The Faint Intergalactic-medium Redshifted Emission Balloon: FIREBall-2 Scientific Camera & Cooling System

Thesis by
Nicole Renate Lingner

In Partial Fulfillment of the Requirements for the
degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2017
Defended April 26, 2017

To Merle Monroe & Bertell

ACKNOWLEDGEMENTS

The work I am presenting here would not have been possible without the help and support of my husband, Bertell Nelson, and many friends and colleagues who have nurtured me throughout this process. “Once upon a time, and a very good time it was”, the talented Michael Schwartz – my first Astronomy professor – gently, yet forcefully made me apply for an internship at the Jet Propulsion Laboratory (JPL) where I met my mentor Amy Mainzer. Michael and his family have since become our dear friends and I will be forever grateful for him. Amy is a force to be reckoned with and I absolutely fell in love with her, her work and work ethic. She introduced me to the world of detectors and I worked with her and many of her friends and colleagues over the last ten plus years. During my undergraduate stint at UC Berkeley I was introduced to Stuart Bale, who was the first scientist to give me free range at his lab, where I learned about oscilloscopes and that nothing makes me happier than experimenting in a basement laboratory. Thank you guys for preparing me to be able to conquer what came next.

Upon arrival at Caltech I joined Chris Martin’s group and immediately started to work in his lab. I need to give a special shout out to Chris! Well, obviously, because he accepted me into this program and his research group, but mostly because he somehow saw that I would be able to pull all this together on my own — when I didn’t think that it was even remotely possible. Now I stand in front of you just ecstatic with my resume, which is beyond any expectation that I might have had when I first arrived at Caltech.

Patrick Morrissey, you took me to the cleanroom and showed me the way. Thanks to you, I will be forever vigilant and organized and try to stay grounded (wink) throughout my lab career.

Stephen Kaye was instrumental in helping me explore my lab capabilities and taught me how to fearlessly take matters, or electricity, into my own hands and experiment away in my little, lucky Lab 13.

Matt Matuszewski was kind enough to lend one or two of his ears to my never-ending science questions (really ANY questions: lab, machine room, food truck, or science, there wasn’t a topic for which he didn’t have an answer). He patiently let me help him built CWI and I learned a ton. You were my travel companion on my first trip to Hawaii and (I think) we had a blast. But seriously, you carried me from

day one, when you took me around campus and introduced me to everyone in our department and beyond. You set me up with study groups, technical support and explained machine room etiquette. Thanks for being you.

Erika Hamden. Oh Erika, I could have never finished this thesis or PhD without you, my sweet and energetic friend. You were always available, helpful and encouraging. I hope to have you in my life forever, or should I say our lives, since you were there throughout my pregnancy and one of the first humans to come visit my daughter Merle. So. Many. Thanks.

Cherie Capri has made my life as an almost full-fledged scientist that much better. Amongst the chaos of traveling and ordering and organizing you kept it all in check and I was able to spend my time and energy building an instrument. You were the other human to hold my hand throughout the pregnancy and thesis writing and greeting Merle, who now calls you grandma.

Dr. John, nobody was as kind, calm and outrageously knowledgeable as you. You explained everything I ever asked about lab equipment and procedures and you did it with wit and charm. You were so wonderfully quirky and said whatever came to your mind, whenever it came to your mind, which made group meetings a lot more fun. I am not alone in missing you dearly and am so deeply saddened that you left this world last December.

Minerva Calderon. I met you on my first day and, because you were such a gentle, happy, and lively person to be around, I wanted to move into your computer dungeon and never leave. We spent a lot of quality time together and you helped me professionally and personally. I even got to tutor your husband in German, which was a lot of fun. I don't have words to describe how sad I am that ALS took you from us.

Gwen Rudie, I always misspelled your name and you still explained astrophysics to me. I was super happy to find out that your science interest and mine aligned, except you were a few years ahead and I could ask and verify a bunch of specific stuff with you. I appreciate your caring honesty.

Jill Burnham, you held my hand through my first attempts at designing an electrical board and you must have laughed so hard at my initial questions. It was so good to have you in my corner, because we all know I needed the help. I had you verify every trace and you agreed to do it too.

