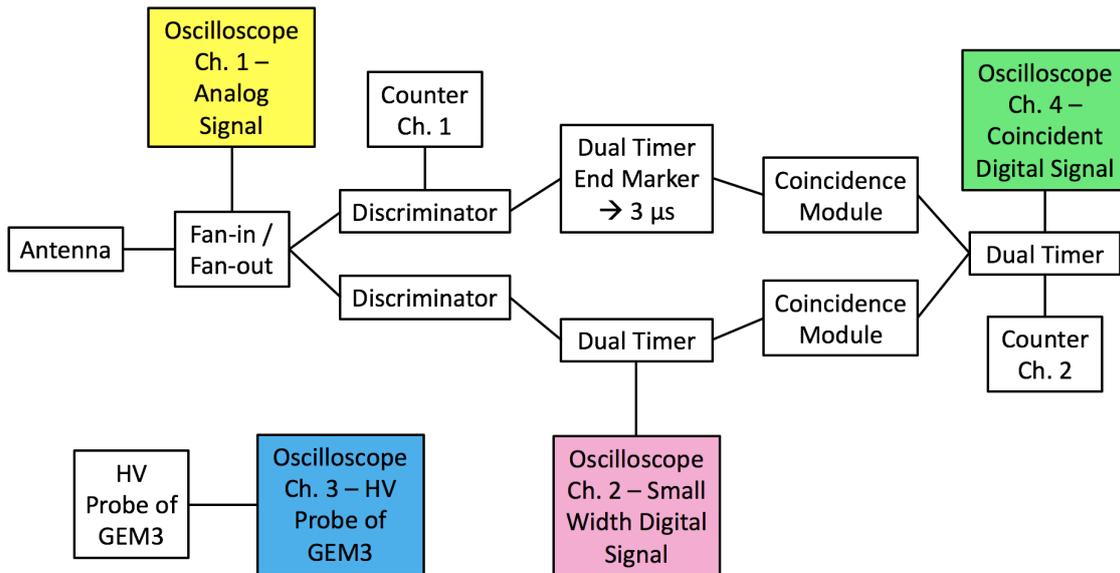


Figure 9.5 – Flow chart of the NIM modules used in the discharge rate study

The choice was made to trigger the oscilloscope on the primary discharge. In order to more accurately detect this, a LEMO cable antenna was placed at the HV pads of the detector and fed into a LeCroy model 821CS linear fan-in/fan-out module. One output of the linear fan-in/fan-out was fed into a LeCroy model 428F quad discriminator, the width and threshold of which can be set by the user. This was set such that the digital signal, triggered by the primary discharge, had a width that encompassed both the primary discharge and any secondary propagations, as seen on the top of Figure 9.6a. This signal was fed to a CAEN model N145 quad scalar and preset counter-timer in order to count the number of primary discharges recorded.

The width of the primary discharge signal is generally on the order of $2 \mu\text{s}$ - this time encapsulates both the initial primary discharge, which is the nearly-instantaneous plasma which causes a short between the two layers of the detector and causes the foil to discharge, and the recharging of the foils through the HV filter, which manifests on the oscilloscope as an exponential tail to the right of the primary voltage spike. To make it so that this digital signal no longer encompassed the primary discharge, we sent it to a CAEN model 2255A dual timer to time-shift the signal by the width of the primary discharge plus some margin of error. In our case, we chose this to be $3 \mu\text{s}$. This resulting signal can be seen on the bottom of Figure 9.6a.

Then, with a second LeCroy model 428F quad discriminator, we took the original signal and once again triggered on the primary discharge, this time with the discriminator's width set to the minimum possible width. The output of this discriminator was then fed to the dual timer, which was used to adjust the width of the digital signal such that it was only large enough to cover a single peak. This produced a digital signal which shows each part of the secondary propagating discharges as individual gates, as seen in Figure 9.6b.

This new signal, along with the time-delayed signal from Figure 9.6a, were then fed into a LeCroy model 465 coincidence unit. The resulting output signal, the coincidence of those two signals, is seen in Figure 9.6c. This signal allowed us to only measure the propagating portions of the discharge without including the initial primary discharge.

In order to measure only a single event and accurately count the number of propagating discharges, this coincident signal was fed into a second CAEN model 2255A dual timer to

stretch the pulse length such that a single gate covered the entire series of propagations, as seen in Figure 9.6d. This signal was then fed into the quad scalar in order to count the number of discharge propagations in comparison to the number of primary discharges.

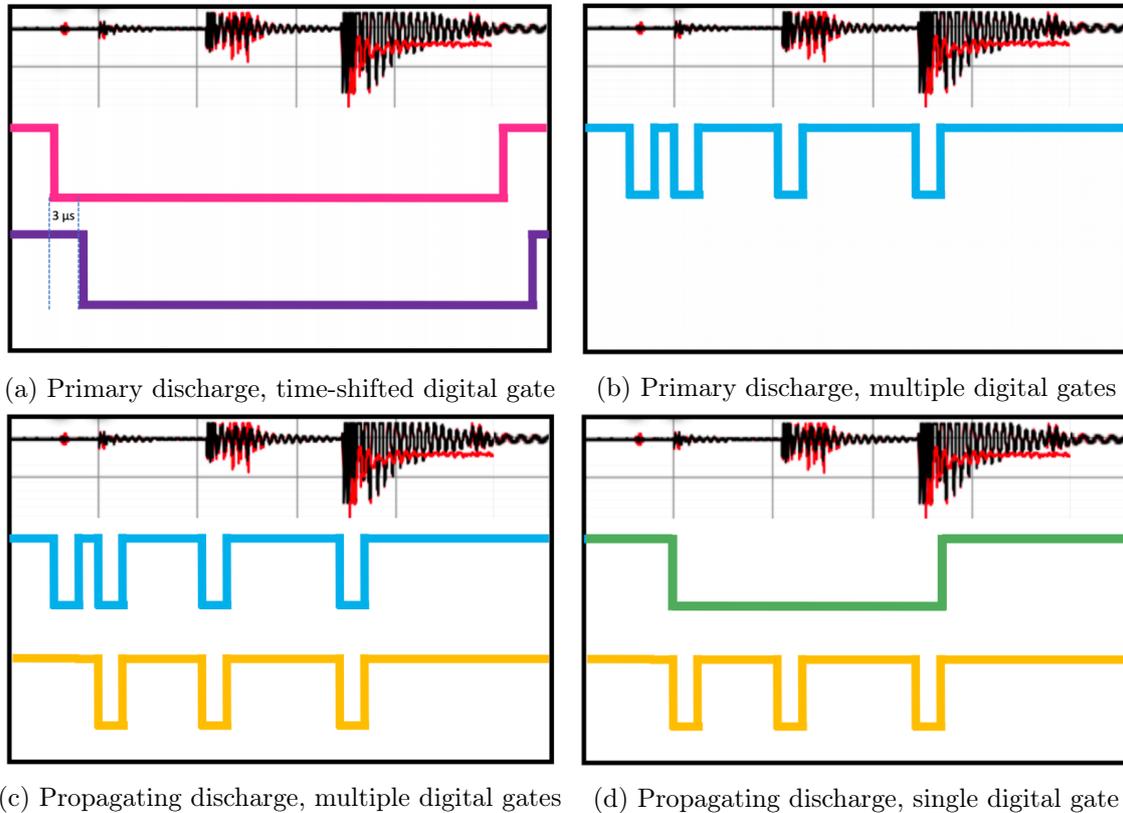


Figure 9.6 – Examples of each step of the oscilloscope-based signal recording

In addition to the quad scalar which counted each primary and propagating discharge, we fed the data from the oscilloscope to a nearby Windows PC, where a LabVIEW program was set up to record both screenshots of the oscilloscope and the corresponding data to file. Examples of a single discharge and a propagating discharge as recorded by this program can be seen in Figures 9.7a and 9.7b, respectively.

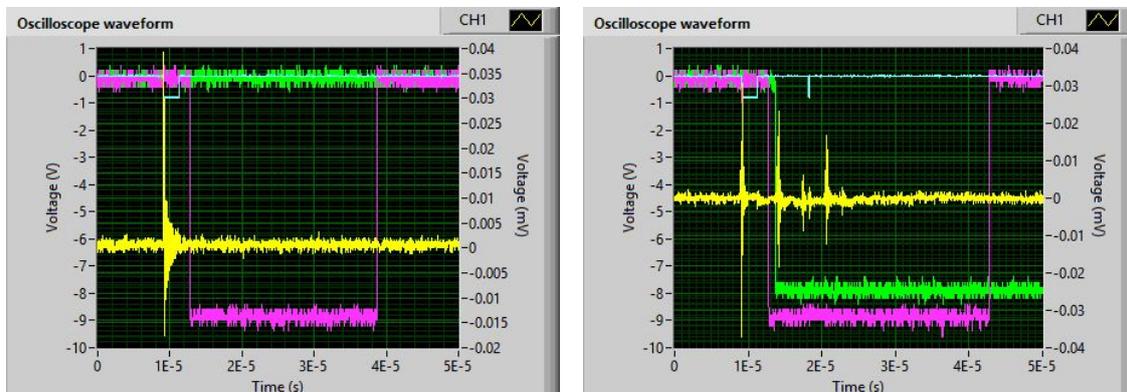


Figure 9.7 – Oscilloscope screenshots of single and propagating discharges in a double-segmented GE1/1 detector

Unfortunately, one key component of this set-up, a Digilent Analog Discovery 2 USB oscilloscope, was unavailable from late November 2019 through February 2020, and the set-up was repurposed for another study. Once the component became available again, the outbreak of the COVID-19 pandemic forced the shutdown of all CERN laboratories, prematurely ending this study. Due to the unexpectedly short experimental time period, not enough statistics were obtained in order to make a clear and precise determination of the rate of propagating discharges within the chamber.

However, preliminary results are nevertheless promising, as seen in Figure 9.8. Here, a comparison is made between a standard $10 \times 10 \text{ cm}^2$ GEM detector, a standard single-segmented GE1/1 detector, and a double-segmented GE1/1 detector. Even with fewer available statistics, the double segmentation provides a clear benefit against single-segmented detectors, showing a lower probability of discharge propagation for induction fields up to $10 \frac{\text{kV}}{\text{cm}}$ than even a simple $10 \times 10 \text{ cm}^2$ detector. Although more precision could most certainly be achieved in regards to the upper limit of discharge propagation probability, the data on the lower end of the induction field scale is clear. Indeed, at the nominal induction field present at CMS, $4 \frac{\text{kV}}{\text{cm}}$, no discharge propagation whatsoever was observed, in comparison to the 70% probability observed in a single-segmented chamber of the same size and type.

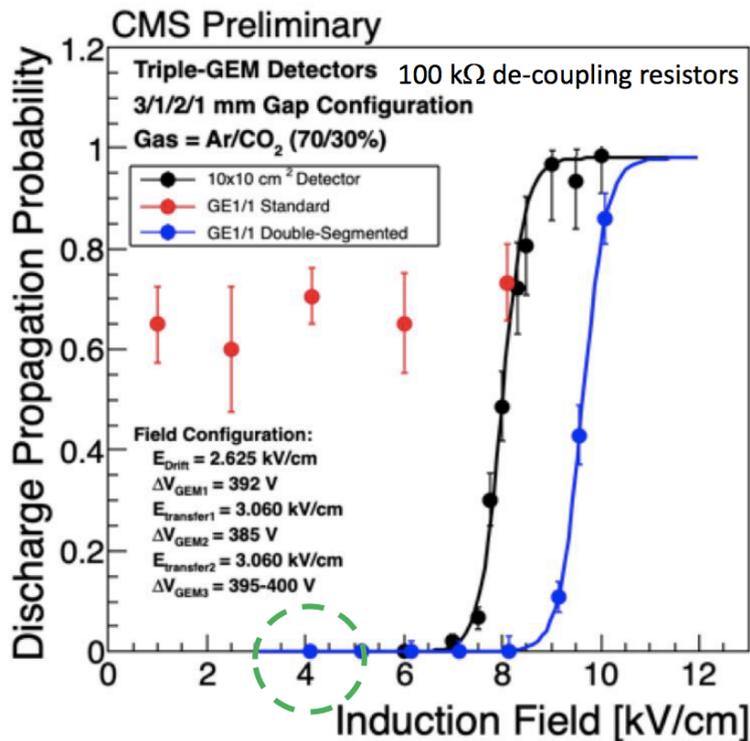


Figure 9.8 – Probability of a propagating discharge by induction field for three GEM detector types: a standard $10 \times 10 \text{ cm}^2$, a single-segmented GE1/1, and a double-segmented GE1/1, with the nominal CMS induction field highlighted [115].

9.2 “Dual-polarity” Cross-Talk Signals

In December 2019, just before the end of year shutdown, an unexpected observation was made while working with an ME0 chamber with double-segmented foils: the presence of a dual-polarity signal in VFATs that were not directly under radiation. After further investigation, this parasitic signal was found to be in all GEM detectors with double-segmented

foils, including GE1/1, GE2/1, and even prototype 10x10 cm² detectors. This, therefore, posed the first threat to the use of double-segmented foils as a viable discharge propagation mitigation strategy and called for significant investigation following the reopening of CERN in early 2020.

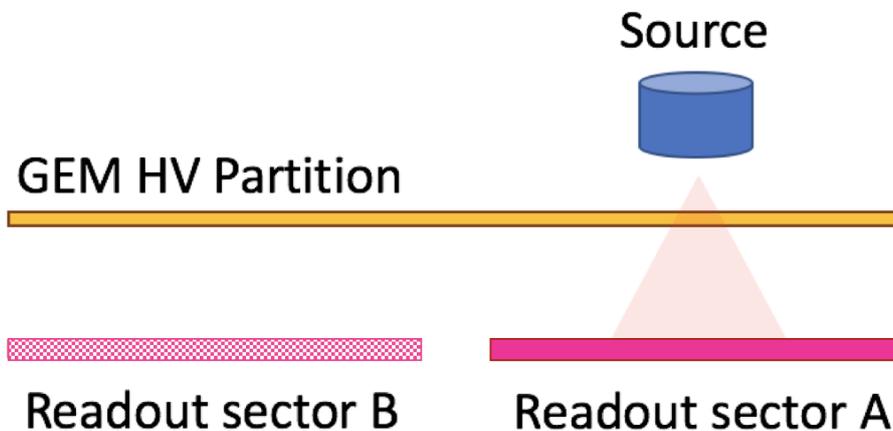


Figure 9.9 – Example set-up which could be used to study the dual-polarity cross-talk signal [116]

For any given detector type, we can consider a set-up which mimics Figure 9.9, in which a source is placed such that it irradiates a given readout sector. For an HV partition of an area on the order of 100 cm², the low capacitance of the partition means that electrical potential fluctuations are propagated along the entirety of the partition. As a result, we see the propagation of true signals (in readout sector A, in solid pink) in the adjacent readout sectors (such as readout sector B, in dotted pink) as a mirrored cross-talk signal via capacitive coupling. This results in the now-infamous image taken by Davide Fiorina of a 5-second oscilloscope persist of the opposite-polarity cross-talk signal and the true signal from which it originated, seen in Figure 9.10. Note that here, the oscilloscope is showing the analog output of an ORTEC pre-amplifier, not the original particle signals - this is what gives the signals their characteristic shape of having a large amplitude followed by a smaller undershoot.

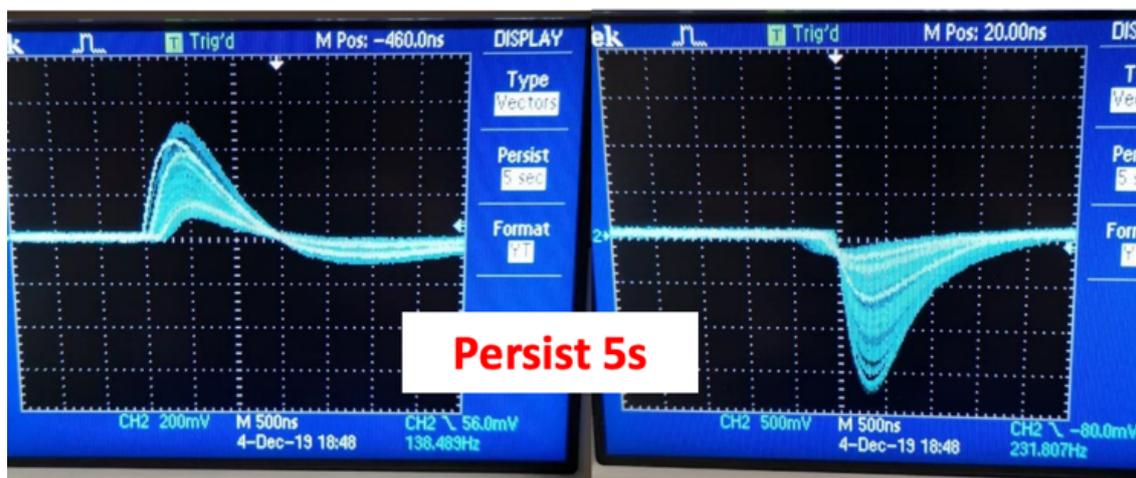


Figure 9.10 – Photos of an oscilloscope set to a 5-second persist, showing the opposite-polarity cross-talk signal (left) and the true signal from which it originated (right) as seen from the analog output of an ORTEC preamplifier [117]

The problem, then, comes from the shape of the VFAT3’s analog signals. The analog output signals of the VFAT3 chip’s preamplifier are of dual polarity, consisting of a large initial positive-polarity portion followed by a smaller negative-polarity portion. This shape is shown in Figure 9.11 for two types of signals - one originating from either x-ray radiation or a minimally-ionizing particle (MIPs) in teal, and the other originating from a HIP in pink. For each signal, the solid line represents the true signal that would be seen in the readout sector directly under radiation by the source (readout sector A in Figure 9.9). In order to detect these signals without being overwhelmed by noise, a discrimination threshold is set in the electronics, shown in the figure as the gold dashed line.

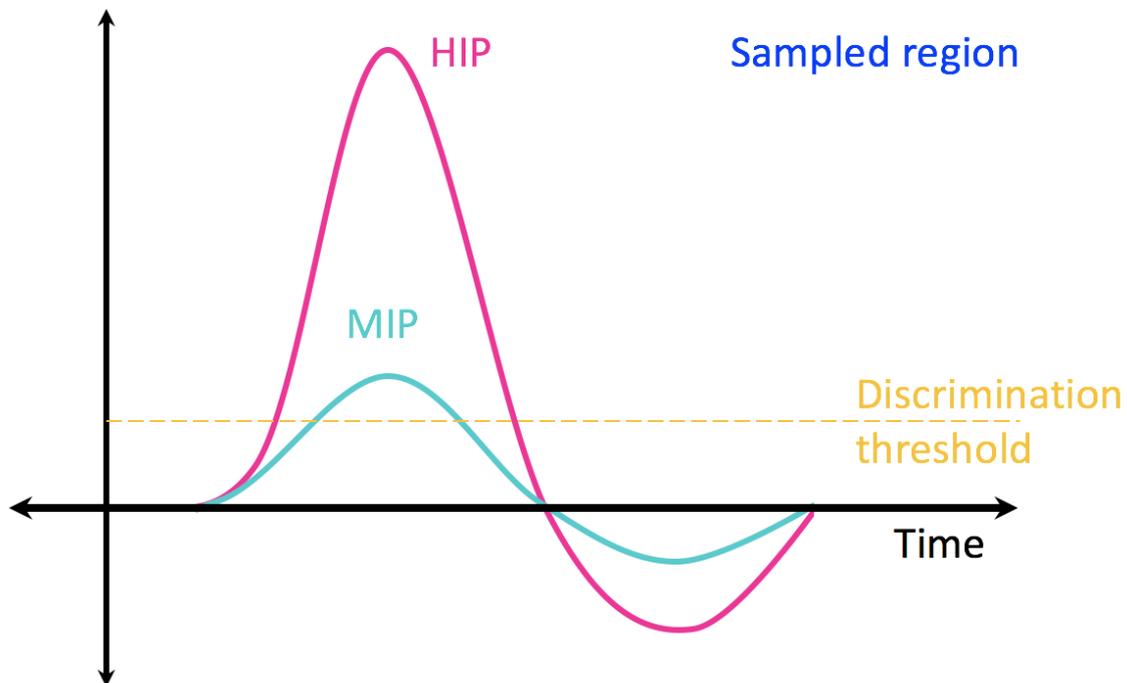


Figure 9.11 – Example of an x-ray/MIP and HIP signal in a GEM detector as seen from the analog output of the VFAT3 chip’s preamplifier

If the reflected signals in the neighboring readout sectors along the same HV partition (readout sector B in Figure 9.9) are then added in, as seen in Figure 9.12, the origin of the problem then becomes apparent. The signal’s negative-polarity tail becomes positive when reflected in the mirror region of the neighboring readout sectors. In cases where the amplitude of the signal tail is high enough, as it is for HIPs, the reflected tail, which is now the same polarity as the main portion of the true signal, can exceed the discrimination threshold and falsely trigger the electronics as if it were also true signal. This then becomes the cross-talk signal which poses a significant threat to detector operations, given that the GEM electronics have a finite bandwidth, which can be saturated by these false signals. This, in turn, introduces dead time into the system, and undesirably reduces the efficiency of the detectors.