This of course brings me to one of my invaluable EMCCD resources, Olivier Daigle. Not only did you come up with an amazing machine that can coax premium signal out of this new breed of sensors (CCCP, not the Soviet Union Wikipedia entry), you are such a nice person to talk to. I am glad that I had so many questions and the opportunity to talk to you seemingly every day there for a while. Thank you for getting back to me, always.

Dean Johnson, I had the absolute pleasure of visiting you on many occasions at your lab and pointing at things to ask countless questions. Thanks to you I have fallen in love with cryo-engineering and I hope to work with you again.

Richard Mewaldt you were so friendly and helpful and offered your infinite knowledge of upper atmosphere particles and everything else concerning high energy and particles. Investigating Cherenkov radiation was actually super fun with you by my side.

Heather Duckworth, my friend and lab confidant. You were integral in getting the scripts for our beloved Nüvü controller together and I loved hanging out with you in the lab and, of course, in snowy Montreal.

An then there are several physicists, engineers, and shops, who have helped me understand and build all the lab and flight hardware, which is now happily integrated into the flight payload. Thank you Bruno Milliard (LAM), David Schiminovich (Columbia, NY), Marc Jaquet (LAM), Marty Crabill (Caltech), Gillian Kyne (Caltech), Michele Yeater (Caltech), Donna Driscoll (Caltech), Hector Rodriguez (Caltech), Roger Smith (Caltech), Accurate Engineering, The Physics Shop (Caltech), Mark Gonzales (Caltech), Jimmy Wade (Sunpower), Herman Cease (Fermilab), Shouleh Nikzad (JPL), Todd Jones (JPL), and Herr Seibert (Carnegie).

And a special thanks goes to the Palomar crew, who always made me feel right at home. I could not wait to get back to the 200” every few months to help with CWI. Thank you Steven Kunsman, Gregord Van Idsinger, Mike Doyle, Jean Mueller, Kajsa Peffer, Dipali and everyone else on that great big hill.

ABSTRACT

The FIREBall-2 instrument provides a means of investigating both the processes responsible for building cosmic structure and the galaxies that trace it. By studying the inflow and outflow of gaseous hydrogen, we can better explain galaxy formation and evolution. Each individual galaxy lies within a dense gaseous region called the circumgalactic medium (CGM), itself surrounded by the diffuse intergalactic medium (IGM). Thin gaseous filaments connect galaxies to each other in space and time and they can be found at nodes of these filaments at the center of dark matter halos.

This large-scale structure, called the cosmic web, has been modeled and corroborated with absorption line studies. To understand the structure of the galaxy environment, we need to map the extremely faint emission from the CGM and these filaments.

FIREBall-2 is a 1-meter-class, balloon-borne, UV telescope that will be launched from Texas in the fall of 2017. The instrument includes a vacuum tank, which houses an all-reflective, wide-field multislit spectrograph, guide camera, and UV optimized EMCCD sensor. The scientific camera sensor has to be operated near -110°C and is cooled with a mechanical cryocooler, which also powers a charcoal getter to maintain high vacuum during flight.

The scientific camera includes a Printed Circuit Board (PCB), which I use as a rigid harness that holds the sensor in place. The sensor (CCD201-20 from e2v) is a delta-doped, electron-multiplying charge-coupled device (EMCCD) that has been modified to be used as a 1K x 2K (1024 x 2048 pixels) sensor. The EMCCD has been developed at the Jet Propulsion Laboratory's Micro Devices Laboratory. I operate this next-generation UV detector (NEXUS) at 10 MHz with read-out electronics from Nüvü Camēras. Nüvü's controller for counting photons (CCCP) was chosen to achieve extremely low detector noise by reducing clock induced charge (CIC) and it is attached to the PCB with a SAMTEC EQCD high-speed coaxial cable. To reduce dark current, the other relevant detector noise source, the EMCCD is supported by a gold-plated copper clamp and cooled with a CryoTel MT mechanical cooler (Stirling engine) from Sunpower. The CCD operating temperature is -110°C, with a required heat lift of about 7 Watts at a 30°C reject temperature. Two 10-Watt Omega heaters are used to regulate the temperature. The cryocooler will be operated at a constant input cooling power of about 70 Watts, which corresponds to 110-Watts of battery