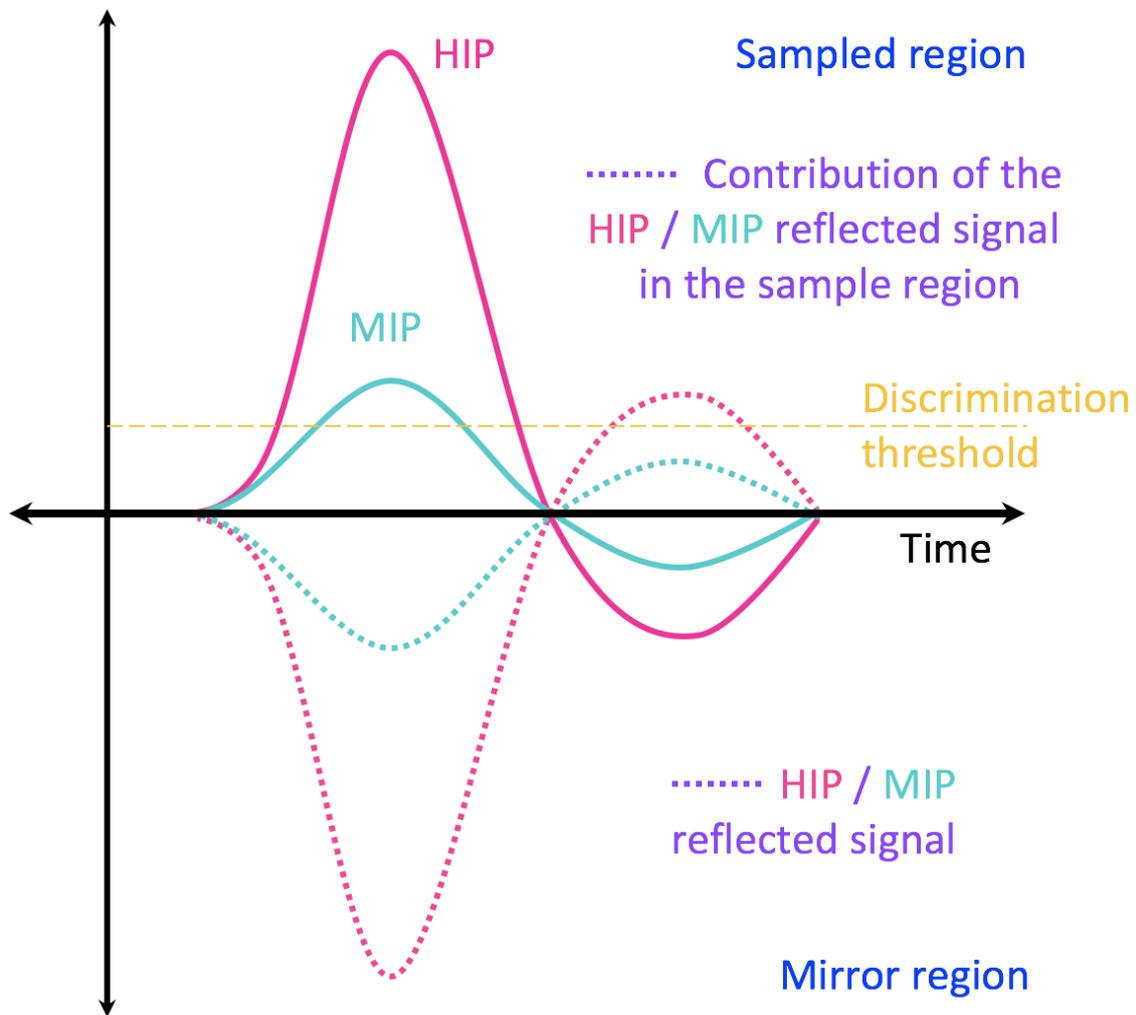


Figure 9.12 – Example of an x-ray/MIP and HIP signal being mirrored in adjacent regions of double-segmented GEM detectors as seen from the analog output of the VFAT3 chip’s preamplifier

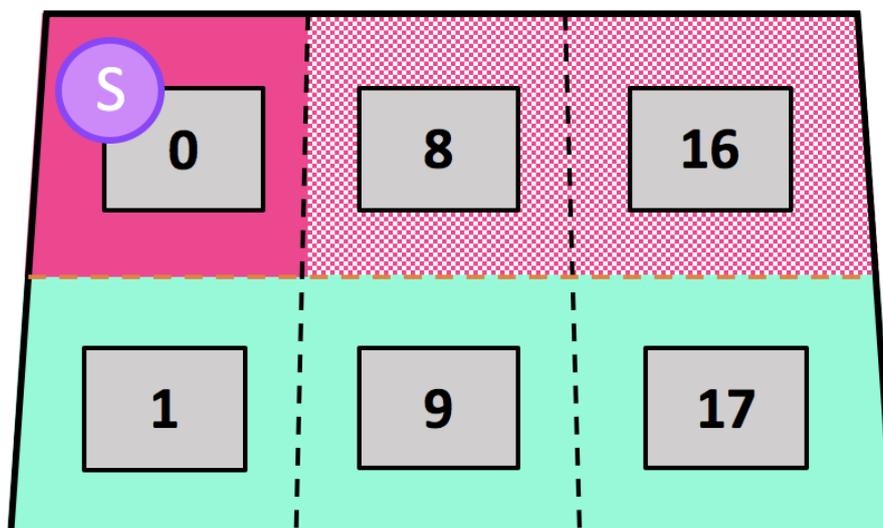


Figure 9.13 – Partial view of the narrow end of a GE1/1 detector, showing the location of the source and HV sectors.

Figure 9.14 is an example of this effect as shown by a set of S-Curves. To get these S-Curves, the source was placed under VFAT #0, at the narrow end of the detector as shown in Figure 9.13. Here, as in Figure 9.9, the solid pink sector represents the readout sector which is directly illuminated by the source and the dotted pink sector represents the readout sectors along the same HV partition which are not directly illuminated by the source. The green portion, in contrast, represents the readout sectors on a separate HV sector from the pink ones.

The signal from this source can be seen in the S-Curve of that same VFAT as the large spike in the center of the plot. The rest of the signals in VFATs #0, #8, and #16 are the result of the cross-talk pulse. However, despite this cross-talk signal being seen throughout the entirety of the HV partition where the source is located, the neighboring HV partition, populated by VFATs #1, #9, and #17, shows no cross-talk contributions. So while the effect can be widespread within a detector, it is only seen in the HV partition(s) where the source is located.

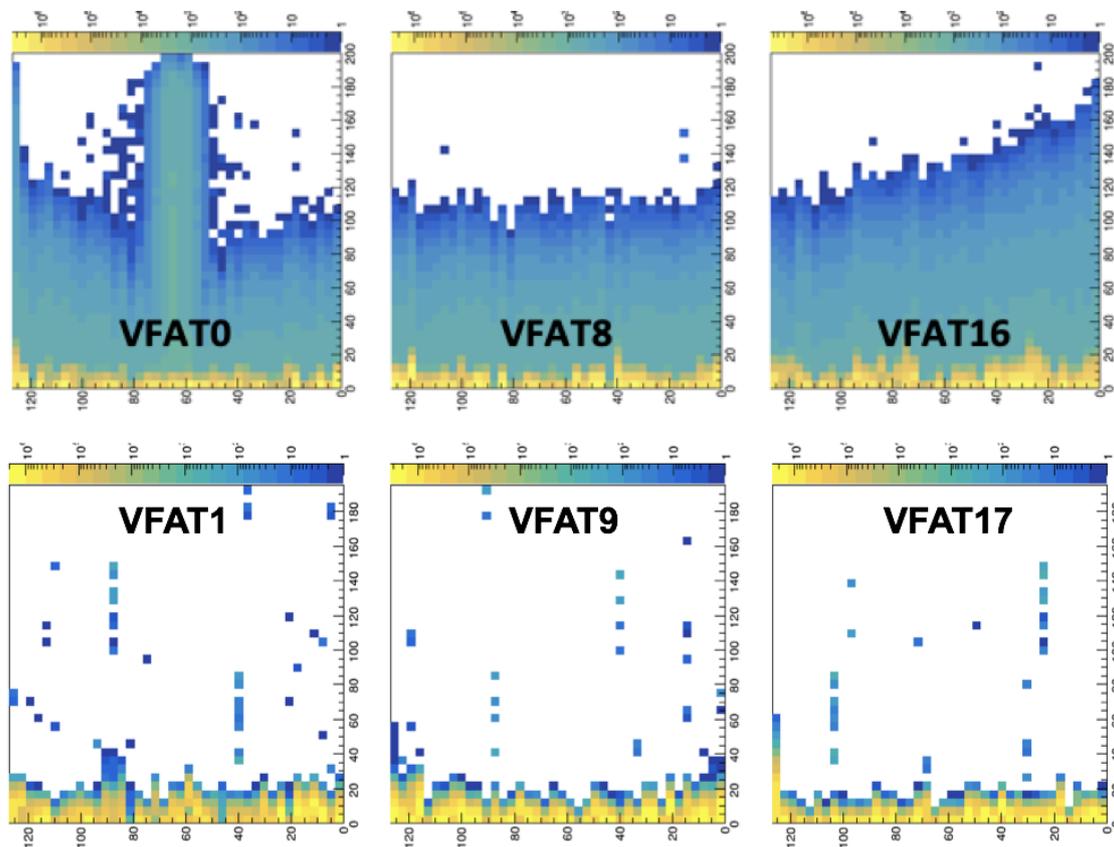


Figure 9.14 – Partial set of S-Curves showing how only VFATs within the same HV partition as the source are affected by the cross-talk signal

These examples are, however, a simplified presentation of the issue, considering only one affected readout sector. However, in a real GEM detector, a given HV partition covers more than simply two readout sectors, and a radioactive event may not occur precisely within the bounds of a single HV partition. Consider, for example, Figure 9.15, a top-down view of a GEM detector’s readout board and foil. Here, our hypothetical source has been placed on the boundary of two HV partitions. The true signal from the radioactive source affects the readout sector whose strips are highlighted in green. As a result of capacitive coupling, the resulting cross-talk signal can then be seen in all of the other

readout sectors on both HV partition 1 and HV partition 2. As a result, the cross-talk signal can be widespread within the detector.

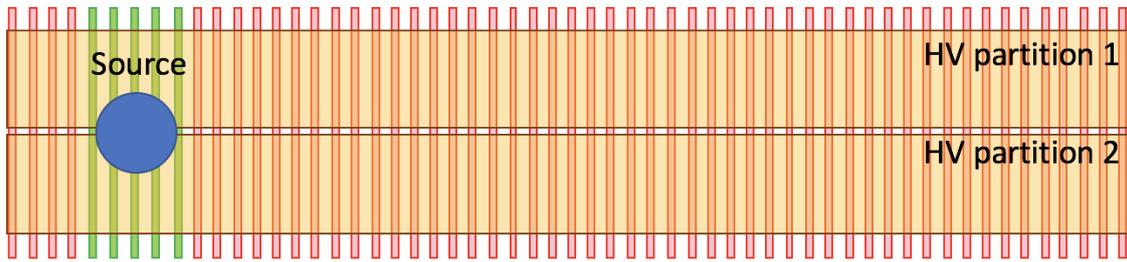


Figure 9.15 – Figure showing how all readout strips within a given HV partition are affected by the cross-talk signal [117]

9.3 GE1/1 Segmentation Size Studies

Although, as mentioned in Section 9.2, this cross-talk signal is seen in double-segmented detectors of all sizes, this does not mean that they are all equal. Thus, in order to study the effect the size of the HV partitions has on the magnitude and prevalence of the signal, a study would be performed using a GE1/1 detector which would be modified to take the same measurements with differently-sized HV partitions. This detector is a prototype GE1/1 detector and initially required significant cleaning within the GEM clean room in order to fix an issue where it displayed significant leakage currents. However, after cleaning, flushing with pure CO₂, and a slow-ramping HV test to check for further leakage currents, the detector was deemed operational and suitable for these tests.

Due to the lack of available CTP7 cards in the laboratory, the detector set-up in question was modified to use a GLIB (see Reference [118]) instead for its communication needs. However, this made virtually no difference in the operations of the detector, as the GLIB was set up to mimic a CTP7 as closely as possible.

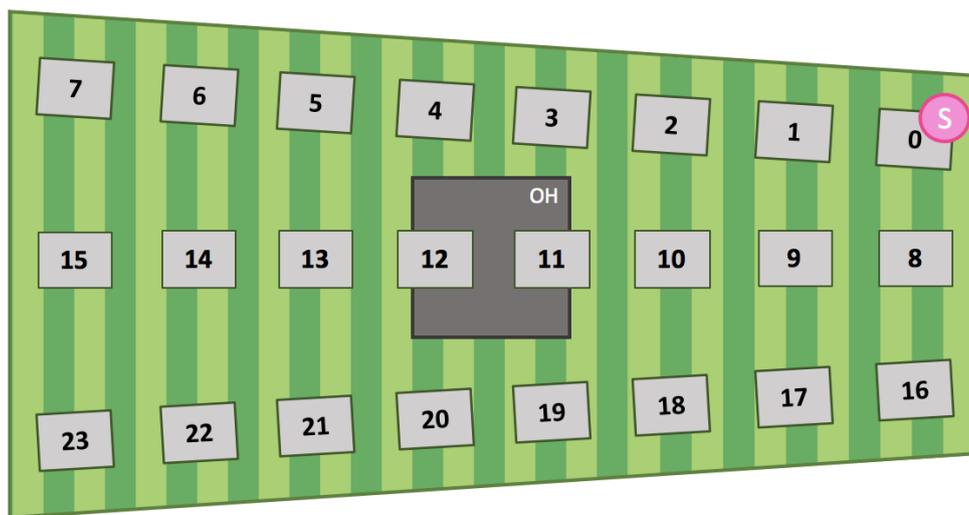


Figure 9.16 – Figure illustrating the prototype GE1/1 detector with unmodified HV partitions and location of the signal-triggering radioactive source

To begin, a set of so-called “fine segmentation” measurements would be taken with an unmodified double-segmented foil, which contains the typical 100 cm² HV partitions shown

in Figure 9.2. A not-to-scale mock-up of this detector with the correct amount of fine HV partitions can be seen in Figure 9.16.

First, a true baseline scan set was taken, with the HV and LV of the detector on, and without any radioactive source present. This set of S-bit threshold scans can be seen in Figure 9.17.

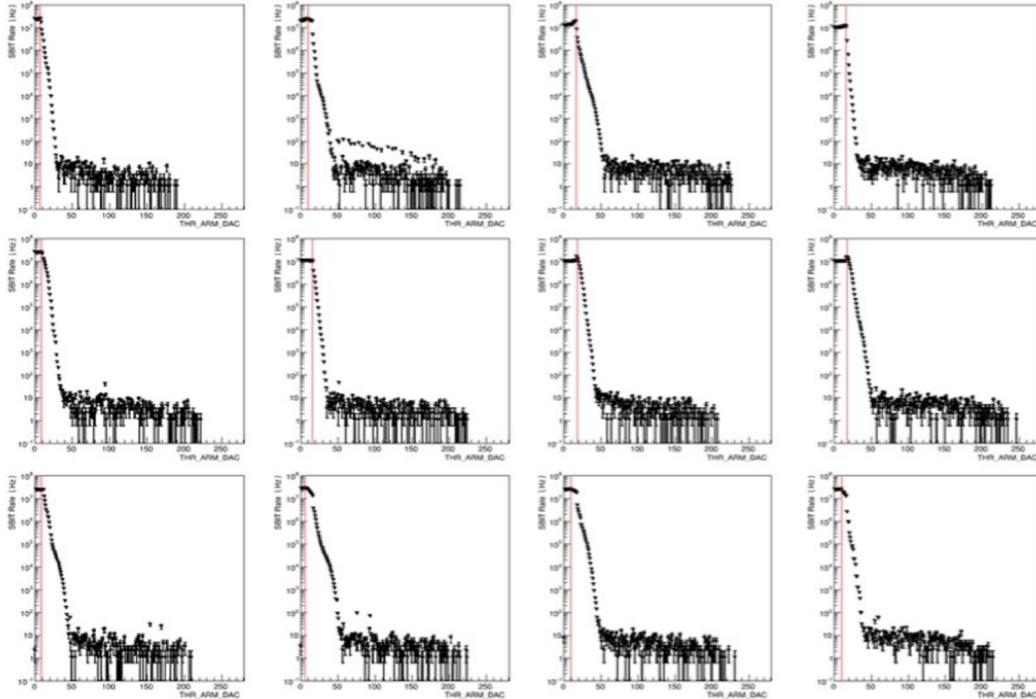


Figure 9.17 – Baseline S-bit threshold scans with LV and HV on, no source present, for the unmodified double-segmented GE1/1 detector

Then a ^{109}Cd source was placed at the position denoted by the “S” in Figure 9.16. This ^{109}Cd would serve as a source of x-rays, with the principal emissions of ^{109}Cd being 22.1 keV x-rays, as seen in the ^{109}Cd emissions spectrum of Figure 9.18.

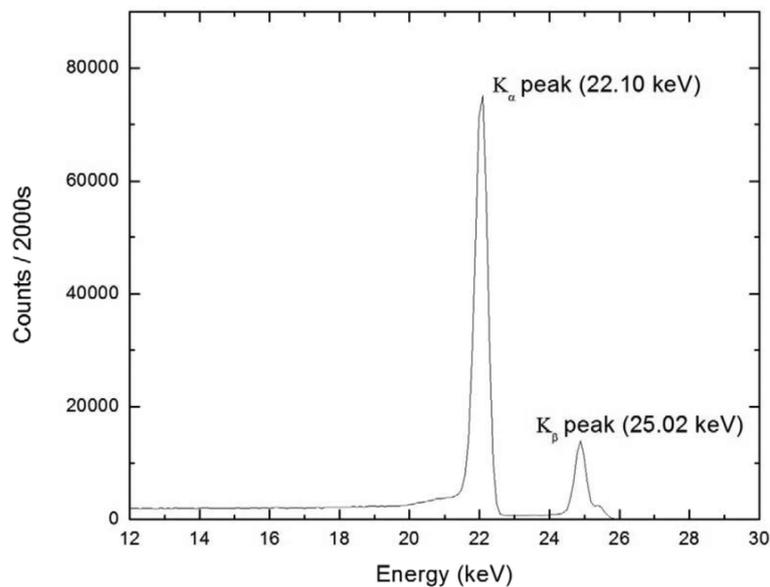
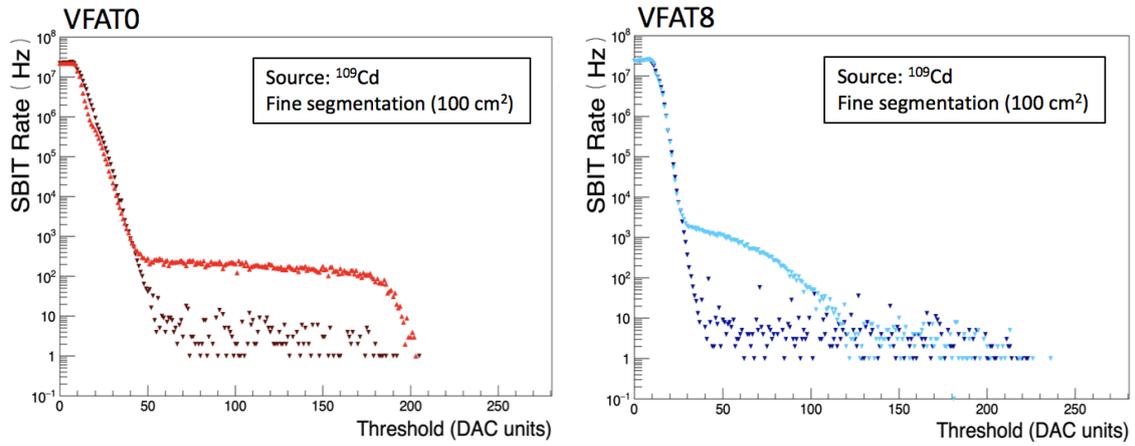


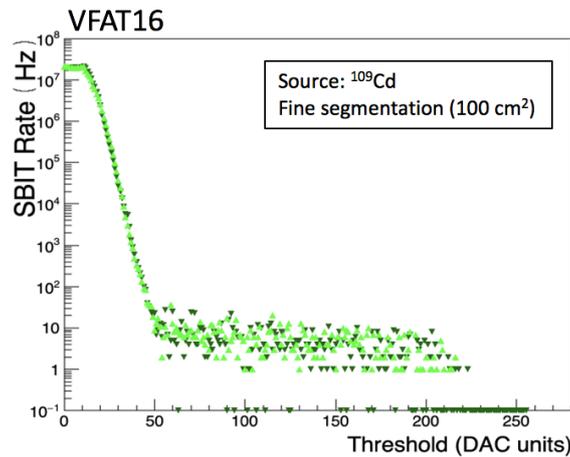
Figure 9.18 – Emissions spectrum of ^{109}Cd

Once the source was installed, an S-bit threshold “OR” scan was taken - chosen as the faster alternative to a per-channel S-bit scan, which would take several hours to run (see Section 3.4.3 for more details). The scans from the three VFATs in the same $i\eta$ section of the detector as the source - VFATs #0, #8, and #16 - are shown in Figure 9.19.



(a) S-bit threshold “OR” scan of VFAT #0

(b) S-bit threshold “OR” scan of VFAT #8



(c) S-bit threshold “OR” scan of VFAT #16

Figure 9.19 – S-bit threshold “OR” scans of VFATs #0, #8, and #16 in the presence of a ^{109}Cd source installed at VFAT #0

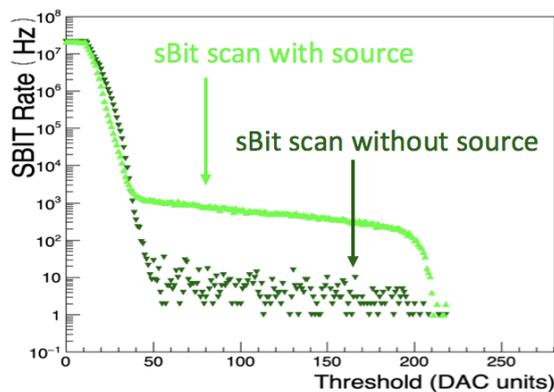
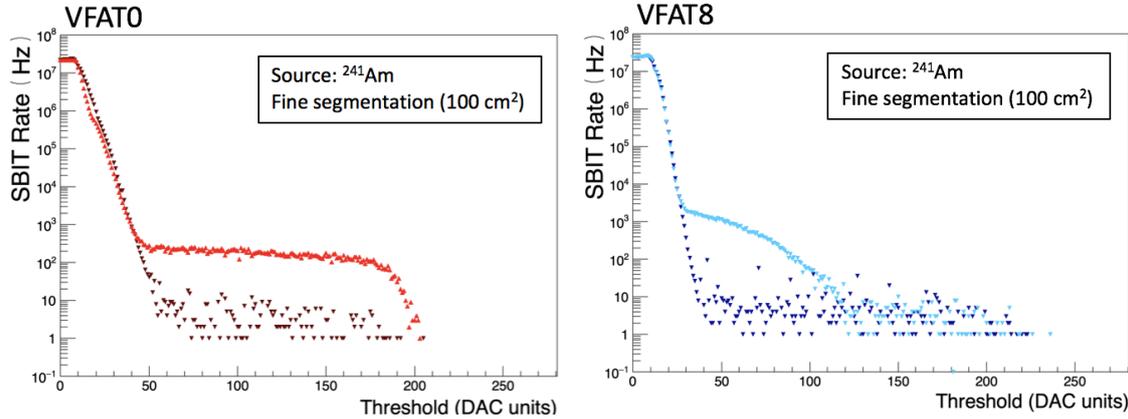


Figure 9.20 – Descriptive explanation of the combination “OR” S-bit threshold scan plots

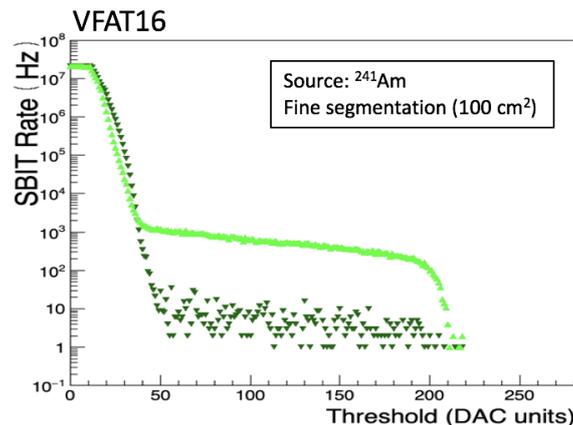
For clarity and conciseness, these figures are presented as a combination of the sourceless baseline measurement and the measurement with the source present. The sourceless baseline measurement is given as the more granular, darker colored, downward-facing triangles, and the source measurement is given as the more solid, lighter-colored, upward-facing triangles. This convention, labeled in Figure 9.20, is maintained throughout all of the following images of the same type.

From these scans, we see that the influence of the ^{109}Cd source rapidly decreases with distance from the source itself, to the point of having virtually no effect on the scan of VFAT #16, just two VFAT positions away from the source. This is expected, as the emissions from the ^{109}Cd fall into the X-ray/MIP category, as described in Section 9.2, and so in this case we did not expect to see any negative-polarity reflected signal that could be interpreted by the electronics as true signal.

The ^{109}Cd was then replaced with an ^{241}Am source, in order to provide a source of α particles. Once again, an S-bit threshold “OR” scan was taken, with the results for those same three VFATs. However, as seen in Figure 9.21, the results for α radiation are significantly different here than they were for x-rays.



(a) S-bit threshold “OR” scan of VFAT #0, fine segmentation (b) S-bit threshold “OR” scan of VFAT #8, fine segmentation



(c) S-bit threshold “OR” scan of VFAT #16, fine segmentation

Figure 9.21 – S-bit threshold “OR” scans of VFATs #0, #8, and #16 in the presence of an ^{241}Am source installed at VFAT #0

In this case, we see that the cross-talk signal has a significant presence in both VFAT #8 and VFAT #16. Interestingly and in contrast to the results with ^{109}Cd , the signal appears to have a much greater presence in VFAT #16 than VFAT #8, despite VFAT

#8 being closer to the source. While it was initially thought that this may be the result of an increased capacitive coupling at the edges of the detector, preliminary research into this suggests that it may in fact be due to the so-called “basin effect”, as this issue is not seen in GE2/1 chambers. This basin effect is the non-uniformity of a GE1/1 detector’s gas gain as a result of the unwanted bending of the PCBs of the detector, where the drift board has a bending of up to 1 mm and the readout board has the same bending but in the opposite direction. This difference is most pronounced in the center of the chamber, where the altered gap height, and the effect it has on the signal, is the most profound.

As the cross-talk signal is clearly seen using the ^{241}Am source, this makes it an ideal candidate for comparison with the results of a larger-segmented HV partition.

With the baseline measurements done, modifications were then made to a GE1/1 detector. As discussed in Section 9 and shown in Figure 9.1, the HV sectors of a GEM foil are separated from each other by a $10\text{ M}\Omega$ resistor which is soldered onto the foil. The simplest way of studying different partition sizes within a single detector, then, is to replace these $10\text{ M}\Omega$ resistors with $0\ \Omega$ resistors on a portion of the foil, in order to simulate those sectors acting as a single HV partition. This replacement can be seen in Figure 9.22, where the surface-mount device (SMD) resistor on the left of the image is the standard $10\text{ M}\Omega$ protection resistor, and the SMD resistor on the right is a replacement $0\ \Omega$ resistor.



Figure 9.22 – Magnified view of a GE1/1 foil showing a $10\text{ M}\Omega$ resistor (left) and $0\ \Omega$ resistor (right) [119]

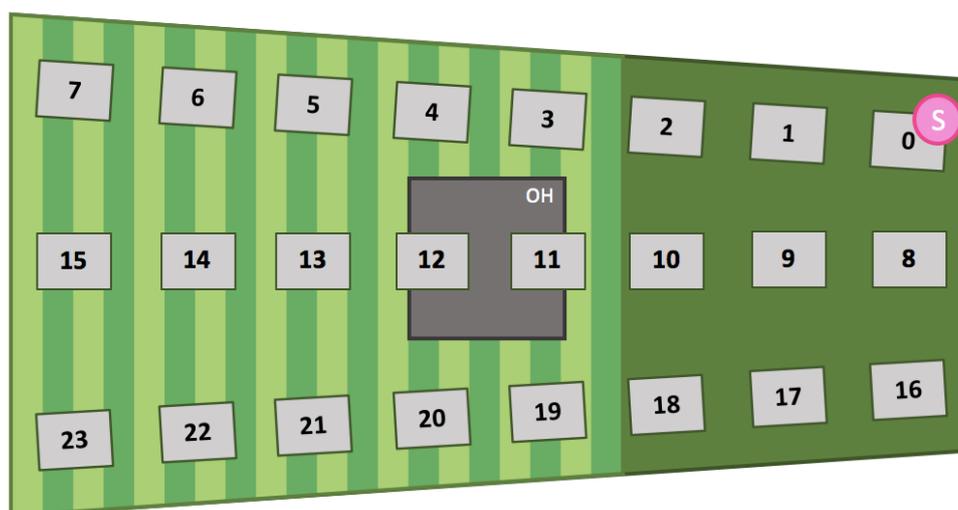


Figure 9.23 – Figure illustrating the prototype GE1/1 detector with modified HV partitions and location of the signal-triggering radioactive source

The foil of the GE1/1 v3 prototype detector was modified in this way, with the ten HV partitions closest to the narrow end of the detector connected via $0\ \Omega$ resistors, giving an area approximately equal to that of a $30 \times 30\ \text{cm}^2$ detector. The remaining twenty partitions, closer to the wide end of the detector, were left unmodified. This set-up can be seen in Figure 9.23.

With these modifications in place, the ^{241}Am source was replaced and another set of S-bit threshold scans was taken. The full scanset is shown in Figure 9.24, with the VFATs affected by the large HV partition boxed in pink, and the location of the source labeled with an “S”. For a more direct comparison, we can look at the S-bit threshold scans of VFATs #0, #8, and #16, and compare them to the fine segmentation scan for the same VFAT. These scans can be seen in Figures 9.25b, 9.26b, and 9.27b.

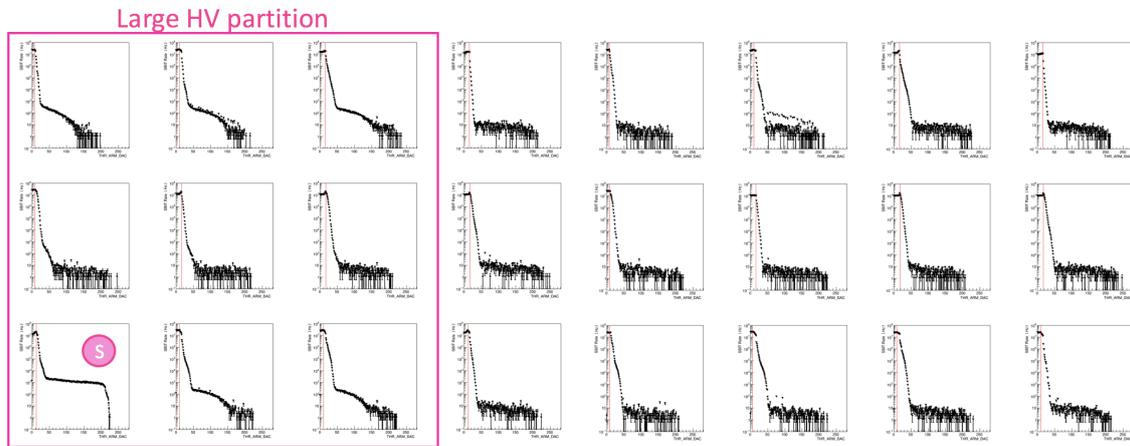
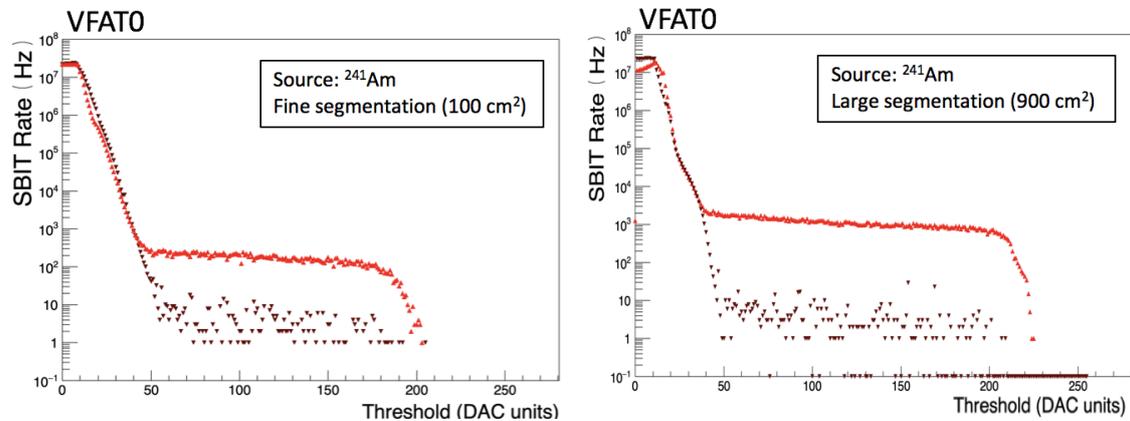
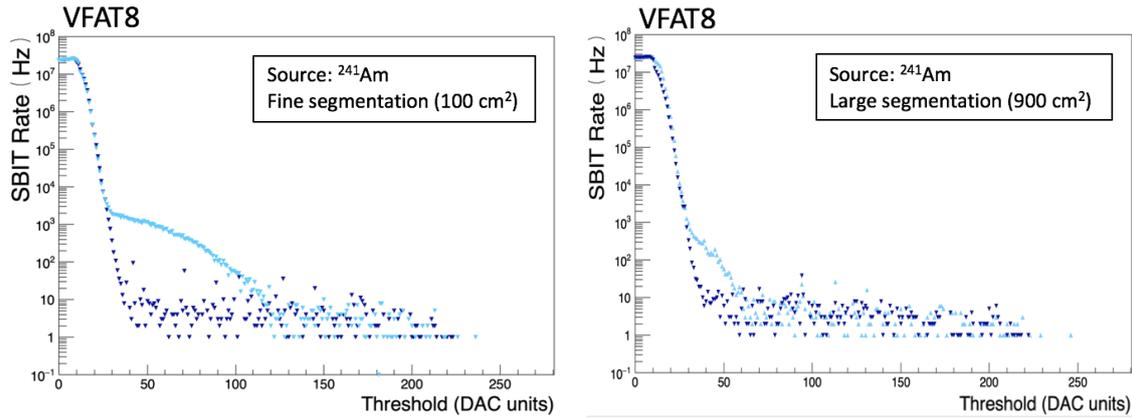


Figure 9.24 – Full set of S-bit threshold scans for the modified GE1/1 detector with an ^{241}Am source near VFAT #0



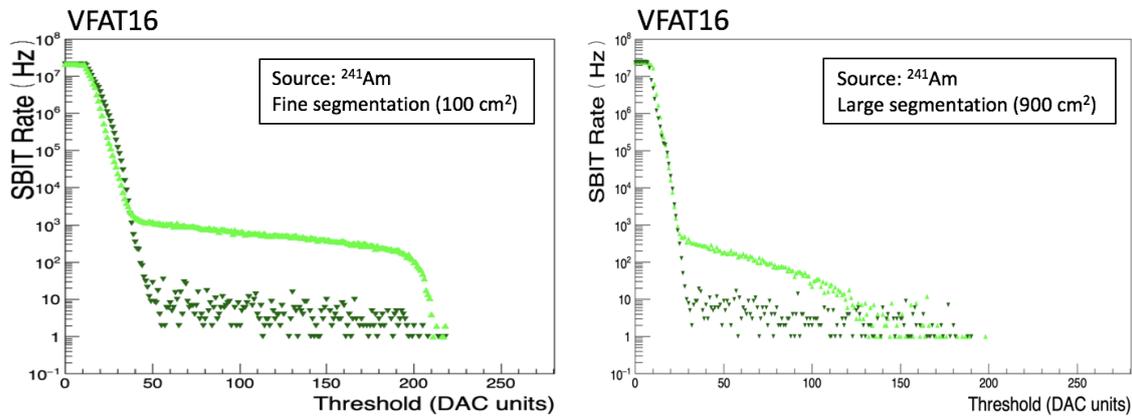
(a) S-bit threshold “OR” scan of VFAT #0, fine segmentation (b) S-bit threshold “OR” scan of VFAT #0, large segmentation

Figure 9.25 – S-bit threshold “OR” scans of VFAT #0 in the presence of an ^{241}Am source installed at VFAT #0



(a) S-bit threshold “OR” scan of VFAT #8, fine segmentation (b) S-bit threshold “OR” scan of VFAT #8, large segmentation

Figure 9.26 – S-bit threshold “OR” scans of VFAT #8 in the presence of an ^{241}Am source installed at VFAT #0



(a) S-bit threshold “OR” scan of VFAT #16, fine segmentation (b) S-bit threshold “OR” scan of VFAT #16, large segmentation

Figure 9.27 – S-bit threshold “OR” scans of VFAT #16 in the presence of an ^{241}Am source installed at VFAT #0

Here, we can look at VFATs #8 and #16 and see a significant reduction in the cross-talk signal for the large segmentation when compared against that of the fine segmentation, with the maximum cross-talk amplitude reduced by approximately a factor of 2. Although this reduction is significant and highly desired, the issue remains that a larger segmentation also comes with a higher probability of discharge propagation. A balance, therefore, needs to be struck between reduction of the cross-talk signal and reduction of the discharge propagation probability.

9.3.1 GE2/1 Studies

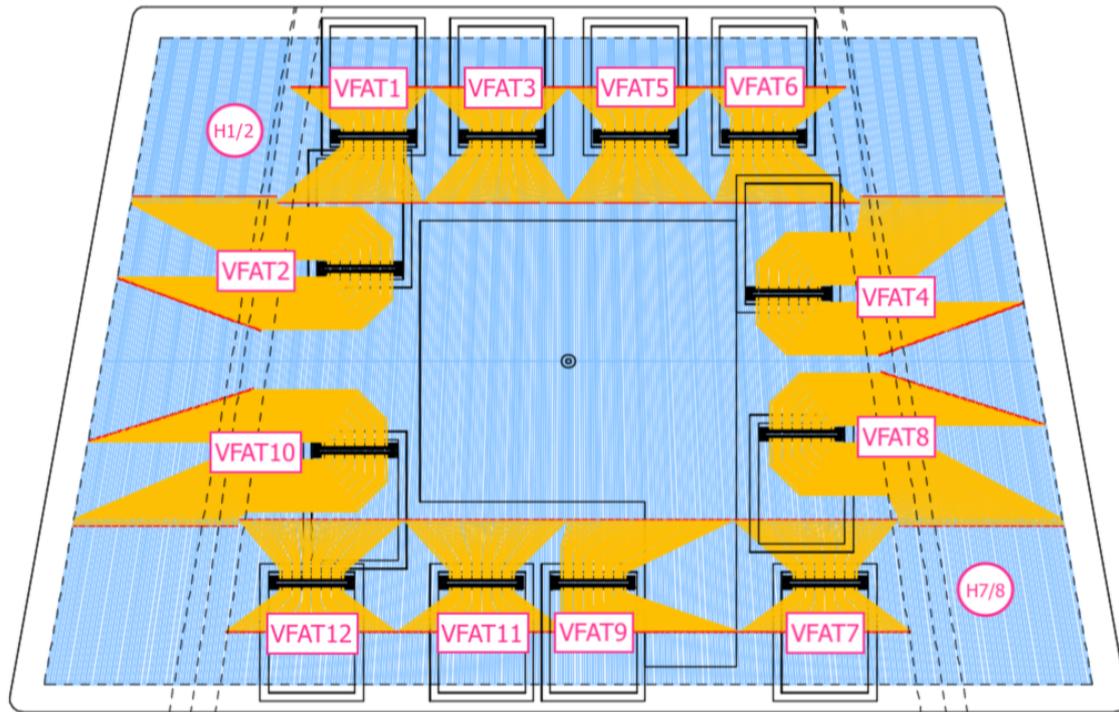


Figure 9.28 – Schematic of the GE2/1 M5 readout board, with VFAT positions labeled [103].

As the usage of double-segmented GEM foils is foreseen only for future GEM projects (See Chapter 8 for more information), with GE1/1 installation well underway at the time of the writing of this dissertation, it logically followed that further studies into these cross-talk signals should be made on GE2/1 and ME0 chambers rather than GE1/1 chambers. Thus, a GE2/1 M5 module, the smallest of the four modules which make up a GE2/1 detector, was procured. This module was chosen not only for its availability, but for the fact that, as the smallest of the four GE2/1 module sizes (along with its sister, the M1 module: see Section 8.1 for more details), it represents the largest HV partition possible for the system.

The technical schematic for this module can be seen in Figure 9.28, taking careful note that the numbering of the VFATs is not necessarily intuitive in nature. Two mm-diameter holes were drilled into the drift board and covered with food-grade plastic wrap and Kapton tape in order to maintain gas tightness, as with the previous study. The approximate location of these holes are labeled in Figure 9.28 according to their closest neighboring VFATs as “H1/2” and “H7/8”. The chamber was operated at a divider current of $I = 700 \mu\text{A}$ and the HV settings found in Table 9.2.

In order to get a baseline measurement of the chamber, an S-bit threshold scan (See Section 3.4.3 for more details) was taken without the presence of a radioactive source. This scan can be seen in Figure 9.29, with each of the VFATs labeled according to Figure 9.28. Unfortunately, the M5 chamber used was a prototype chamber with several known issues, poorly-responsive VFAT #1 position, so this position must not be considered as having any scientific merit when discussing study results.

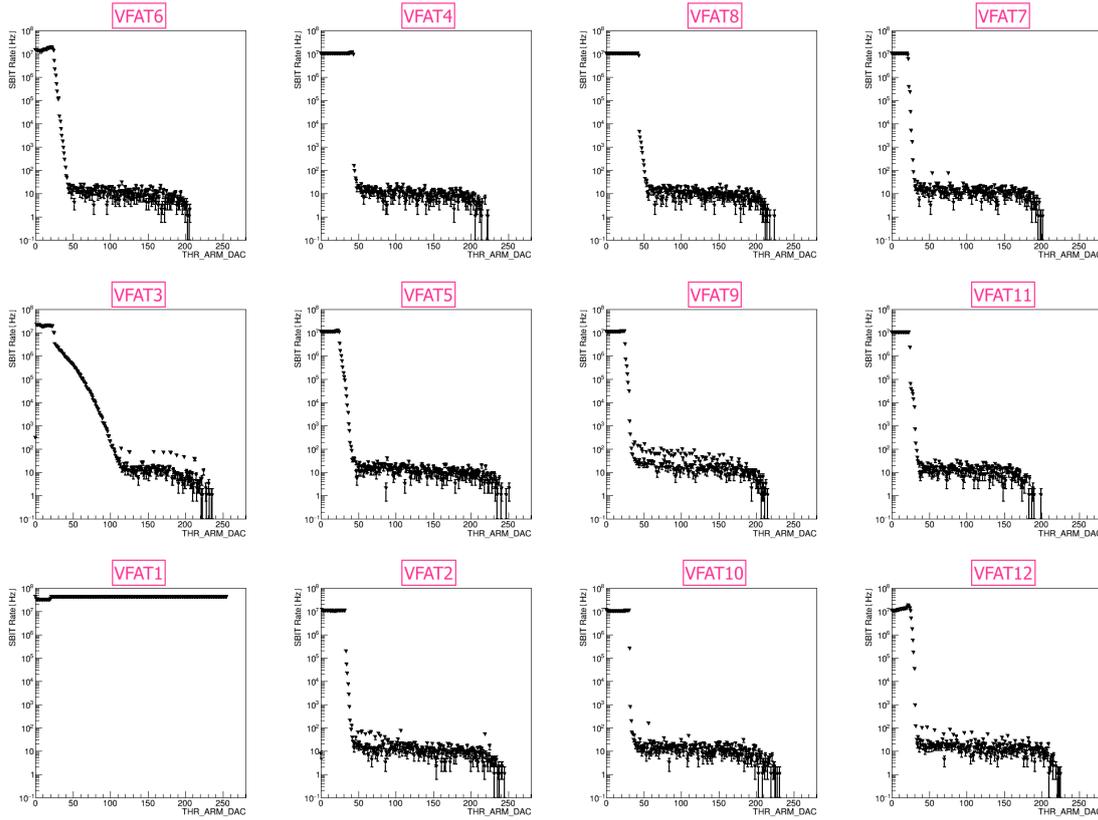


Figure 9.29 – Baseline S-bit threshold scan for the modified double-segmented GE2/1 M5 chamber

Detector Layer	V_0 (V)
Drift	787.50
G1 _{top}	392.00
G1 _{bottom}	306.60
G2 _{top}	385.00
G2 _{bottom}	612.50
G3 _{top}	367.50
G3 _{bottom}	437.50

Table 9.2 – HV settings for the GE2/1 M5 dual-polarity cross-talk study

A ^{109}Cd source was then taped to the bottom of the detector at the H1/2 position to serve as an x-ray source. Another S-bit scan, seen in Figure 9.30, was then taken at the same HV settings as the baseline scan. Here, the presence of the x-ray emissions from the ^{109}Cd source is clearly seen by VFAT #2, with the plateau of the S-bit rate increasing by approximately three orders of magnitude, but the remainder of the S-bit threshold scans remain largely unchanged. This is in line with what was seen in the previous sections, where only the closest readout sector of the same HV partition as the ^{109}Cd source sees any cross-talk signals.