power. The cooler is mounted to the side of the vacuum tank and I have incorporated the Sunpower active vibration cancellation system (AVC) to reduce the vibrational noise. A flexible copper ribbon conducts heat from the bottom of the CCD via the solid copper clamp. The detector is operated in a vacuum of about 10^{-6} Torr, initially started by a turbo pump, with the pressure then lowered and maintained by activated charcoal adsorption during flight. The flight hardware has been integrated into the flight vacuum tank and is currently undergoing testing.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

Chapter 1

INTRODUCTION

1.1 Galaxies and the Cosmic Web

Introduction

Observations at different wavelengths have shown that galaxies are not uniformly distributed throughout space, but instead reside in dense regions of large-scale structure. The origin of these structures also yields large filaments of gas which inflow back to the galaxies themselves at the center of dark matter halos. Models predict that the dark matter seeded by primordial quantum fluctuations formed the architecture of the universe, creating a cosmic web of sheets and filaments of dark and normal (baryonic) matter. If we can understand how galaxies evolve in the cosmic web environment, we can better understand our cosmic environment and the universe at large. New instrumentation is poised to answer questions about our cosmic origins, and we learn more with every new observational tool that comes online. Guided by previous measurements and simulations, we seek to develop a novel instrument that will investigate galaxy evolution, reaching beyond current capabilities. By expanding our understanding of the cosmic web, we will come to know better the universe in which we live.

Large Scale Dark-Matter Simulations

Dark matter simulations allow us to model and understand the clustering of galaxies and their formation and evolution within the parent dark matter halo. It is encouraging that we have observed galaxy groups confirming the structures that these simulations predict on large scales. To understand the mechanisms that drive galaxy formation and evolution, however, we need to look in more detail at galaxies and their gaseous environments. That is, we need to get a better picture of the diffuse intergalactic medium (IGM), which spans the vast majority of space, and the circumgalactic medium (CGM) that occupies the dark halos at the interface of galaxies and IGM. Observing the gas that surrounds galaxies essentially amounts to observations of hydrogen, which makes up three quarters of the baryonic mass in the universe and is the main component of the IGM and CGM .

If we want to understand galaxy formation and evolution, we must trace the inflow and outflow of hydrogen. It is also important to classify these hydrogen clouds to