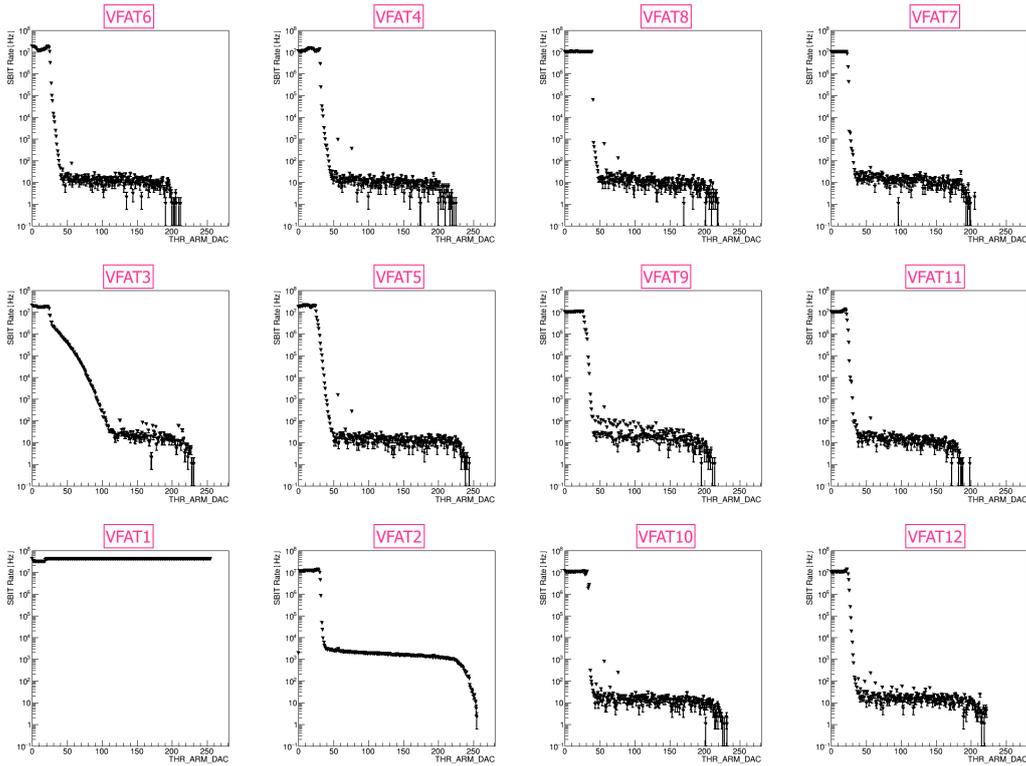


Figure 9.30 – S-bit threshold scan for the GE2/1 M5 chamber with ^{109}Cd source at position H1/2

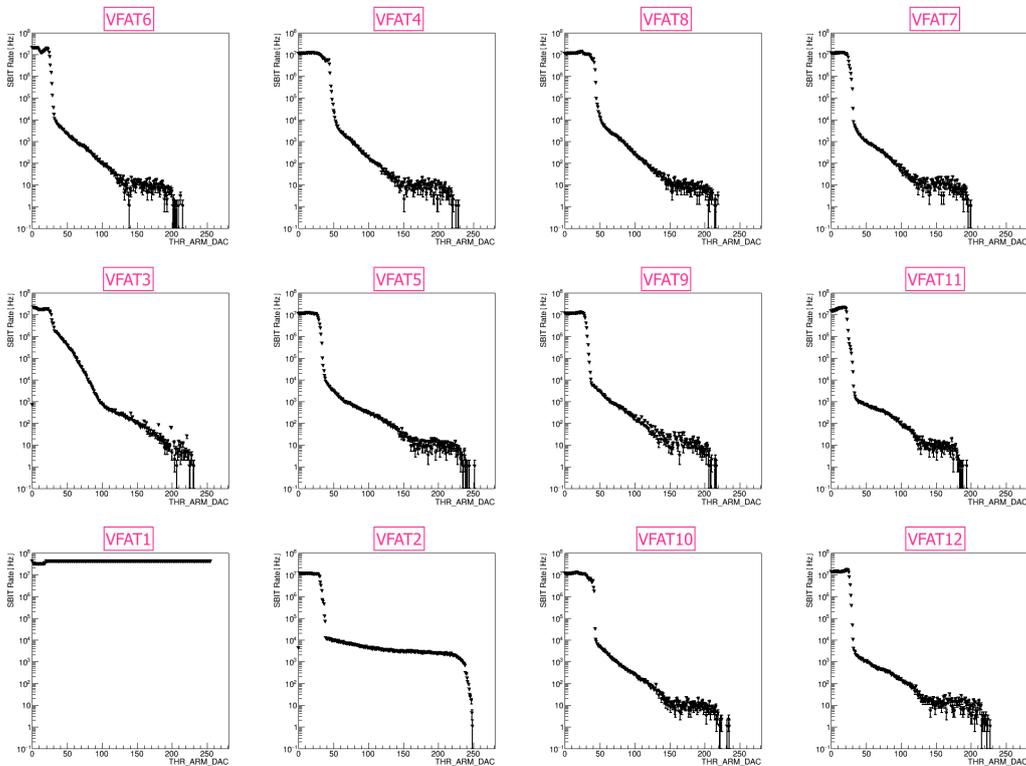


Figure 9.31 – S-bit threshold scan for the GE2/1 M5 chamber with ^{241}Am source at position H1/2

One important thing to note here is the curve that exists between the noise plateau

and signal in VFAT #3. It is tempting to see this as a cross-talk contribution, given its similarity to the shape of the contribution seen in Figure 9.27. However, this is not the case in this instance. Referring back to the baseline scan of Figure 9.29, we see that this is merely a quirk of this particular prototype detector and can safely be ignored for the purposes of measuring the cross-talk signal.

The ^{109}Cd source was then replaced with an ^{241}Am source, the same used in Section 9.1, in order to provide a source of α particles. Another S-bit scan, seen in Figure 9.31, was then taken at the same HV settings as the previous two scans.

Unlike with the ^{109}Cd source, the ^{241}Am source is not only seen directly in VFAT #2, but all of the other VFAT positions are seen to be affected as well, with a heightened S-bit rate between the noise plateau and the cosmic muon plateau at higher THR_ARM_DACs. This is evidence that the dual-polarity cross-talk signal is being seen in all VFAT sectors of the detector, and represents a significant roadblock for the use of double-segmented foils in GE2/1: previously, it was determined that enlarging the size of the HV partition can lower the presence of the cross-talk signal. But the M5 module, by nature, has the smallest size limit of this partition increase. With the presence of the cross-talk signals seen in all of the readout sectors in this test scan, we are left with the necessity of pursuing other mitigation strategies.

In order to perform a test of one such mitigation strategy, a 3.3 nF blocking capacitor was installed on the detector between G3 and ground, and another S-bit threshold scan was taken. The results, seen in Figure 9.32, were not promising, with the effect being seen at even higher THR_ARM_DACs than without the capacitor.

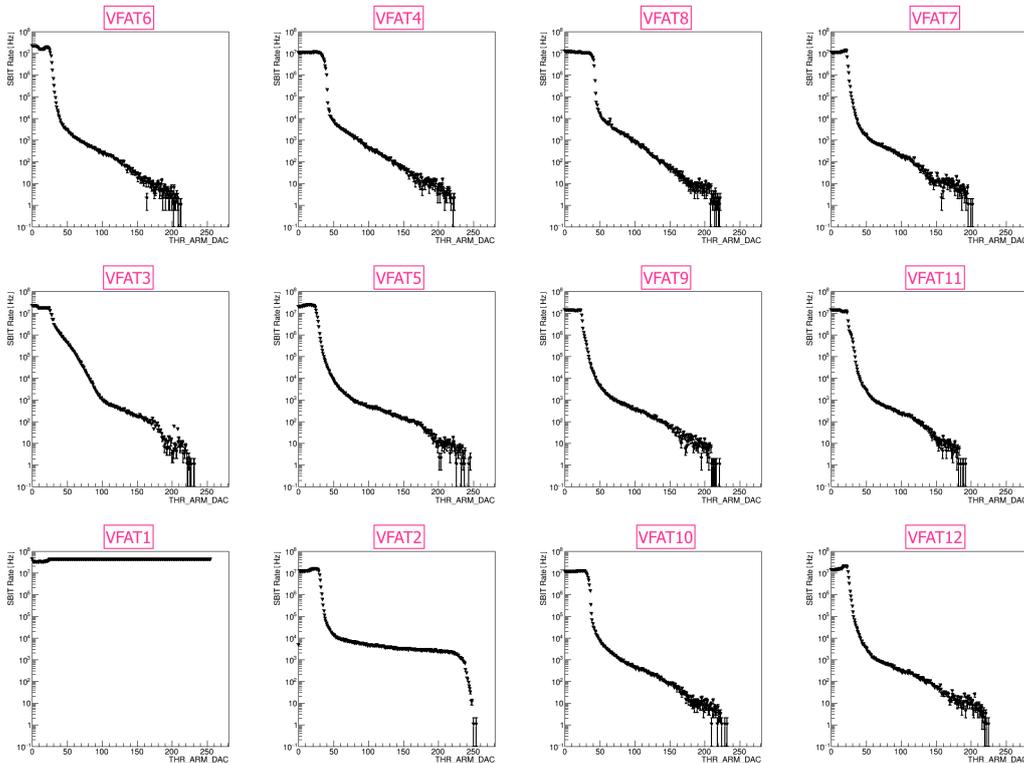


Figure 9.32 – S-bit threshold scan for the GE2/1 M5 chamber, with a 3.3 nF blocking capacitor installed between G3 and GND, with ^{241}Am source at position H1/2

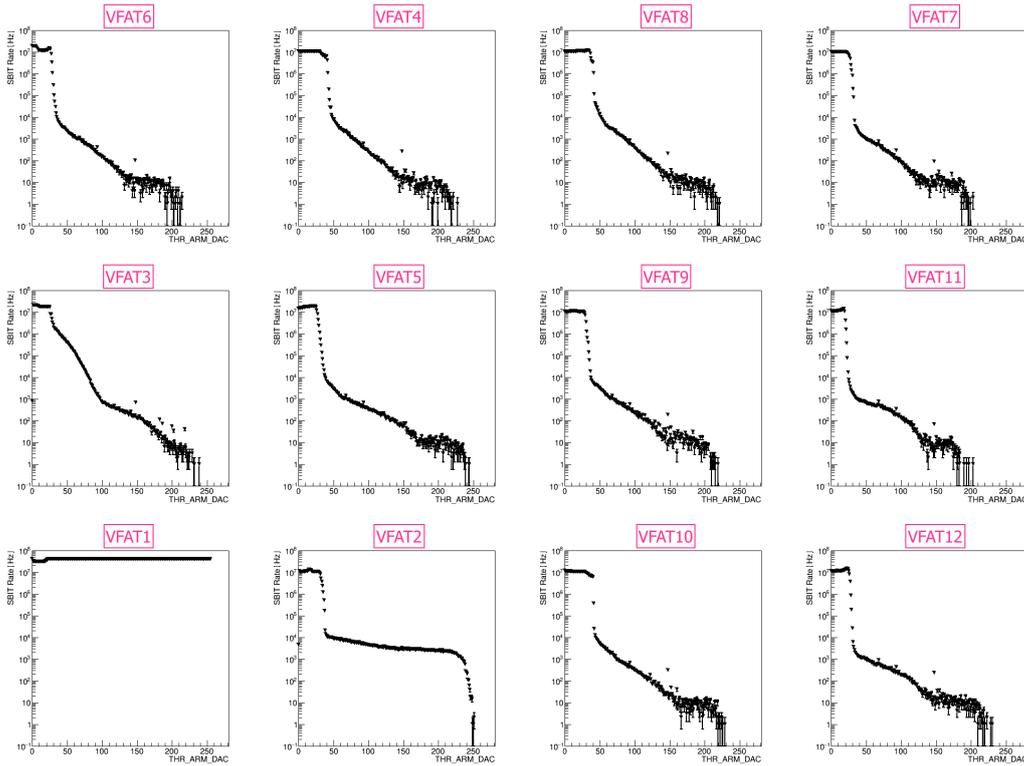


Figure 9.33 – S-bit threshold scan for the GE2/1 M5 chamber, with a 3.3 nF blocking capacitor and 100 k Ω resistor installed between G3 and GND, with ^{241}Am source at position H1/2

The detector was then modified one final time, with a 100 k Ω resistor installed between the 3.3 nF blocking capacitor and ground. Again, an S-bit threshold scan was taken, with the results seen in Figure 9.33. Although this did ameliorate the problem to an extent, it was not nearly enough to mask the problem to a workable degree.

Unfortunately, at this point, the outbreak of the COVID-19 pandemic forced the shutdown of all CERN laboratories, and as a result, the study was prematurely ended here. Nevertheless, we can still draw some conclusions from the data that we were able to obtain. From the scans performed with the ^{241}Am source present, we see that the cross-talk within the detector as a result of HIPs is ubiquitous, appearing in every VFAT. This poses a severe threat to the operational uptime of the readout electronics.

Take, for example, the example signals shown in Figure 9.34. We first have the detection of a HIP, shown in pink, as it exceeds the discrimination threshold. Then, after some time, this signal passes back below the discrimination threshold as it transitions into its undershoot, before eventually returning to the signal baseline. This signal is mirrored into the cross-talk signal, shown in blue. The time in which the amplitude of the cross-talk signal is above the discrimination threshold represents time in which no true signal can be detected, regardless of whether or not any particles are passing through that readout sector. This is the so-called “blind deadtime” induced by the cross-talk signal, and represents the time in which the VFAT is totally blind to any new signal in that location.

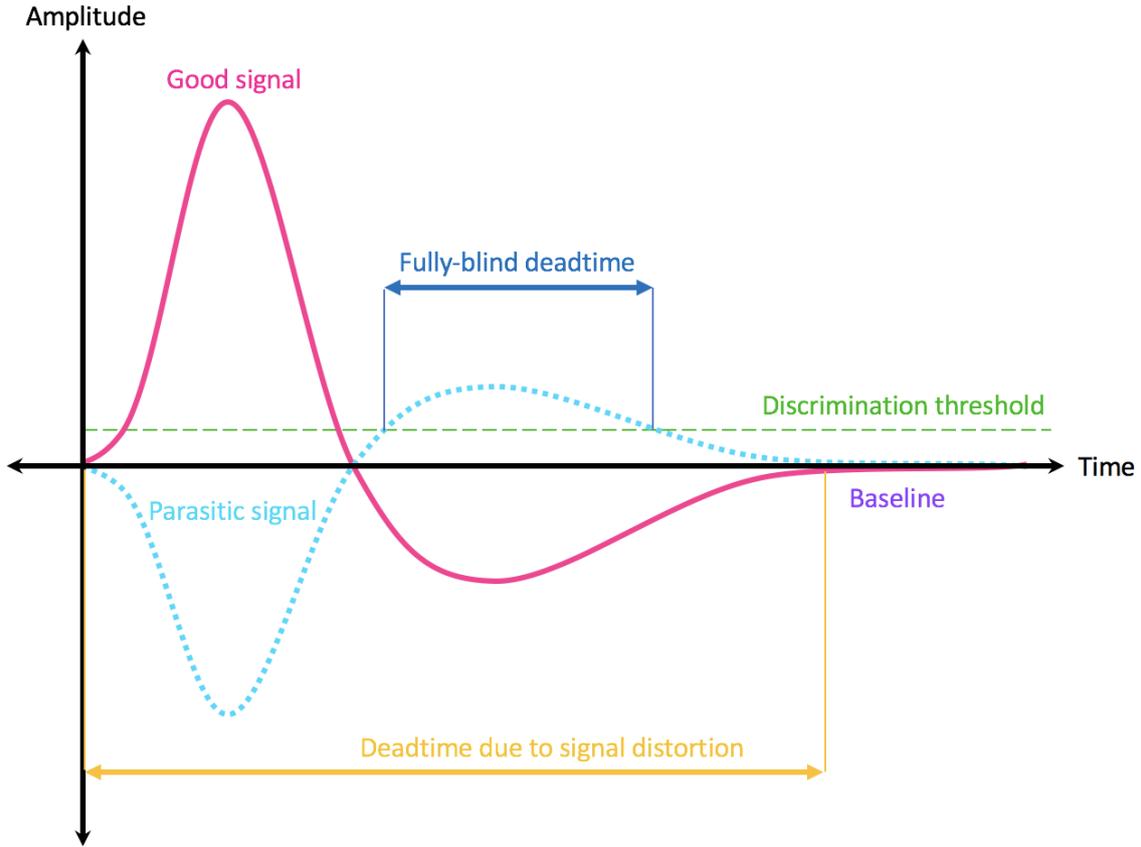


Figure 9.34 – Example of a parasitic cross-talk signal, as seen from the analog output of the VFAT3 chip’s preamplifier, inducing deadtime into the readout electronics

However, this does not represent the sum total of deadtime introduced by the cross-talk signal. If, for example, a MIP was to be detected during the same time as the initial negative-polarity part of the cross-talk signal, even if it were to pass the discrimination threshold, the cross-talk signal would distort it to the point of it becoming unusable. It would be delayed in time from its true arrival time such that it would be impossible to assign it to the correct bunch crossing, and therefore it would be useless for particle reconstruction. Thus, we must consider the deadtime induced by the cross-talk signal to be the entire length of the signal, from the time it first appears until the time it returns to the baseline.

In general, for a given signal, the maximum amount of time to return to this baseline, and thus the maximum time it takes for the VFAT to return to a state in which it can detect a new signal, is 500 ns, equivalent to 20 LHC bunch crossings. So these cross-talk signals can each induce 500 ns of deadtime into the system. And, in the case of HIPs, in which the cross-talk signal is seen in every VFAT on the detector, this equates to 500 ns of deadtime during which the entire detector is essentially blind to physics.

In order to determine how much total deadtime this introduces into the GE2/1 system, we can look at the rate of HIPs expected in that area of CMS. Figure 9.35 shows the rate of HIPs of $E_{deposit} > 30$ keV at the location of the GE2/1 system, mostly protons and nuclei, as simulated in GEANT. This rate is shown by the stepped red line, with each step representing the location of one η partition of a GE2/1 chamber. At its highest, this rate is $\sim 200 \frac{\text{kHz}}{i\eta}$.

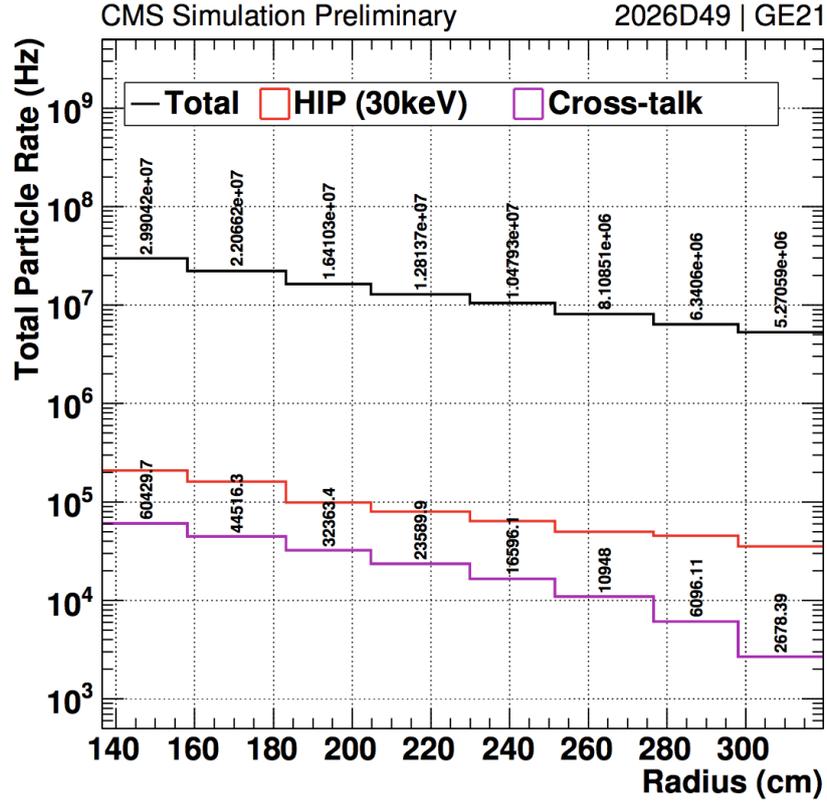


Figure 9.35 – GEANT-simulated rate of HIPs of $E_{deposit} > 30$ keV at the location of the GE2/1 system and resulting cross-talk rate in GE2/1 [102]

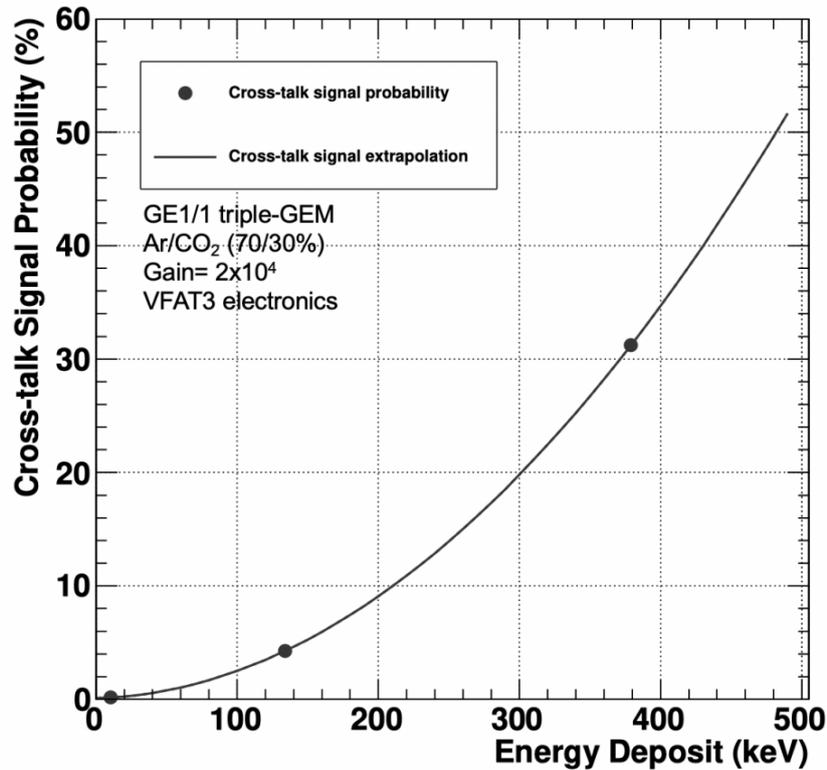


Figure 9.36 – Experimental probability of induced cross-talk in a GE1/1 detector per amount of energy deposited by the HIP [102]

In order to obtain a predicted cross-talk signal rate, this hit rate was then convoluted with the probability of inducing a cross-talk signal within the detector. This probability is dependent on the energy deposited by the HIP, and is shown in Figure 9.36.

These values were obtained using the setup shown in Figure 9.16, with one or more PVC attenuators placed between the detector and the alpha source in order to obtain a range of deposited energies. In order to get the probability of inducing a cross-talk signal, the S-bit rate with the source present was compared to the S-bit rate without the source. This comparison was made in two different ways: the first, shown in orange in Figure 9.37, looked at the cross-talk rate at a threshold which corresponded to 100 Hz of noise (ΔR_1), while the second, shown in pink in Figure 9.37, looked at the cross-talk rate at a fixed threshold of 100 DAC units (ΔR_2). Both methods resulted in consistent values.

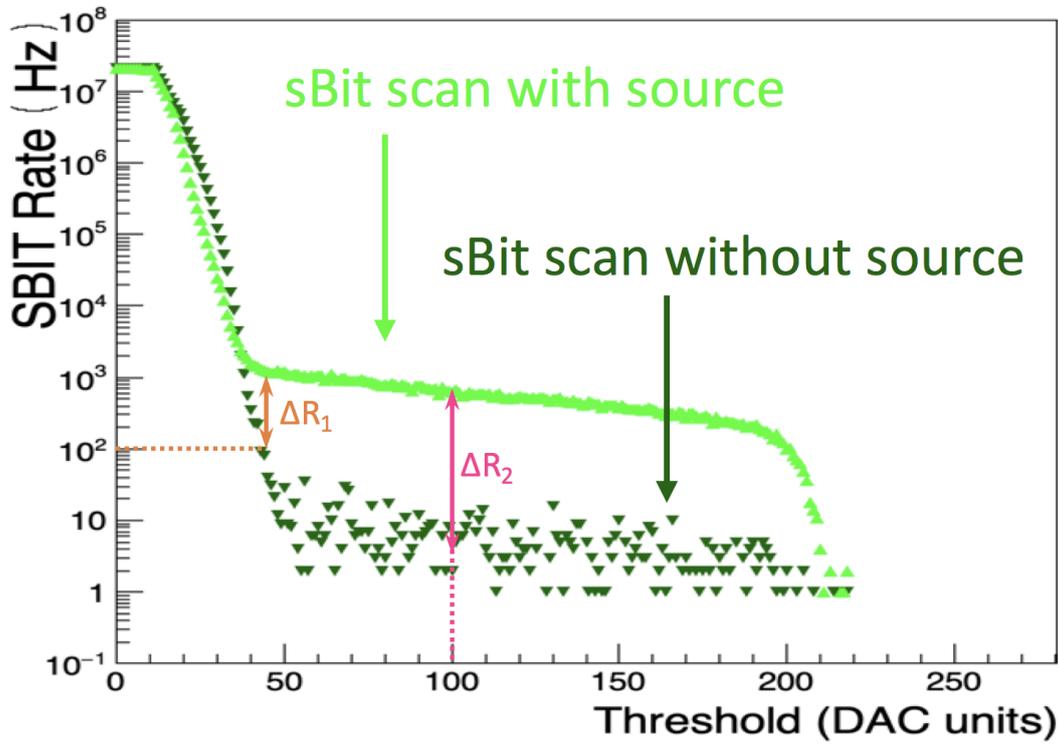


Figure 9.37 – Measurement of the cross-talk rate in a GE1/1 detector as compared to the sourceless baseline S-bit rate measurements

The resulting predicted cross-talk signal rate is shown in Figure 9.35 by the purple line, once again with each step representing a single GE2/1 η partition. We see here, then, that in the region of lowest background, the cross-talk rate is simulated to be $2 - 3 \frac{kHz}{i\eta}$, whereas in the region of highest background, this rate skyrockets to $60 \frac{kHz}{i\eta}$.

Taking this maximum value of 60 kHz and multiplying it by the maximum deadtime of the VFAT3 chip per cross-talk signal, we get a maximum overall deadtime of $(60 \cdot 10^3 \text{ Hz}) \cdot (500 \cdot 10^{-9} \text{ s}) = 0.03$, or 3% overall deadtime.

While this is unacceptably high, efforts are currently ongoing to further reduce this inefficiency, and it should be noted that this 3% is likely an overestimation of the deadtime. In addition to utilizing the maximum values for both the rate of cross-talk and the deadtime of the VFAT3 itself, it also utilizes high values of $E_{deposit}$. From Figure 9.36, we see that the cross-talk signal probability was considered over a range of energies, up to $E_{deposit} = 380 \text{ keV}$. However, it is expected that HIPs within the CMS detector will only have $E_{deposit}$

≈ 100 keV, and since the probability is $E_{deposit}$ -dependent, this should decrease the overall deadtime.

9.3.2 ME0 Studies

Parallel to the GE2/1 studies described in Section 9.3.1, a separate study with ME0 was designed which would utilize carefully-controlled signals from a pulse generator rather than a natural radioactive source. To this extent, an experimental set-up, seen in Figures 9.39, was put together using a double-segmented ME0 chamber: the prototype chamber ME0-CERN-0002-proto, with $5\text{ M}\Omega$ protection resistor foils. Each readout sector was connected to a flexPCB and a Panasonic-to-LEMO adapter which was terminated with a $50\ \Omega$ resistor. The output of a Wavetek 395 100 MHz synthesized arbitrary waveform generator was connected via LEMO cable to the $i\eta, i\phi = 8,3$ readout sector of the board, chosen to mimic the sister set-up located at Florida Institute of Technology (FIT), and split to display on channel 1 of a Tektronix TDS 2024C four channel digital storage oscilloscope, where it was terminated with a $50\ \Omega$ resistor. Channel 2 of the oscilloscope was then connected to a given readout sector and also terminated with a $50\ \Omega$ resistor.

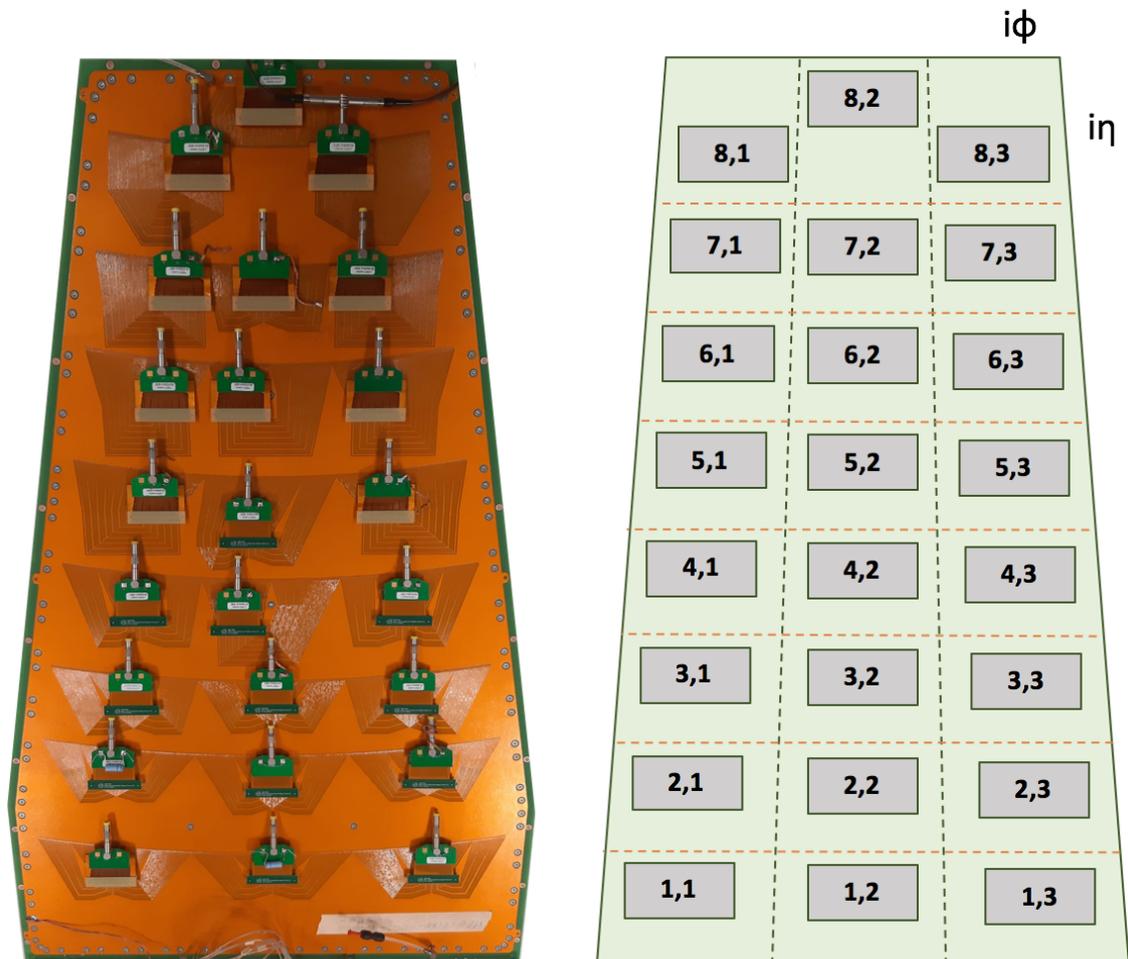


Figure 9.38 – ME0 detector and corresponding $i\eta, i\phi$ readout sector labels

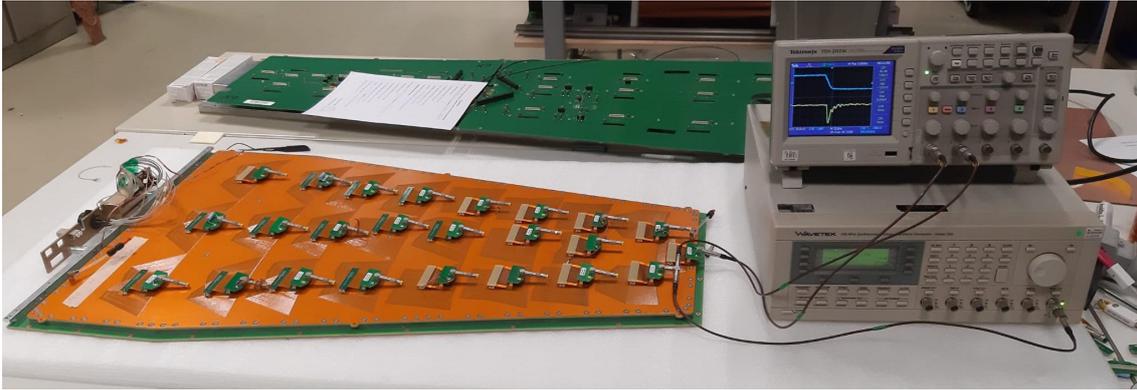


Figure 9.39 – Initial experimental set-up for the double-segmented ME0 cross-talk study

Due to the closure of CERN during the early days of the COVID-19 pandemic, the study was put on hold for several months. In that time, a Digilent Analog Discovery 2 USB oscilloscope and BNC adaptor board became available. As a result, after the reopening of the laboratory, the initial set-up was replaced with the set-up seen in Figure 9.40, where the TDS 2024C was replaced with the Analog Discovery 2 for more precise measurements. More thorough grounding measures were also put into place, grounding each of the Panasonic-to-LEMO adapters to the shared main ground of the laboratory, in order to reduce any signal pick-up within the detector.

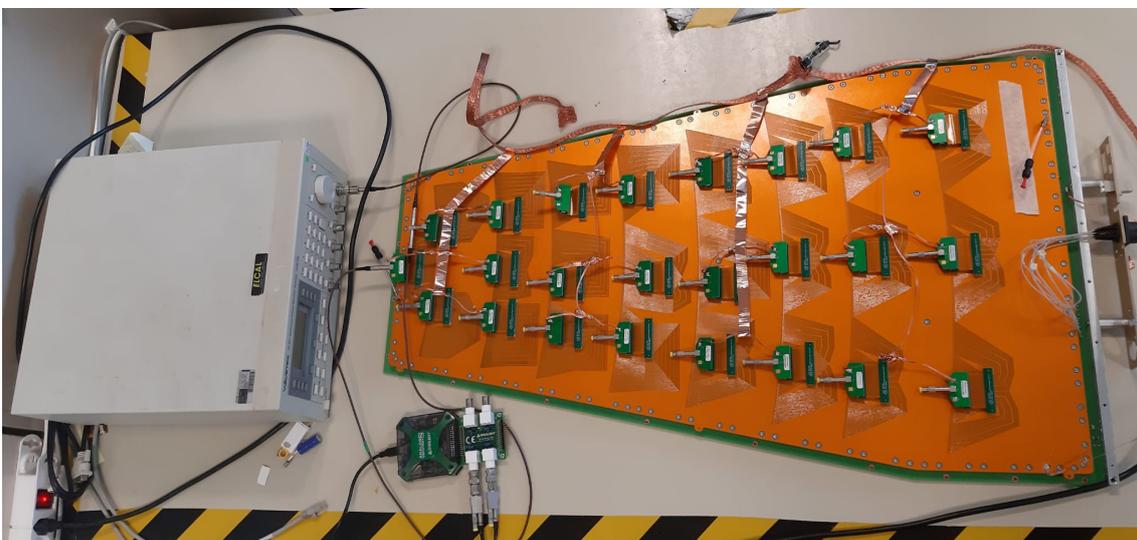


Figure 9.40 – Modified experimental set-up for the double-segmented ME0 cross-talk study

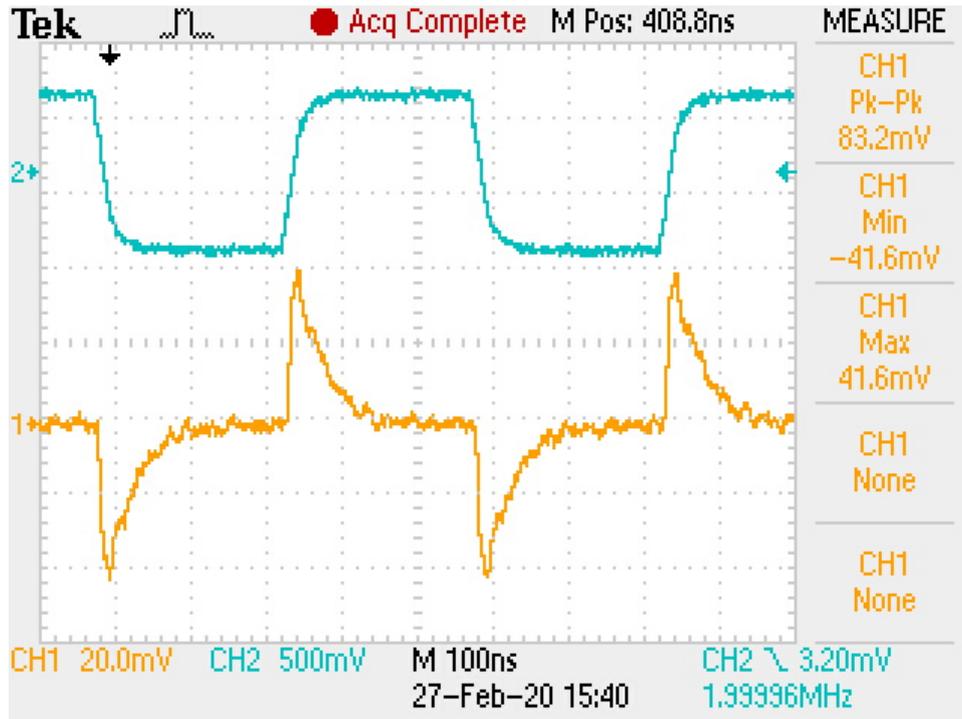


Figure 9.41 – Example Tektronix TDS 2024C readout of a cross-talk pulse resulting from a 550 mV, 2 MHz square wave in sector 7,3

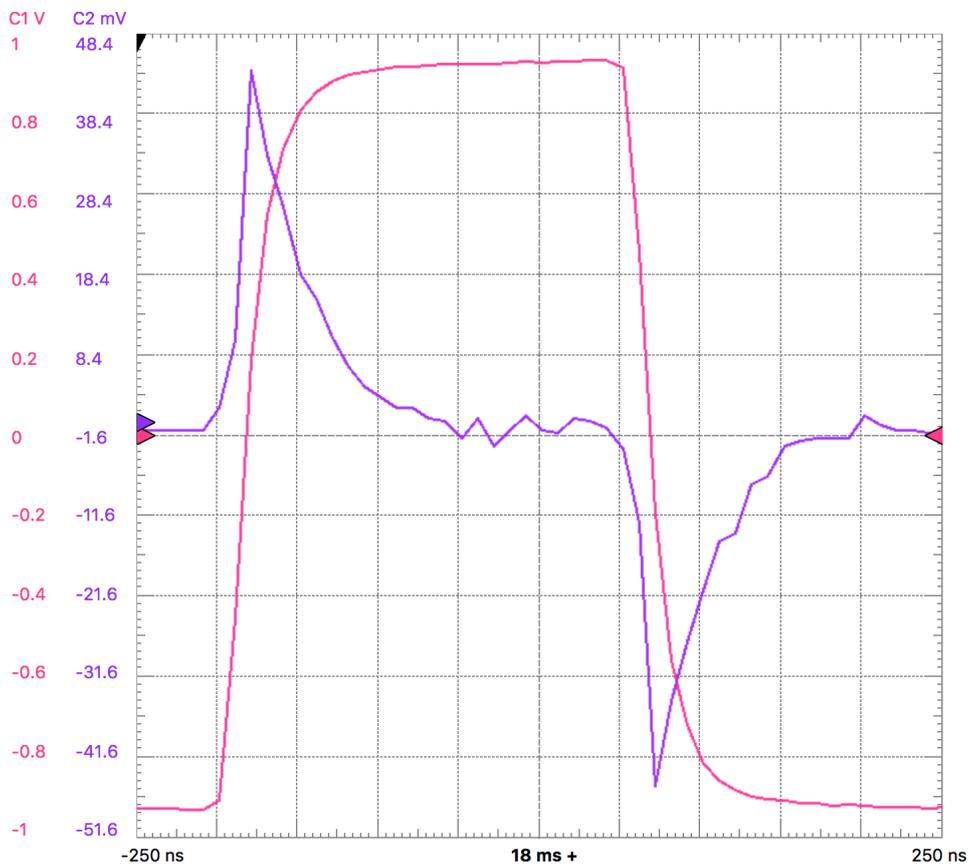


Figure 9.42 – Example Digilent Analog Discovery 2 readout of a cross-talk pulse resulting from a 1 V, 2 MHz square wave in sector 8,2

For each measurement, a 2 MHz square wave of a given amplitude was outputted from the Wavetek 395 generator and injected into $i\eta, i\phi = 8,3$. An example of the incoming square wave and resulting cross-talk pulse in the initial Tektronix set-up can be seen in Figure 9.41, and an example from the Digilent set-up can be seen in Figure 9.42. Interestingly, it is figures such as these that led us to stop referring to the cross-talk signal as “dual-polarity” cross-talk signals, as here we see that the positive-amplitude cross-talk pulse follows the rising edge of the square wave, and the negative-amplitude cross-talk pulse follows the falling edge of the square wave. Of course, physicists are creatures of habit (see: conventional current and other idiosyncrasies of physics), so this misnomer is likely to be found in other existing and future studies on the subject.

The pulse amplitude of the resulting cross-talk pulse was then recorded at each of the other readout sectors, and normalized to a percentage of the incoming pulse amplitude. The resulting plot can be seen in Figure 9.43. Note that in this figure, only 7 of the 23 possible readout locations are shown. This is due to the fact that in the other 16 locations, any potential cross-talk signal was indistinguishable from the baseline noise in the chamber, and is therefore not considered.