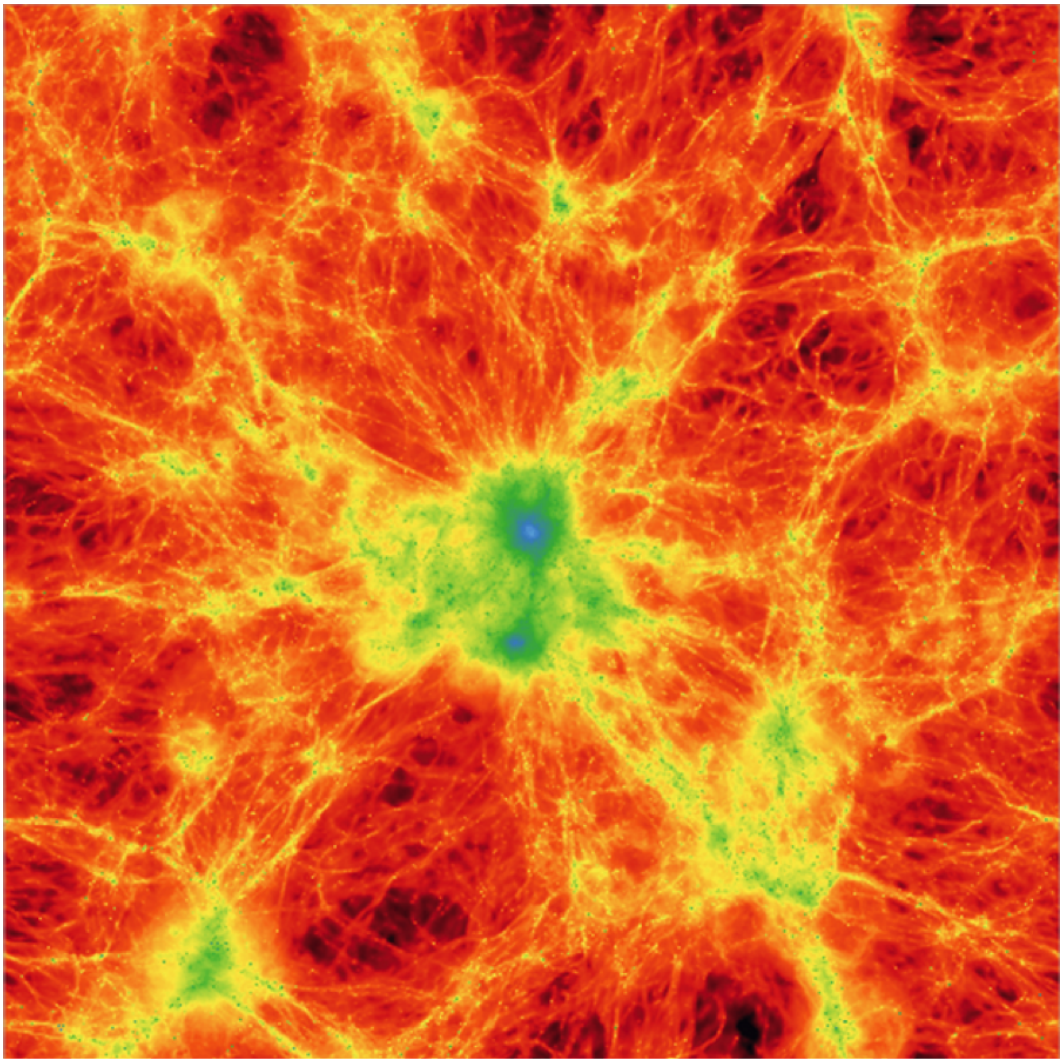

Figure 1.1: Simulation of the gas density field around a central elliptical galaxy. Credit: Illustris Project

distinguish between different regions and ultimately describe the structure of the galaxy halo environment: is it composed of individual discrete clouds, a uniform fog, or distinct filaments as in the simulation above? Gas in this context is best described by measuring column density, because it has a direct physical meaning: the number of ions per cm^2 (Mihos, 2005). For cosmic web observations, we focus in particular on neutral hydrogen column density. Because current state-of-the-art techniques rely on absorption line studies (Section 1.2), regions with high column densities (significant density of H_I) are referred to as absorbers—i.e., a "strong absorber" refers to a blob of high column density.