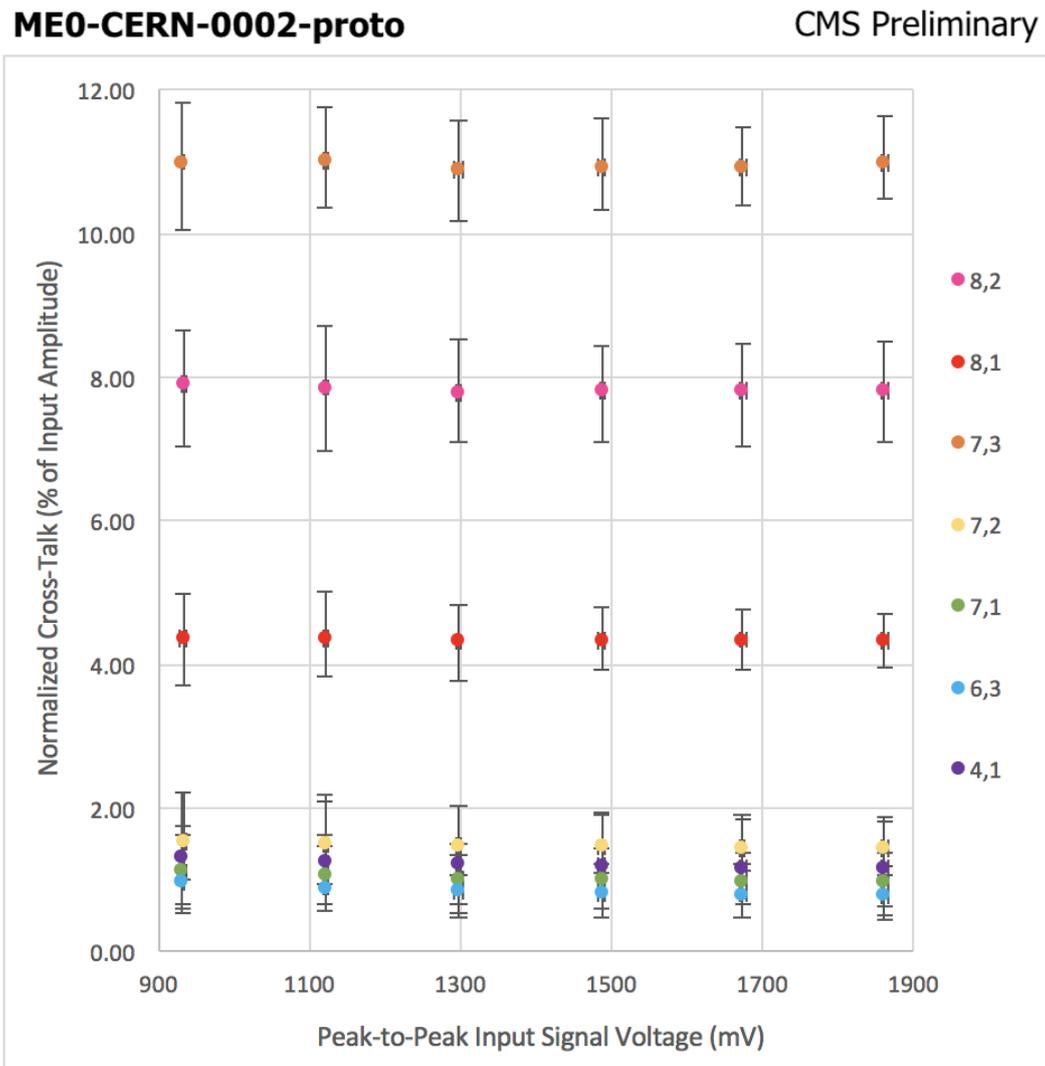


Figure 9.43 – Scatter plot of normalized cross-talk amplitudes in ME0-CERN-0002-proto for each readout sector vs. the incoming pulse amplitude

By normalizing the amplitude of the cross-talk pulse to the amplitude of the incoming square wave, we see that the amplitude of the cross-talk pulse is directly proportional to the amplitude of the incoming pulse. Thus, we can consider the average normalized value rather than having to consider individual values. By visualizing this averaged value over the physical area of the detector, as in Figures 9.44 and 9.45, where the $i\eta$ and $i\phi$ sectors correspond to those shown in Figure 9.38, a number of both expected and unexpected results emerge.

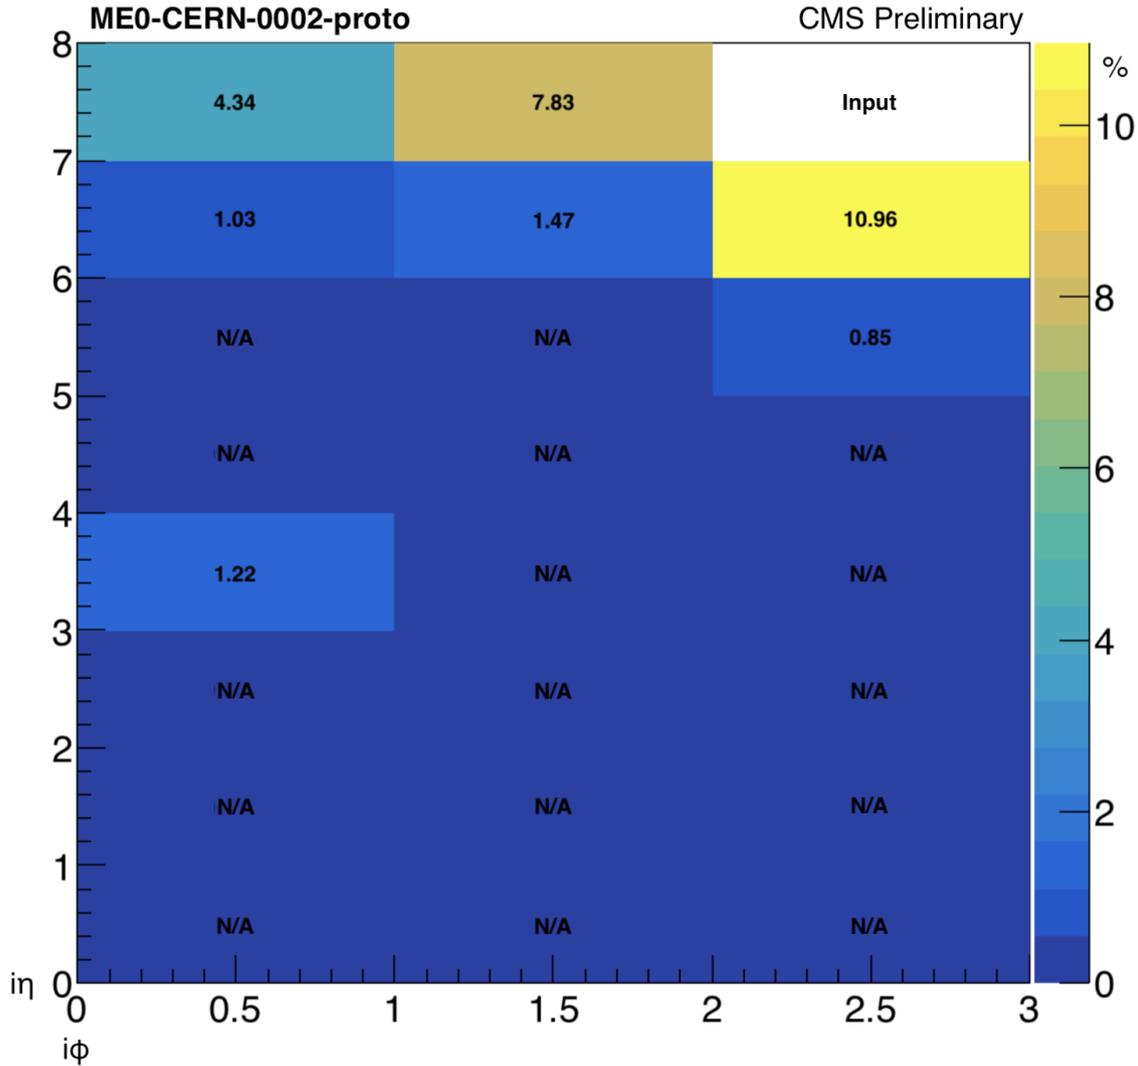


Figure 9.44 – Average normalized cross-talk amplitudes in ME0-CERN-0002-proto

As expected, the two readout sectors which are physically closest to the input sector 8,3 or within the same HV sector show the largest cross-talk pulse amplitudes, with sector 7,3 at 10.96%, sector 8,2 at 7.83%, and sector 8,1 at 4.34%. As sector 7,3's readout strips are all directly above sector 8,3's readout strips, rather than merely being adjacent to one side as in the case of sector 8,2, we see the highest level of cross-talk here. Beyond these sectors, the cross-talk decreases dramatically until it is indistinguishable from the background noise, labeled as "N/A" in Figure 9.44.

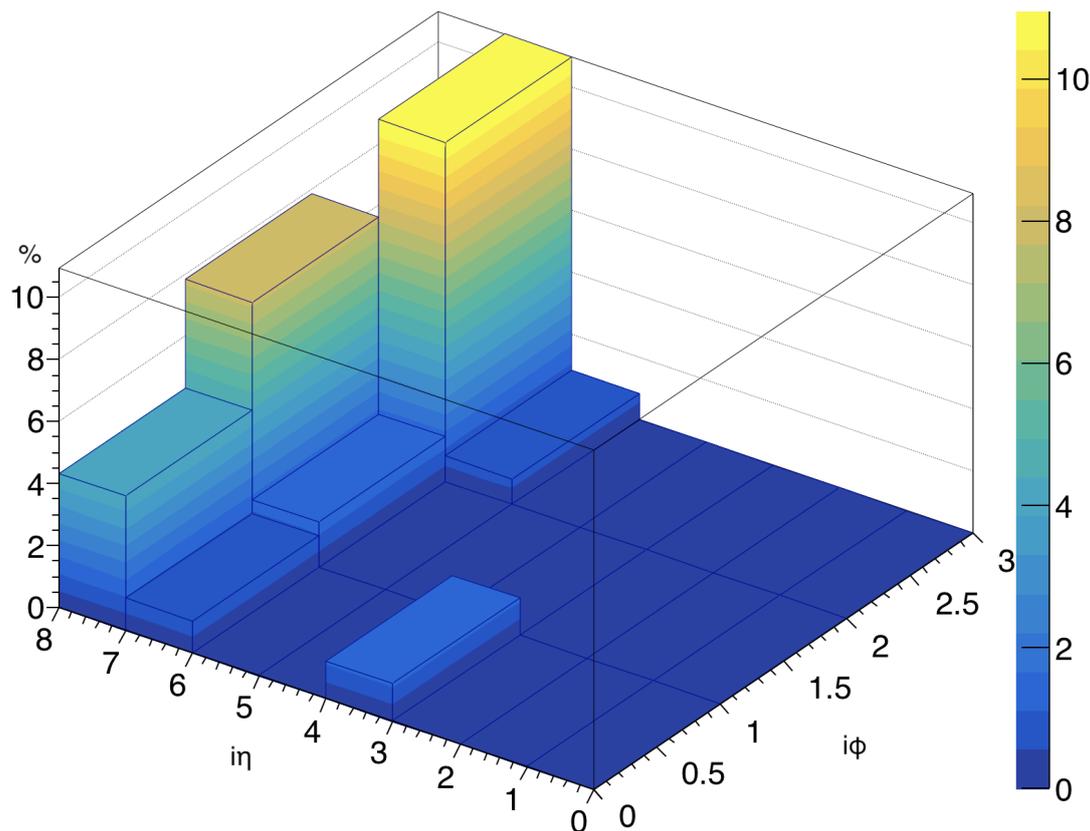


Figure 9.45 – Lego plot of average normalized cross-talk in ME0-CERN-0002-PROTO

However, in sector 4,1 we see a completely unexpected result. Despite being far from the pulse source, this sector shows abnormally high cross-talk amplitudes, outside of the expected pattern of behavior. Here, the Digilent Analog Discovery 2 oscilloscope became crucial in understanding this unexpected result, in that it provided a much clearer picture of the signal than the Tektronix TDS 2024C did. An example of a pulse as seen by the Analog Discovery 2 can be seen in Figure 9.46. When compared with Figure 9.42, we see that the so-called “cross-talk” pulse in this sector has a much more complex structure than in the other sectors. And, crucially, the part of the pulse which lines up with the rising and falling edges of the incoming square pulse has the opposite polarity than what is expected. This allows us to conclude that this is not a true cross-talk pulse as in the other sectors, but rather is the result of pulse pick-up as a result of a lack of proper grounding.

Pulse pick-up as a result of a lack of proper grounding was seen in the set-up at FIT and was eliminated via the use of a dedicated copper shielding plate, seen in Figure 9.47 on top of the ME0 detector. Each of the Panasonic-to-LEMO connectors was connected to this shielding plate via copper tape, and the set-up as a whole was connected to earth via the use of braided copper cables. Unfortunately, such a shielding plate is unique to FIT, having been designed for GE1/1 by the author of this dissertation in 2015, and was therefore not available for use at CERN.

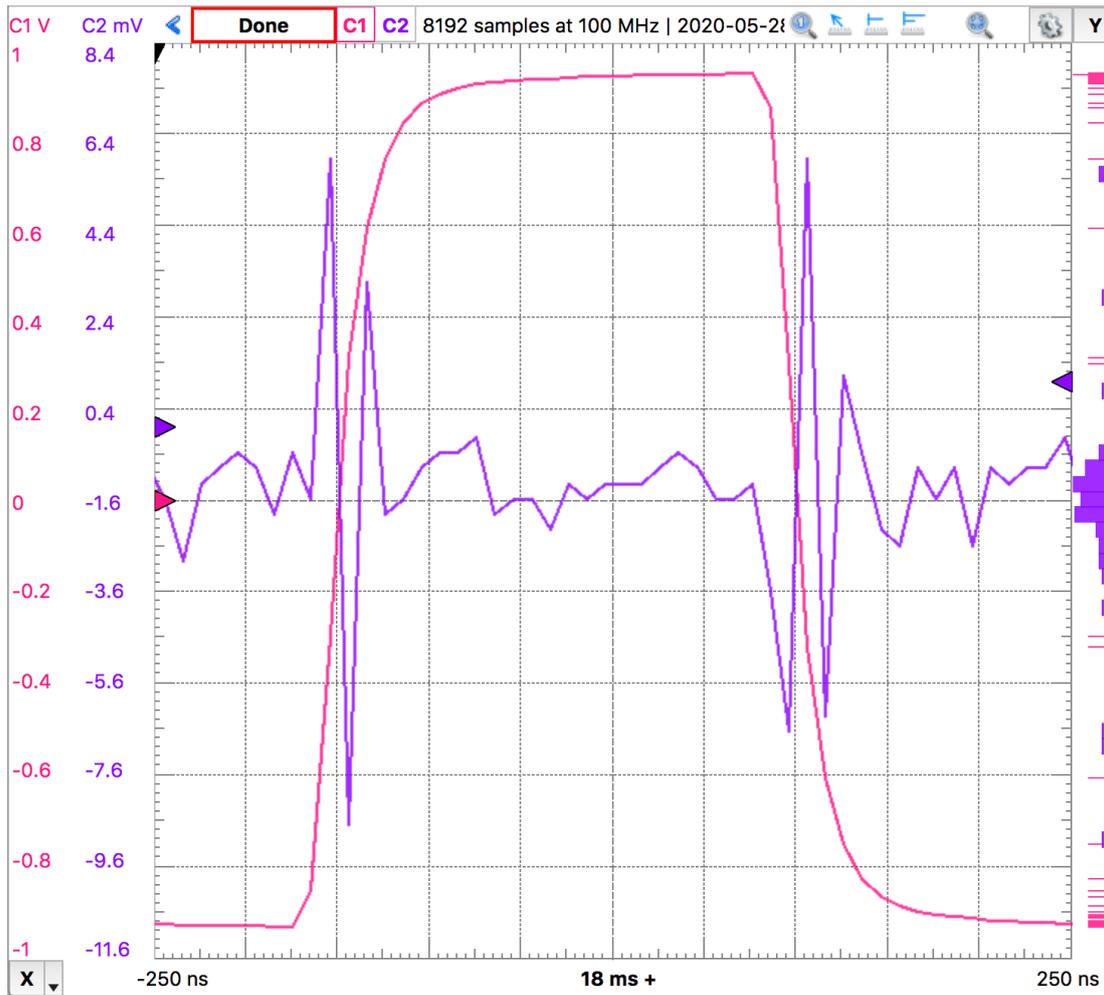


Figure 9.46 – Example Analog Discovery 2 readout of a cross-talk pulse resulting from a 1 V, 2 MHz square wave in sector 4,1

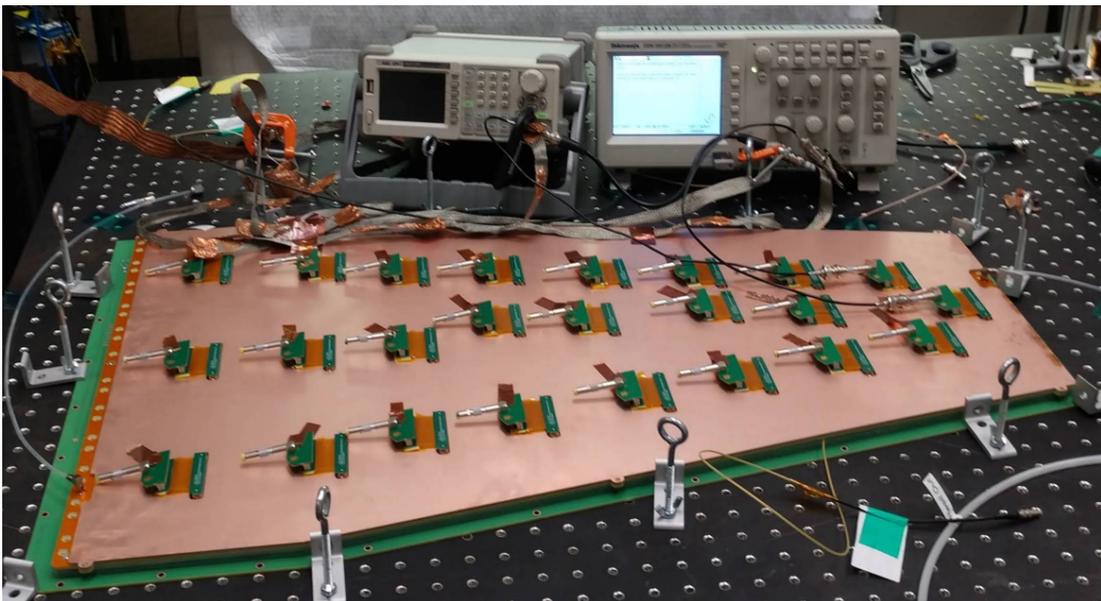


Figure 9.47 – Photo of the sister ME0 cross-talk set-up at FIT, illustrating the use of a dedicated copper shielding plate

The theory that this errant signal is the result of grounding-related pulse pick-up is further supported by the data that was first obtained using the initial setup seen in Figure 9.39, prior to the addition of the grounding measures seen in Figure 9.40. Although these results are not shown due to their obsolescence and lack of precision, it is worth noting that they showed cross-talk signal in two unexpected sectors - 5,1 and 3,1 at 6.16% and 9.06% respectively. These sectors, once properly grounded, showed no cross-talk at all, and the signal seen in sector 4,1 after grounding was of significantly lower amplitude at just 1.22%, supporting the idea of this being a grounding-related issue rather than true cross-talk.

Although further studies are planned for the set-up, including reducing the impedance $G3_{bottom}$ to ground, increasing the area of the HV segments of $G3_{bottom}$, and adding bypass capacitors onto $G3_{bottom}$, these studies have by necessity been put on hold for the time being. As a result of the strict restrictions put into place following the COVID-19 closure and staged reopening, only essential and time-sensitive work as declared on approved work packages may be performed. Any and all GE2/1 and ME0 work, as of June 2020, has not been declared essential, and although the above studies were able to be performed clandestinely, any further studies would require far more resources (such as use of the clean room) and therefore are not possible under the current scheme.

9.3.3 “Final” Configuration Choice

Despite the bureaucratic stoppage of work on GE2/1 and ME0 as a result of the COVID-19 restrictions, the timeline for production decisions to be made for these two systems remains nevertheless short. As a result, despite the lack of completeness of the studies discussed in this chapter, a decision had to be made in order to not further delay the production schedule of these systems. Although this is liable to change, the current “final” configuration choice as of the writing of this dissertation is described in this subsection.

One important consideration in deciding the final configuration is that of budgetary restrictions, such that any proposed solution should endeavor to be as cost-efficient as possible. In the case of the foils, one of the largest contributors to the total cost is the mask used to produce the foils. Therefore, an ideal solution would be one which would not require a separate mask to be produced for a portion of the foils needed for a given detector type. One such solution is depicted in Figure 9.48. Here, G1 and G2 would be double-segmented foils, with each HV partition separated by a 100 k Ω resistor. G3 would be made with the same mask, but the HV partitions of $G3_{bottom}$ would be connected via 0 Ω resistors in order to form the equivalent of a single solid plane. This plane would then be protected by a protection resistor of at least 100 k Ω . In this way, the detectors would benefit from the protection of double-segmented foils for G1 and G2, without running into the issue of cross-talk in G3, and without requiring two different foil masks be made for a single production.

However, this solution is not without its drawbacks. Moving back to the large induction capacitance on $G3_{bottom}$ of a single-segmented solution may reintroduce the problem of propagating discharges. As such, some tests are still needed in order to validate this as a workable solution, including a minimum of two weeks of discharge tests, followed by a minimum of two weeks of cross-talk studies, before a final decision could be made. Although the final results of these tests will be outside of the scope of this dissertation simply due to time constraints, Figure 9.49 presents a preliminary look at the discharge propagation probability of a detector which utilizes this solution.

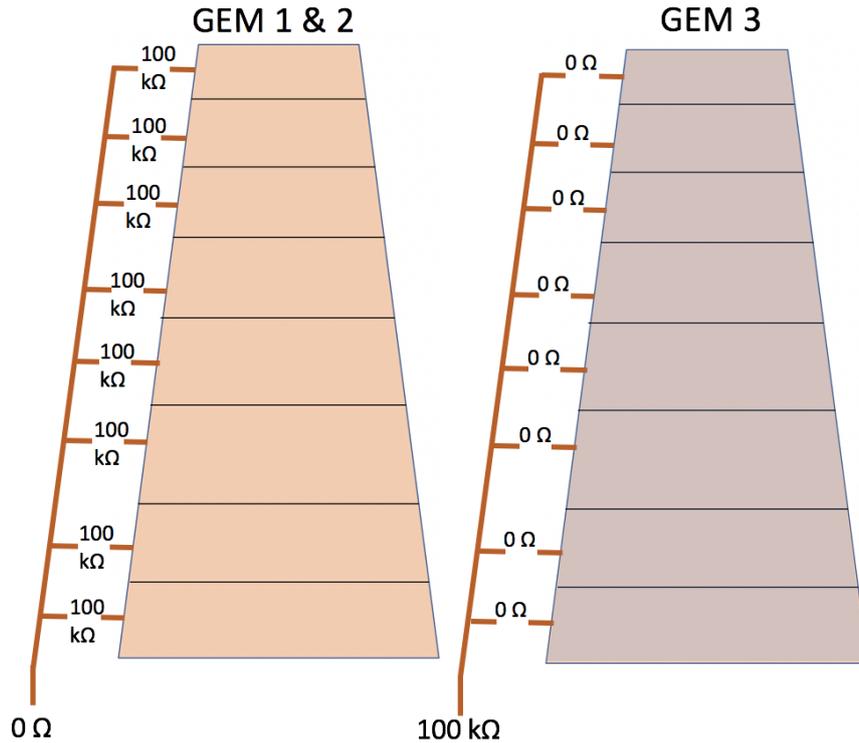


Figure 9.48 – Figure of the proposed solution for double-segmented foils in GE2/1 [117]

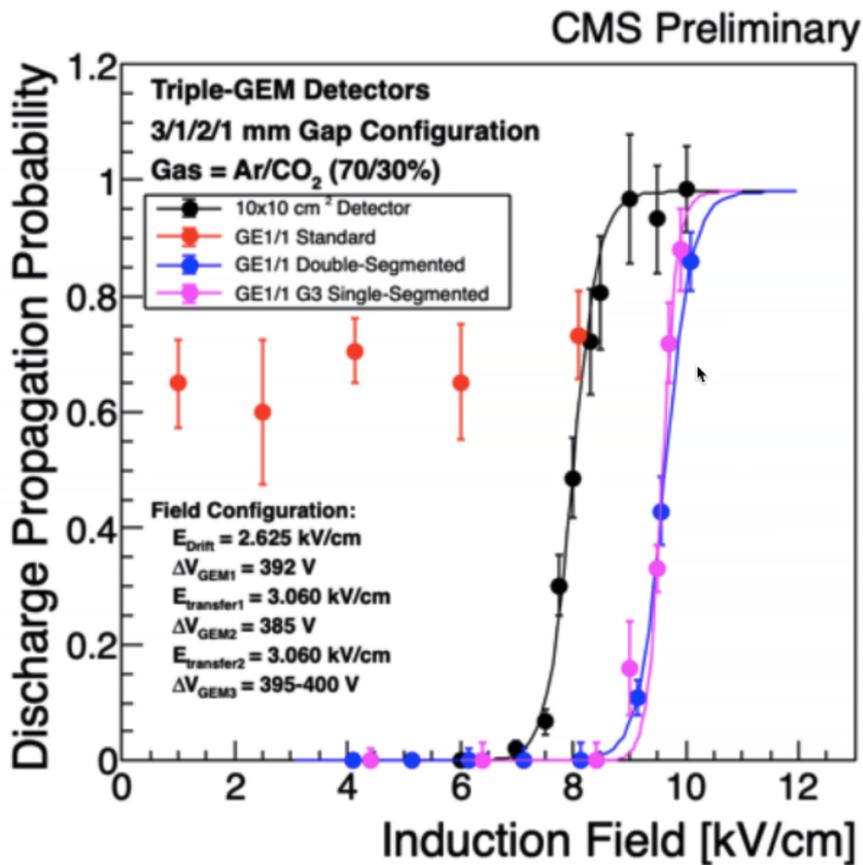


Figure 9.49 – Probability of a propagating discharge by induction field for four GEM detector types: a standard 10x10 cm², a single-segmented GE1/1, a double-segmented GE1/1, and a GE1/1 with double-segmented G1 and G2 but single-segmented G3

Here, we see very promising initial results, with this new configuration's discharge probability curve closely resembling that of the fully double-segmented chamber, albeit with a steeper slope at the point of sharply increased probability. Although highly preliminary, these results point towards this being a viable solution for robust and efficient GE2/1 and ME0 systems.

Chapter 10

Summary and Conclusions

10.1 The CMS GE1/1 System

The CMS GEM project represents a monumental achievement for the CMS detector, being the first entirely new subsystem, and entirely new detector technology, to be introduced into CMS since it was first built over a decade ago.

The detector technology in question, GEM detectors, are a class of micropattern gas detectors which operate on the principle of electron multiplication through cascading Townsend avalanches. Their primary feature is the presence of GEM foils, 50 μm -thick polyimide sheets that are coated with 5 μm of copper on each side and etched with tapered holes (diameter of 70 - 140 μm , pitch of 50 μm) using a technique of masked photolithography. These layers are stacked on top of a PCB readout and sealed within a gas-tight volume that is flooded with Ar:CO₂ 70:30. A high voltage is applied to the foils, which creates electric fields between the foils and through their holes. When a primary ionizing particle such as a muon passes through the active gas volume, the ionized electrons encounter these fields and multiply at each foil layer, for a total gain of approximately 10^4 at the readout board. GEM detectors are a robust and versatile technology, with a particle rate capability of up to $4.5 \frac{\text{kHz}}{\text{cm}^2}$, a per-chamber detection efficiency higher than 97% for detecting minimally-ionizing particles, an angular resolution of 300 μrad for $\Delta\Phi = \Phi_{\text{GE1/1}} - \Phi_{\text{ME1/1}}$, a timing resolution of 8 - 10 ns per detector, a gas gain uniformity of 30% across a single detector and between detectors, the ability to tolerate an integrated charge of $\sim 18 \frac{\text{mC}}{\text{cm}^2}$ without suffering damage or performance loss, and a discharge rate which does not impede performance or operations.

The GE1/1 system, the first of three CMS GEM systems to be installed in the CMS endcaps over the next several years, is composed of 144 meter-long triple-GEM detectors. These are coupled into two-detector superchambers, and arranged as 18 short ($1.61 \leq \eta \leq 2.18$) and 18 long ($1.55 \leq \eta \leq 2.18$) alternating superchambers per endcap, in order to accommodate the existing structure of CMS. In conjunction with the existing ME1/1 CSC chambers, the GE1/1 detectors effectively multiply a muon's path length within the first muon station by 2.4 - 3.5 compared to that of the ME1/1 chambers alone, and allow for the exploitation of the measurement of the bending angle between the two stations in an area where muons emerge at an angle of around 10° relative to the beam axis. This significantly improves the momentum resolution of the L1 standalone muon trigger and drastically reduces its disproportionately-large contribution to the overall L1 muon trigger rate. By allowing the L1 muon trigger thresholds to be maintained at low p_T values, the GE1/1 project will also significantly increase the efficiency of capturing interesting physics processes featuring soft leptons. Finally, when used for position matching with the new CMS track trigger, they will allow the standalone muon trigger to provide a high efficiency for displaced muons and other exotic particles in addition to serving as a backup for the combined muon trigger.

10.2 The Slice Test and Resulting Changes

Much of this dissertation took place in the context of the slice test, a demonstrator system of 10 GE1/1 prototype detectors which were installed in the CMS -1 endcap from 2017 through 2018. This demonstrator system would serve as a proof of concept for the full GE1/1 system, and also serve as a learning opportunity for how the GEM detectors behave under their full operational conditions with enough time remaining before LS2 to make any necessary changes to the final detectors as a result.

A task force called the sustained operations team was created to monitor the health and behavior of the slice test detectors through the use of time-series S-curve scans, which show the noise of each channel of a given VFAT. A channel which has been burnt or otherwise damaged such that the capacitance of the strip can no longer be seen manifests in the S-curve plots as having noise between 0.0414 and 0.109 fC.

In the context of this task force, the author performed in-depth studies of the behavior of the slice test GEM detectors within the CMS environment. This behavior included the electrode-by-electrode voltage and current stability of those detectors which were powered by a multi-channel power supply; the detection of discharges as they manifested in the voltage and current measurements of those detectors which were powered by a single-channel power supply, where those discharges were distorted and/or masked by the presence of an HV divider and the sampling speed of the electronics; and the correlation of the discharges with the environmental conditions within CMS at the time, including the magnetic field, beam luminosity, temperature, and pressure, as well as the operational conditions of the detectors, including their gain and voltage settings.

During the slice test, a significant percentage of the front-end readout channels were lost as a result of propagating electrical discharges within the detectors. As a result of the detailed studies performed by the author as listed above, it was determined that these discharges were caused in part by the sparking of contamination on the foils, which was present due to a lack of cleaning built into the quality control procedures, and in part due to incorrect operational values which, when combined with the presence of beam in the LHC, led to Townsend avalanches evolving into discharges within the detector.

In order to preserve the lifetime of the full GE1/1 system and prevent such rampant discharge damage from occurring in the future, several preventative steps were taken. First, new quality control procedures were implemented which clean the foils at multiple points during the chamber validation process. This includes foil cleaning and HV stability tests in QC4, and an intricate procedure for HV training of all foils, first in CO₂ and then in AR:CO₂ 70:30, both during QC8 and at P5. Second, all chambers are coupled into superchambers that share the same gain profiles, so no mix-up of operational parameters between detector layers can occur. And finally, a detail R/D campaign was launched to develop an external protection circuit for the VFAT3 hybrids to protect the VFATs from any discharges which may still occur despite the former two precautions.

Due to space and time constraints, only two protection circuits were considered: a set of resistors in series called HV3b_v3, and a set of diodes in parallel called HV3b_v4. Ultimately, the resistors were chosen, as the diodes failed to protect the front-end electronics as desired. The HV3b_v3 with 470 Ω of resistance was shown to be able to withstand 500 electrostatic discharges at 470 μ J per discharge without the VFAT sustaining any damage, and was shown to be durable against discharges of energies up to 1.5 mJ. This, however, comes at the cost of injecting a non-negligible amount of noise and cross-talk.

While the external protection circuit provides protection against propagating discharges, a secondary measure was put into place to prevent them from propagating in the first place. In cases of complex discharge propagation, much of the stored energy that fuels the re-ignition is stored in the HV filter. By modifying the HV filter with an additional 200 k Ω resistor between the GEM electrode and the default 10 k Ω resistor, we are able to sufficiently reduce the probability of discharge reignition without drastically reducing the rate capability of the detector. Between this measure and the HV3b_v3 protection circuit, the probability of damage to the VFAT3 was reduced from $92.8\% \pm 6.8\%$ to $2.86\% \pm 1.99\%$.

10.3 Future CMS GEM Systems: GE2/1 and ME0

Following the successful installation of GE1/1, two additional GEM systems will be added to the CMS endcaps in the coming years. The first, GE2/1, sits approximately two meters behind GE1/1 and interfaces with the ME2/1 CSC detectors, and has an acceptance range of $1.62 \leq \eta \leq 2.43$. It will consist of 72 triple-GEM detectors with an opening angle of 20° , arranged as 36 two-detector superchambers between the two endcaps. Each of these detectors are split into four individual modules which are bound together by an aluminum frame, each of which has its own optohybrid board. The VFAT3 chips will be installed on plug-in cards, whose larger area allow for larger, more robust external protection circuits to be utilized. Measurements of efficiency and protection capabilities for different protection circuits are currently ongoing.

The addition of GE2/1 increases the level of hits for a traversing muon from six to eight, a necessity for obtaining a robust track reconstruction at the L1 trigger. By increasing the local lever arm by more than a factor of two, and by having two independent directional measurements within the endcap, we will be able to make a standalone momentum measurement at the L1 trigger. In addition, by exploiting the measurement of the muon local bending angle at the GE1/1-ME1/1 and GE2/1-ME2/1 stations, we can greatly reduce both the fraction of soft muons with an overestimated momentum and the trigger rate for a given p_T threshold while still maintaining a high trigger efficiency.

The second future CMS GEM system, ME0, sits in front of GE1/1, and will act as the very forward muon tagger, with a geometric acceptance of $2.03 \leq \eta \leq 2.80$, higher than any current subdetector system in the muon endcaps. It will consist of 216 triple-GEM detectors with an opening angle of 20° , arranged as 36 six-detector stacks.

10.4 Double Segmentation of Future CMS GEM Detectors

One contributor to the complex propagation of discharges within large GEM detectors is the single segmentation of the GEM foils. In GE1/1 detectors, the top each foil is segmented into $< 100 \text{ cm}^2$ sectors which are separated by 10 M Ω resistors, while the bottom of each foil is left as one continuous conductor. This results in a high total capacitance on the bottom of the foil which leads to a higher probability for discharge propagation. Segmenting the bottom of the foils in much the same way as the top foils can reduce this probability significantly, at the cost of injecting an unwanted cross-talk signal into the system.

In CMS GEM detectors, a given signal has a large negative-polarity signal followed by a smaller, positive-polarity tail. The discrimination threshold is set such that it is triggered on the negative portion of the signal. However, in double-segmented detectors, the low capacitance of the 100 cm^2 partitions means that electrical potential fluctuations

are propagated along the entirety of the partition, and as a result, the adjacent readout sectors see a mirror image of this signal via capacitive coupling. If the original signal is large enough, as it is in the case of highly-ionizing particles, the tail of this mirrored signal can pass the discrimination threshold and trigger the front-end electronics as if it were real signal. This introduces dead time into the detectors by saturating the detectors' finite bandwidth with false signals, preventing the VFAT3 chips from detecting true signals, and ultimately reduces the efficiency of the detectors. At the worst estimate, this signal detection inefficiency results in a 3% total deadtime of the GE2/1 system.

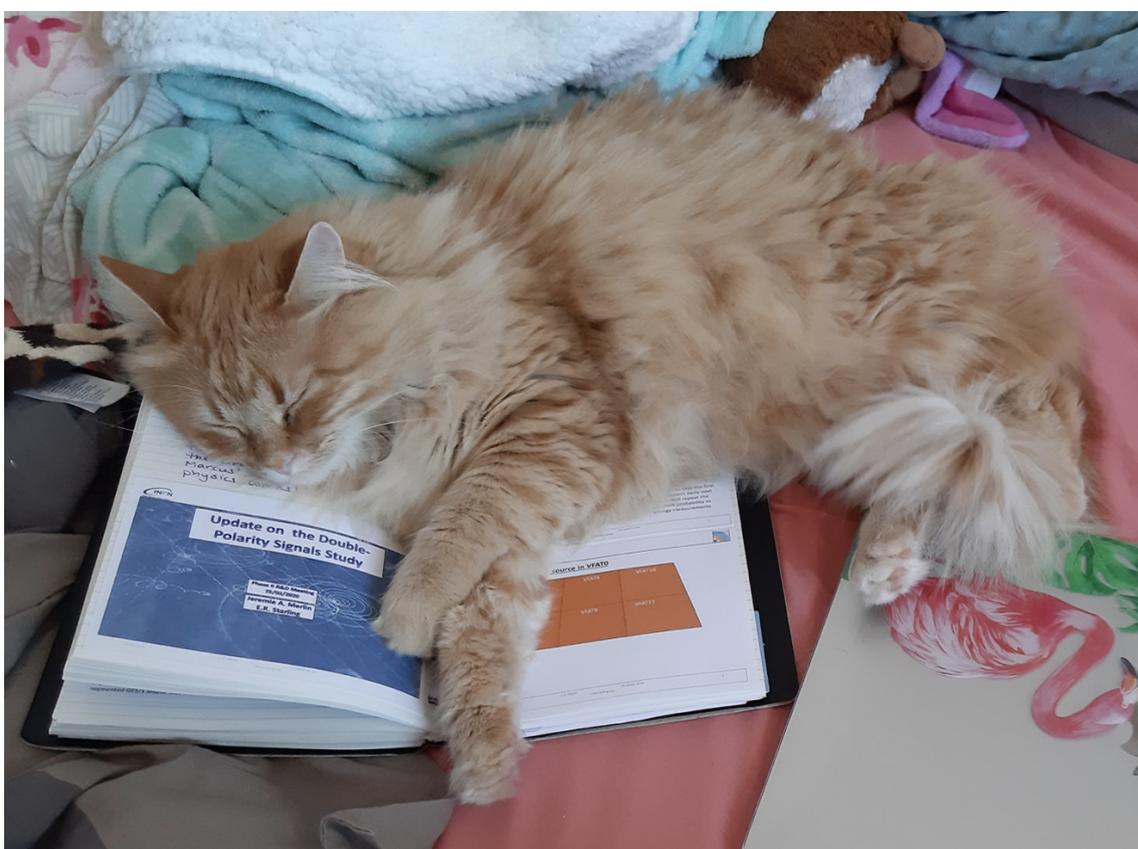
Studies were performed to investigate the distribution of the cross-talk signal within an ME0 detector, which found that the amplitude of the cross-talk signals are greatest in the η sectors directly above and below the input sector, as well as in the ϕ sectors along the same HV partition as the input sector. Additional studies were performed to investigate the effect the size of the partitions has on the cross-talk signal, which found that increasing the size of the partition decreased the presence of cross-talk, as expected.

Ultimately, the decision was made to implement double-segmentation on foils G1 and G2, while leaving G3 single-segmented. In order to minimize costs of production, the same photolithic mask would be applied to all foils, and then the segments on the bottom of the G3 foil would be connected with 0Ω resistors to imitate a single-segmented foil. Preliminary results show this to be a potentially viable solution, though further testing is still required.

10.5 Final Remarks

Ultimately the slice test proved to be an invaluable endeavor, as the knowledge of propagating discharges within the detectors prevented the catastrophic total destruction of the GE1/1 system's readout electronics. This knowledge kickstarted a campaign to improve the robustness against discharge damage not only for GE1/1, but for GE2/1 and ME0 as well. So far the preventative measures taken have been extremely promising, and further measures will continue to improve the operational lifespans of the future CMS GEM systems.

The results presented in this dissertation have far-reaching consequences which extend far beyond the CMS GEM systems, impacting not only the worldwide GEM community but the more expansive MPGD community as well. We have demonstrated for the first time that, with care to prevent discharge damage, large-area triple-GEM detectors can be successfully used in the harsh environment of an LHC experiment - a breakthrough which represents a great leap forward in the use of MPGDs in large-scale high energy physics experiments. Ultimately, we can look forward to a long and happy life for the CMS GEMs and a great, big, beautiful tomorrow for MPGDs in high energy physics.



The end...

Bibliography

- [1] B. Vormwald, “The CMS Phase-1 Pixel Detector: experience & lessons learned from 2 years of operation,” in *PIXEL 2018*, 2018, <https://tinyurl.com/ERSRef3J>.
- [2] L. Caminada, “The Phase 1 upgrade of the CMS pixel detector,” *Proceedings of Science*, vol. 254, no. 17, 2015, <https://pos.sissa.it/254/017/pdf>.
- [3] “Detector Overview,” 2011, <http://cms.web.cern.ch/news/detector-overview> (defunct as of 04 December 2020).
- [4] A. Branson, J. Shamdasani, and R. McClatchey, “A Description Driven Approach for Flexible Metadata Tracking,” in *7th ESA International Conference on Ensuring Long-Term Preservation and Adding Value to Scientific and Technical Data*, 2013, <https://arxiv.org/abs/1402.5953>.
- [5] naezmi, “Hadron calorimeter,” 2007, <https://tinyurl.com/ERSRefZ>.
- [6] M. Hoch, “CMS Solenoid Magnet during Long Shutdown 1 (LS1),” 2013, <https://cds.cern.ch/record/2017311>.
- [7] M. Bianco, “Upgrade of the CMS Muon Spectrometer in the forward region with the GEM technology,” in *International Conference on Technology and Instrumentation in Particle Physics*, 2017, <https://bit.ly/2TKe4QG>.
- [8] CMS Collaboration, “The Phase-2 Upgrade of the CMS Muon Detectors,” CERN, Geneva, Tech. Rep. CERN-LHCC-2017-012. CMS-TDR-016, Sep 2017, <https://cds.cern.ch/record/2283189>.
- [9] CIEMAT, “The CMS Drift Tube Chambers,” 2020, <http://cms.ciemat.es/drift-tubes1>.
- [10] G. M. I. Pugliese, “GHG Emission Reduction Strategy,” 2020, <https://tinyurl.com/ERSRef2Z>.
- [11] I. De Bruyn, “Muons Status Report,” in *CMS Remote Week Bangkok*, 2019, <https://tinyurl.com/ERSRef3A>.
- [12] R. Guida, B. Mandelli, and G. Rigoletti, “Performance studies of RPC detectors operated with new environmentally friendly gas mixtures in presence of LHC-like radiation background,” in *International Conference on High Energy Physics*, 2020, <https://tinyurl.com/ERSRef3B>.
- [13] Carbon Dioxide Information Analysis Center, Environmental Sciences Division, Oak Ridge National Laboratory, “CO2 emissions (metric tons per capita) - United States,” 2020, <https://tinyurl.com/ERSRef3D>.
- [14] R. R. Shinde and E. Yuvaraj, “Fabrication and Characterization of Glass Resistive Plate Chamber (RPC),” *Report for: Summer Research Fellowship Programme of India’s Science Academies*, <https://tinyurl.com/ERSRef3Q>.