1.2 Detecting Absorption vs Emission

Absorption Line Studies of the CGM and IGM

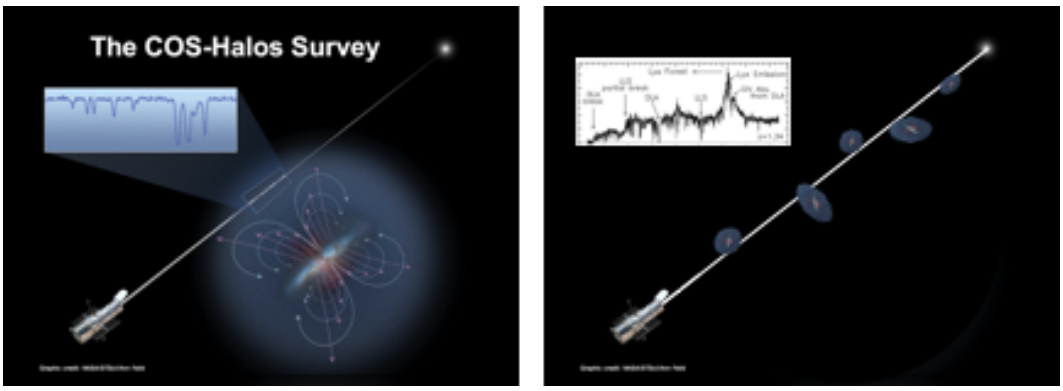

Figure 1.2: Left: This artist's conception depicts an absorption line study: The spectrum of a bright quasar in the background is captured by the spectrograph, and as the light from the quasar passes through an absorber surrounding an intervening galaxy, absorption lines are imprinted on the quasar emission spectrum. Right: This cartoon demonstrates how the light from this quasar, along the line of sight, could intersect multiple absorbers surrounding galaxies at different distances. Because these target galaxies reside at different redshifts, the Lyman-α line is shifted, leaving the Lyman-α forest spectral signature (Graphic credit: NASA/STScI/Ann Field, inset from Charlton and Churchill 2000).

Most of what we know about these absorbers comes from absorption line studies, which is an indirect method of looking at the gas surrounding galaxies (see Gwen Charlena Rudie (2013)). This procedure requires a bright background source, like a quasar, and an intervening foreground absorber (in many cases there are multiple absorbers). We can then identify a given absorber by looking at the absorption features in the background quasar spectrum. The neutral hydrogen Lyman-α line (at 1215.67 Å) is the most abundant transition and an indicator of the presence of hydrogen in the absorber. The absorption system that we find in a quasar emission profile looks like it has been spread out into a "forest" of lines. This occurs because each line is redshifted by a different amount corresponding to the absorber's distance to us, creating the "Lyman-α forest." Depending on the composition of the absorber, these absorption lines will have additional spectral signatures that tell us more about the metallicity, temperature, and density of the medium (Rauch and Murdin, 2006).

Until very recently, investigating the gas within a galaxy's CGM has primarily been approached by looking at these absorption features in the light from bright background quasi-stellar objects (QSOs).

One shortcoming of this kind of analysis is its failure to capture the three-dimensional distribution of the gas. We do not have a complete picture of what is going on around a given galaxy and, just as important to understanding galaxy evolution, we have a limited idea of how the gas is distributed along the line of sight. While these surveys are carefully prepared, there is generally only one quasar behind a given galaxy, restricting our view to one tiny region somewhere around this galaxy. The same quasar intersects the environment of galaxies at other redshifts, but altogether this incomplete data cannot reveal how those galaxies might connect.

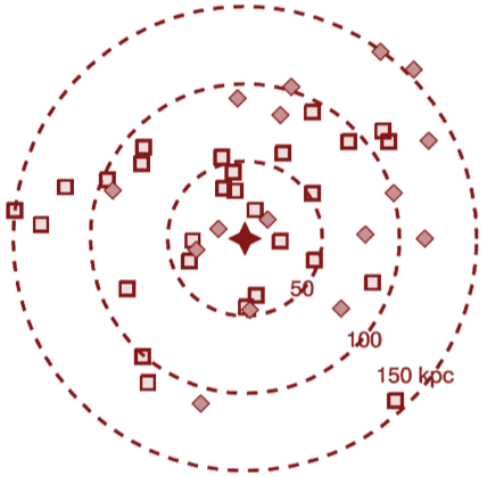

Figure 1.3: Distribution of QSOs on the sky with respect to the target galaxies. This is a Stacked Image and each galaxy has been shifted to the center and the radial coordinate is physical kpc at a given galaxy redshift. (J. Tumlinson et al., 2013)

Figure 1.3 illustrates the statistical information that can be obtained by probing the environment of one specific galaxy in only one small area. By combining sparse data from multiple galaxies, statistical statements can be made about galaxy types in general, without describing the environment of any particular galaxy. Even so, we have learned a great deal about the IGM and CGM from these absorption line maps and they have contributed a great deal to our understanding of the cosmic web. Gwen Rudie and Chuck Steidel at Caltech completed a large study of the CGM and IGM around galaxies with the Keck Baryonic Structure Survey (KBSS) in 2012 (Gwen C Rudie et al., 2012). They looked at the neutral hydrogen gas around galaxies and then combined their findings to reach a statistical understanding of the environment around galaxies at high redshift. This investigation has helped shape how we see galaxies and the medium in which they live.