- [15] S. Costantini et al., “Radiation background with the CMS RPCs at the LHC,” *Journal of Instrumentation*, vol. 10, 2015, <https://doi.org/10.1088/1748-0221/10/05/C05031>.
- [16] R. Breedon and D. Wenman, “Images of the CMS Muons - Cathode Strip Chambers (CSC),” 2008, <https://cds.cern.ch/record/1431505/>.
- [17] CMS Collaboration, “Performance of the CMS Cathode Strip Chambers with Cosmic Rays,” *Journal of Instrumentation*, vol. 5, 2010, <https://doi.org/10.1088/1748-0221/5/03/T03018>.
- [18] —, “The CMS trigger system,” *Journal of Instrumentation*, vol. 12, 2017, <http://doi.org/10.1088/1748-0221/12/01/P01020>.
- [19] K. Buńkowski, “RPC PAC muon trigger of the CMS detector,” in *Physics at Future Colliders*, 2006, <https://tinyurl.com/ERSRef3F>.
- [20] D. Ungaro, “The Link-Board Control in the RPC Trigger System for the CMS Experiment,” Ph.D. dissertation, University of Helsinki, 2004, <https://tinyurl.com/ERSRef3H>.
- [21] C. Fernández-Bedoya and G. Masetti, “Commissioning of the CMS DT electronics under magnetic field,” *Proceedings of the Topical Workshop on Electronics for Particle Physics*, pp. 81–85, 2007, <https://doi.org/10.5170/CERN-2009-006.81>.
- [22] J. Hauser, “The CMS Level 1 Muon Trigger,” 2007, <https://tinyurl.com/ERSRef3I>.
- [23] J. Erö, C. Deldicque, M. Galánthay, H. Bergauer, M. Jeitler, K. Kastner, B. Neuherz, I. Mikulec, M. Padrta, H. Rohringer, H. Sakulin, A. Taurok, C.-E. Wulz, A. Montanari, G. M. Dallavalle, L. Guiducci, G. Pellegrini, J. Fernández de Trocóniz, and I. Jiménez, “The CMS Drift Tube Trigger Track Finder,” *Journal of Instrumentation*, vol. 3, 2008, <https://doi.org/10.1088/1748-0221/3/08/P08006>.
- [24] G. Rakness, J. Hauser, and D. Wang, “Synchronization of the CMS Cathode Strip Chambers,” *Proceedings of the Topical Workshop on Electronics for Particle Physics*, pp. 80–84, 2007, <https://doi.org/10.5170/CERN-2007-007.80>.
- [25] “Gas Detectors Development Group,” 2019, <http://gdd.web.cern.ch/GDD/>.
- [26] J. Merlin, “GEM single-mask characterization and influence of the GEM foil orientation,” in *Micropattern Gas Detectors Conference 2017*, 2017, <https://tinyurl.com/ERSRefN>.
- [27] M. Villa, S. Duarte Pinto, M. Alfonsi, I. Brock, G. Croci, E. David, R. De Oliveira, L. Ropelewski, H. Taureg, and M. Van Stenis, “Progress on large area GEMs,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 628, pp. 182–186, 2011, <https://doi.org/10.1016/j.nima.2010.06.312>.
- [28] F. Fallavolita, “Triple-Gas Electron Multiplier technology for future upgrades of the CMS experiment: construction and certification of the CMS GE1/1 detectors and longevity studies,” Ph.D. dissertation, Università Degli Studi Di Pavia, 2019, <https://cds.cern.ch/record/2658126>.

- [29] H. Raether, *Electron Avalanches and Breakdown in Gases*. Washington, USA: Butterworths, 1964.
- [30] D. Pinci, “A triple-GEM detector for the muon system of the LHCb experiment,” Ph.D. dissertation, Università Degli Studi Di Cagliari, 2002, <https://tinyurl.com/ERSRefO>.
- [31] C. Richtera, A. Breskin, R. Chechik, D. Mormann, G. Gartya, and A. Sharma, “On the efficient electron transfer through GEM,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 478, pp. 538–558, 2002, [https://doi.org/10.1016/S0168-9002\(01\)00896-8](https://doi.org/10.1016/S0168-9002(01)00896-8).
- [32] F. Sauli, “The gas electron multiplier (GEM): Operating principles and applications,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 805, pp. 2–24, 2016, <https://doi.org/10.1016/j.nima.2015.07.060>.
- [33] S. Bachmann, A. Bressan, L. Ropelewski, F. Sauli, A. Sharma, and D. Mörmann, “Charge amplification and transfer processes in the gas electron multiplier,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 438, pp. 376–408, 1999, [https://doi.org/10.1016/S0168-9002\(99\)00820-7](https://doi.org/10.1016/S0168-9002(99)00820-7).
- [34] G. Bencivenni, G. Felici, F. Murtas, P. Valente, W. Bonivento, A. Cardini, A. Lai, D. Pinci, B. Saitta, and C. Bosio, “A triple GEM detector with pad readout for high rate charged particle triggering,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 488, pp. 493–502, 2002, [https://doi.org/10.1016/S0168-9002\(02\)00515-6](https://doi.org/10.1016/S0168-9002(02)00515-6).
- [35] G. Bencivenni, W. Bonivento, A. Cardini, C. Deplano, P. De Simone, G. Felici, D. Marras, F. Murtas, D. Pinci, M. Poli Lener, and D. Raspino, “A Fast Multi-GEM based Detector for High-Rate Charged-Particle Triggering,” *2001 IEEE Nuclear Science Symposium Conference Record*, vol. 1, pp. 462–466, 2002, <https://doi.org/10.1109/nssmic.2001.1008499>.
- [36] F. Sauli, *Principles of operation of multiwire proportional and drift chambers*. Geneva, Switzerland: CERN Yellow Report, 1977.
- [37] M. Abbas et al., “Performance of prototype GE1/1 chambers for the CMS muon spectrometer upgrade,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 972, 2020, <https://doi.org/10.1016/j.nima.2020.164104>.
- [38] J. Lemaire, “Exploring FPGA hardening solutions at the detector level for the future high luminosity phase of the CMS experiment at the LHC,” Master’s thesis, Université libre de Bruxelles.
- [39] D. Abbaneo et al., “Layout and Assembly Technique of the GEM Chambers for the Upgrade of the CMS First Muon Endcap Station,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 918, 2018, <https://doi.org/10.1016/j.nima.2018.11.061>.
- [40] G. De Lentdecker, “Lessons learned from GE1/1,” in *CMS Forward Muon Upgrade Workshop*, 2018, <https://tinyurl.com/ERSRef2K>.

- [41] M. Zubair Asyraf Bin Jalal, “Triple-Gas Electron Multiplier technology for the future upgrades of the CMS Muon Spectrometer: construction of the CMS GE1/1 SuperChambers.”
- [42] “Noise Investigation,” 2016, <https://tinyurl.com/ERSRef2O>.
- [43] P. Aspell, “VFAT: A 128 channel chip for charge sensitive readout of multi-channel silicon & gas particle detectors,” in *External KTT Network Meeting*, 2009, <https://tinyurl.com/ERSRef2S>.
- [44] —, “VFAT2 - Operating Manual,” 2011, <https://tinyurl.com/ERSRefA>.
- [45] P. Aspell, G. Anelli, P. Chalmet, J. Kaplon, K. Kloukinas, H. Mugnier, and W. Snoeys, “VFAT2 : A front-end system on chip providing fast trigger information, digitized data storage and formatting for the charge sensitive readout of multi-channel silicon and gas particle detectors,” *2008 IEEE Nuclear Science Symposium Conference Record*, 2008, <http://doi.org/10.1109/NSSMIC.2008.4774696>.
- [46] P. Aspell, “Preliminary VFAT3 User Manual,” 2018, <https://tinyurl.com/ERSRefM>.
- [47] P. Moreira, “GBTx Manual, V0.16,” 2018, <https://tinyurl.com/ERSRef3K>.
- [48] A. Gabrielli, K. Kloukinas, P. Moreira, A. Marchioro, S. Bonacini, and F. Sousa, “GBT-SCA: The Slow Control Adapter for the GBT System,” 2011, <https://tinyurl.com/ERSRef3L>.
- [49] A. Caratelli, K. Kloukinas, C. Paillard, S. Bonacini, A. Marchioro, R. De Oliveira, and P. Moreira, “GBT-SCA: The Slow Control Adapter for the GBT System,” 2015, <https://tinyurl.com/ERSRef3M>.
- [50] F. Vasey, L. Olantera, C. Soos, J. Troska, S. Kwan, A. Prosser, A. Xiang, J. Ye, T. Huffman, and T. Weidberg, “The Versatile Link Application Note,” 2017, <https://tinyurl.com/ERSRef3O>.
- [51] T. Lenzi, “Development of the DAQ System of Triple-GEM Detectors for the CMS Muon Spectrometer Upgrade at LHC,” Master’s thesis, Université libre de Bruxelles.
- [52] A. Colaleo, A. Safonov, A. Sharma, and M. Tytgat, “CMS Technical Design Report for the Muon Endcap GEM Upgrade,” Tech. Rep. CERN-LHCC-2015-012. CMS-TDR-013, Jun 2015, <https://cds.cern.ch/record/2021453>.
- [53] E. R. Starling, “Status of the Readout Electronics for the Triple-GEM Detectors of the CMS GE1/1 System and Performance of the Slice Test in the 2017-18 LHC Run,” *Proceedings of Science*, vol. 343, no. 132, 2019, <https://doi.org/10.22323/1.343.0132>.
- [54] CERN DC-DC Project, “FEASTMP_CLP,” 2014, rev. 1.0.
- [55] E. R. Starling, “Electronics System of the CMS GE1/1 Muon Upgrade and Performance of the Slice Test During the 2017-18 LHC Run,” in *Topical Workshop on Electronics for Particle Physics*, 2019, <https://tinyurl.com/ERSRef2Q>.
- [56] E. Hazen, A. Heister, C. Hill, J. Rohlf, S. Wu, and D. Zou, “The AMC13XG: a new generation clock/timing/DAQ module for CMS MicroTCA,” *Journal of Instrumentation*, vol. 8, 2013, <http://doi.org/10.1088/1748-0221/8/12/C12036>.

- [57] A. Svetek et al., “The Calorimeter Trigger Processor Card: The Next Generation of High Speed Algorithmic Data Processing at CMS,” in *Topical Workshop on Electronics for Particle Physics*, 2015, <https://tinyurl.com/ERSRef2F>.
- [58] A. Svetek, M. Blake, M. Cepeda Hermida, S. Dasu, L. Dodd, R. Fobes, B. Gomber, T. Gorski, Z. Guo, P. Klabbbers, A. Levine, I. Ojalvo, T. Ruggles, N. Smith, W. Smith, J. Tikalsky, M. Vicente, and N. Woods, “The Calorimeter Trigger Processor Card: the next generation of high speed algorithmic data processing at CMS,” *Journal of Instrumentation*, vol. 11, 2016, <http://doi.org/10.1088/1748-0221/11/02/C02011>.
- [59] E. Juska, “GEM backend electronics,” in *USCMS Phase-2 Forward Muon Upgrade Workshop*, 2016, <https://tinyurl.com/ERSRef2E>.
- [60] T. Lenzi, “Development of the DAQ System of Triple-GEM Detectors for the CMS Muon Spectrometer Upgrade at LHC,” Ph.D. dissertation, Université libre de Bruxelles, 2016, <https://tinyurl.com/ERSRef2D>.
- [61] —, “A micro-TCA based data acquisition system for the Triple-GEM detectors for the upgrade of the CMS forward muon spectrometer,” *Journal of Instrumentation*, vol. 12, 2016, <https://tinyurl.com/ERSRef2G>.
- [62] J. Mendez, “GBT-FPGA Tutorial,” in *Topical Workshop on Electronics for Particle Physics*, 2016, <https://tinyurl.com/ERSRef2H>.
- [63] E. R. Starling, “Electronics System of the CMS GE1/1 Muon Upgrade and Lessons Learned From the Slice Test During the 2017-2018 LHC Runs,” *Proceedings of Science*, vol. 370, no. 118, 2020, <https://doi.org/10.22323/1.370.0118>.
- [64] E. R. Starling and G. De Lentdecker, “Status of the Readout Electronics for the Triple-GEM Detectors of the CMS GE1/1 System and Performance of the Slice Test in the 2017-18 LHC Run,” in *Topical Workshop on Electronics for Particle Physics*, 2018, <https://tinyurl.com/ERSRef2R>.
- [65] G. Mocellin, “Trimming / Threshold Study,” in *GEM Operations and Online Meeting*, 2019, <https://tinyurl.com/ERSRef2U>.
- [66] R. King, “Real-Time Data Quality Monitoring System for the CMS Muon GEM Detector at the LHC,” Bachelor’s Thesis, Texas A&M University.
- [67] COMPASS Collaboration, “COMmon Muon Proton Apparatus for Structure and Spectroscopy,” 2020, <http://wwwcompass.cern.ch/>.
- [68] B. Ketzer, “Micropattern gaseous detectors in the COMPASS tracker,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 494, pp. 142–147, 2002, [https://doi.org/10.1016/S0168-9002\(02\)01457-2](https://doi.org/10.1016/S0168-9002(02)01457-2).
- [69] A. Colaleo, A. Safonov, A. Sharma, and M. Tytgat, “Construction of GEM Detectors for the COMPASS Experiment,” Tech. Rep., Dec 2011, <https://tinyurl.com/ERSRef3E>.
- [70] L. Ropelewski, “TOTEM GEM detectors TOTEM GEM detectors for tracking and triggering,” in *TOTEM T2 Engineering Design Review*, 2006, <https://tinyurl.com/ERSRef3F>.

- [71] C. Altunbasa, M. Capéans, K. Dehmelt, J. Ehlers, J. Friedrich, I. Konorov, A. Gandib, S. Kappler, B. Ketzer, R. De Oliveira, S. Paul, A. Placci, L. Ropelewski, F. Sauli, F. Simon, and M. Van Stenis, “Construction, test and commissioning of the triple-gem tracking detector for compass,” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 490, pp. 177–203, 2002, [https://doi.org/10.1016/S0168-9002\(02\)00910-5](https://doi.org/10.1016/S0168-9002(02)00910-5).
- [72] A. Colaleo, A. Safonov, A. Sharma, and M. Tytgat, “TOTEM Technical Design Report,” Tech. Rep. CERN-LHCC-2004-002. TOTEM-TDR-001, Jan 2004, <https://cds.cern.ch/record/704349>.
- [73] M. Alfonsi, G. Bencivenni, W. Bonivento, A. Cardini, P. De Simone, F. Murtas, D. Pinci, M. Poli Lener, D. Raspino, and B. Saitta, “The triple-GEM detector for the M1R1 muon station at LHCb,” *2005 IEEE Nuclear Science Symposium Conference Record*, <https://doi.org/10.1109/NSSMIC.2005.1596379>.
- [74] M. Poli Lener, “Triple-GEM detectors for the innermost region of the muon apparatus at the LHCb experiment,” Ph.D. dissertation, Università Degli Studi Di Roma, “Tor Vergata”, 2005, <https://tinyurl.com/ERSRef3G>.
- [75] D. Abbaneo et al., “Upgrade of the CMS muon system with triple-GEM detectors,” *Journal of Instrumentation*, vol. 9, 2014, <https://doi.org/10.1088/1748-0221/9/10/C10036>.
- [76] CERN Web Services, “Welcome to CMS RPC EndCap,” 2018, <https://tinyurl.com/ERSRef2A>.
- [77] M. Ressegotti, “Micropattern Gas Detectors for the CMS Experiment’s Muon System Upgrade: Performance Studies and Commissioning of the first GEM Detectors,” Ph.D. dissertation, Università Degli Studi Di Pavia, 2019, <https://tinyurl.com/ERSRef2B>.
- [78] M. Ressegotti and I. Vai, “CMSGEMDCS,” 2018, <https://tinyurl.com/ERSRef2C>.
- [79] “CMS DQM GUI,” <https://cmsweb.cern.ch/dqm/online>.
- [80] L. Tuura, G. Eulisse, and A. Meyer, “CMS data quality monitoring web service,” *Journal of Physics: Conference Series*, vol. 219, 2010, <https://doi.org/10.1088/1742-6596/219/7/072055>.
- [81] T. McCauley, “A browser-based event display for the CMS Experiment at the LHC using WebGL,” *Journal of Physics: Conference Series*, vol. 898, 2017, <https://doi.org/10.1088/1742-6596/898/7/072030>.
- [82] “GEMDPGPublic,” <https://tinyurl.com/ERSRef2Y>.
- [83] M. Abbas et al., “Performance of a Demonstrator Triple-GEM Detector in p p Collisions at the CMS Detector,” *Prepared for: Journal of Instrumentation*, currently in preparation for publication.
- [84] V. Peskov, M. Abbrescia, and P. Fonte, *Resistive Gaseous Detectors: Designs, Performance, and Perspectives*. Wiley, 2018.

- [85] J. A. Merlin, “Study of Discharges and Their Effects in GEM Detectors,” in *Micropattern Gas Detectors Conference 2019*, 2019.
- [86] V. Peskov, P. Fonte, M. Danielsson, C. Iacobaeus, J. Ostling, and M. Wallmark, “The Study and Optimization of New Micropattern Gaseous Detectors for High-Rate Applications,” *IEEE Transactions on Nuclear Science*, vol. 48, no. 4, pp. 1070–1074, 2001, <https://doi.org/10.1109/23.958725>.
- [87] F. Ivone, “Discharge mitigation strategies for GE11 detectors in the CMS Experiment,” Master’s thesis, Università Degli Studi Di Bari Aldo Moro.
- [88] V. Peskov, “Discharge Phenomena in Gaseous Detectors,” in *RD51 Collaboration Meeting and the "MPGD Stability" workshop*, 2018, <https://tinyurl.com/ERSRefG>.
- [89] Y. P. Raizer, *Gas Discharge Physics*, 1991.
- [90] J. Merlin, “Effects of Discharges on GEM Detectors - Preliminary Tests and Plans,” in *RD51 Collaboration Mini Week*, 2018, <https://tinyurl.com/ERSRefH>.
- [91] A. Utrobičić, “Usage of Single Hole THGEM Foil for Delayed Discharge Propagation Analysis,” in *RD51 Collaboration Meeting and the MPGD Stability Workshop*, 2018, <https://tinyurl.com/ERSRefE>.
- [92] A. Utrobičić, M. Kovacic, F. Erhardt, M. Jercic, N. Poljak, and M. Planinic, “Studies of the delayed discharge propagation in the Gas Electron Multiplier (GEM),” *Nuclear Instruments and Methods in Physics Research Section A*, vol. 940, pp. 262–273, 2019, <https://doi.org/10.1016/j.nima.2019.06.018>.
- [93] D. Fiorina, “GEM Spark Protection Studies,” in *24th GEM Workshop*, 2019, <https://tinyurl.com/ERSRefF>.
- [94] J. Merlin, “Study of the Discharge Propagation Effect in Triple-GEM Detectors,” in *Phase II R&D Meeting*, 2019.
- [95] CMS Collaboration, “Nonfactorization in Van der Meer scans in Run 2,” *CMS Performance Notes*, 2019, <https://tinyurl.com/ERSRef2V>.
- [96] M. Forge, “QC6 Test: Stability and I-V Characteristics of GEM foils,” in *GEM Student Day*, 2019, <https://tinyurl.com/ERSRef2X>.
- [97] B. Dorney, “HV Settings for GE1/1 Detectors,” 2017, <https://tinyurl.com/ERSRef2W>.
- [98] M. Tytgat, “Quality Control for the First Large Areas of Triple GEM Chambers for the CMS Endcaps,” in *Micropattern Gas Detectors*, 2015, <https://tinyurl.com/ERSRefD>.
- [99] F. Licciulli and G. De Robertis, “Protection Resistor Sustained Test,” 2018.
- [100] F. Licciulli, “VFAT3 Internal Protection,” 2018.
- [101] P. Aspell, “CMS GEMs Sustained Operation Studies, Front-End Hybrid R&D and Production,” 2018.

- [102] J. A. Merlin, “Discharge Protection, Cross-Talk, and Rate Capability,” 2020, muon Annual Review.
- [103] A. Mohammed and A. Ovcharova, “GE2/1 Readout Boards,” 2018, <https://tinyurl.com/ERSRefJ>.
- [104] M. Matveev, “Readout and Trigger Electronics for the Triple-GEM Detectors of the CMS GE2/1 System,” *Proceedings of Science*, vol. 370, no. 54, 2020, <https://doi.org/10.22323/1.370.0054>.
- [105] R. University, “GE2/1 Optohybrid Board Draft Specification 2.3,” 2020, <https://tinyurl.com/ERSRefR>.
- [106] Texas A&M University, “GE2/1 Backend System Interfaces,” 2019, <https://tinyurl.com/ERSRefT>.
- [107] E. Juska, “GE2/1 & ME0 backend electronics,” in *CMS GEM Workshop*, 2019, <https://tinyurl.com/ERSRefS>.
- [108] C. Calabria, “Measurement and simulation of the background in the CMS muon detectors,” in *Computing in High-Energy Physics*, 2018, <https://tinyurl.com/ERSRef3P>.
- [109] A. Peck and E. Juska, “ME0 Electronics Prototypes: Preliminary Testing Results,” 2019, <https://tinyurl.com/ERSRefU>.
- [110] J. Carlson, N. McColl, and A. Peck, “The ME0 ASIAGO Board (ME0-type Optohybrid),” 2020, <https://tinyurl.com/ERSRefV>.
- [111] “lpGBTv1 Manual,” 2011, <https://tinyurl.com/ERSRefX>.
- [112] F. Vasey, J. Troska, D. Ricci, S. Machado, A. Prosser, T. Huffman, T. Weidberg, J. Ye, and P. Gui, “The Versatile Link PLUS Project,” 2015, <https://tinyurl.com/ERSRef3N>.
- [113] F. Vasey, “Versatile Link PLUS: Status and Plans,” in *Sixth Common ATLAS CMS Electronics Workshop for LHC Upgrades (ACES)*, 2018, <https://tinyurl.com/ERSRefY>.
- [114] P. Moreira, “The lpGBT: a radiation tolerant ASIC for Data, Timing, Trigger and Control Applications in HL-LHC,” in *Topical Workshop on Electronics for Particle Physics*, 2019, <https://tinyurl.com/ERSRefW>.
- [115] F. Simone, “GEM sparks: status of validation mitigation measures (VFAT protection, foil segmentation),” in *Phase II Upgrade Group Review*, 2019.
- [116] J. Merlin, “Introduction to the Double-Polarity Signal Issue,” 2020, <https://tinyurl.com/ERSRefL>.
- [117] —, “Cross-talk Studies with VFAT3,” in *GEM Phase II Upgrade Workshop (GE2/1 and ME0)*, 2019.

-
- [118] P. Vichoudis, S. Baron, V. Bobillier, S. Haas, M. Hansen, M. Joos, and F. Vasey, “The Gigabit Link Interface Board (GLIB), a flexible system for the evaluation and use of GBTbased optical links,” *Journal of Instrumentation*, vol. 5, 2010, <https://doi.org/10.1088/1748-0221/5/11/C11007>.
- [119] J. Merlin and E. R. Starling, “Update on the Double-Polarity Signals Study,” 2020, <https://tinyurl.com/ERSRefK>.