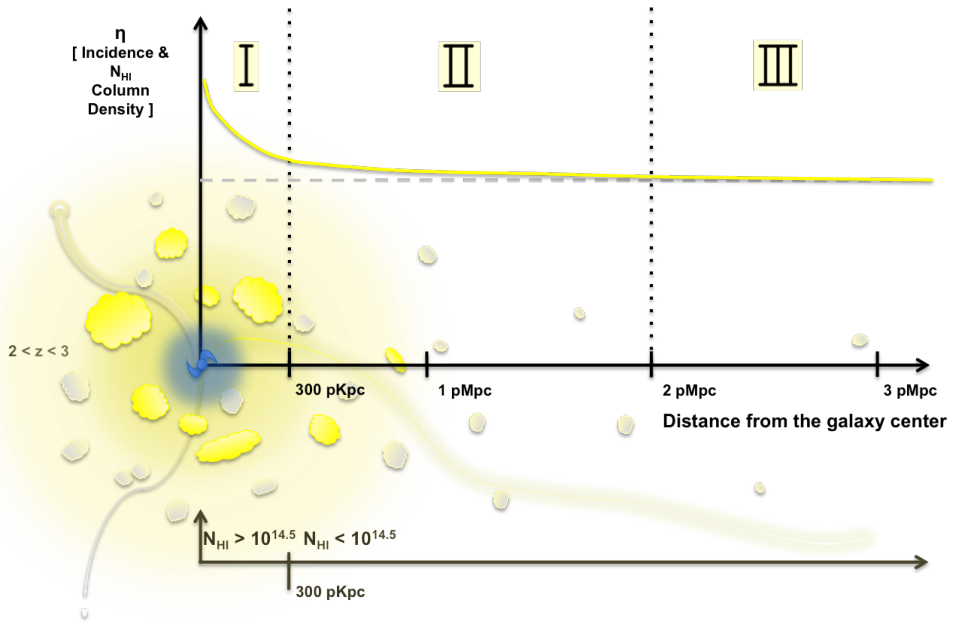

Figure 1.4: Visualization of our current understanding of the cosmic web and galaxy halos. The clouds represent puffy overdensities. There are three distinct zones in the gaseous envelopes around galaxies.

Figure 1.4 is a summary of what we know so far about the medium that surrounds galaxies and how it extends to other galaxies. A schematic model, this cartoon provides a way to visualize specific questions about the cosmic web. As described thoroughly in Gwen C Rudie et al., 2012, there seems to be a standard structure to most galaxy halos. By combining galaxy environment information from almost 900 galaxies, they found three distinct zones in the gaseous envelopes around galaxies. The density of neutral hydrogen gas falls off at about 300pkpc, and we can refer to this first zone as the circumgalactic medium. The second zone sits between 300 pkpc and 2 pMpc. Neutral hydrogen with high column density is not as abundant as in the first zone; it plateaus to a roughly constant value, but is still higher than in the third zone. There is no sharp boundary between the second and third zones, and the distribution of absorbers between 2 pMpc and 3 pMpc is identical to that found anywhere in the IGM, even outside galaxy clusters.

These same three zones can be identified in the velocity distribution of absorbers along the line of sight of a QSO observation. In the cosmic web context, it is probably easiest to think of galaxies not as closed systems, but as interconnected, living and breathing objects. While these studies have provided a valuable general picture of

a galaxy halo, absorption line observations of the cosmic web cannot distinguish among rival models of CGM structure: a uniform fog, a cloud containing clumps, or filamentary structure. This data does not reveal much about the actual distribution of material surrounding any one galaxy, crucial information needed to answer specific questions about cosmic structure:

- Is all of space filled with diffuse gas, with some regions denser than others, or is this medium spherical and more defined?
- How is the gas around this galaxy connected to the surrounding gas both in time and space?
- Is matter blown out of the galaxy by supernovae explosions, active galactic nuclei, and star formation?
- If there are filaments, are they scattered and lost from such a blast? Do they reconnect over time, or form new paths?
- Do galactic winds inject energy, metals, and momentum into the CGM, which could suppress star formation?
- Does this material eventually fall back in and help star formation start up again?

These fundamental questions can only be answered by directly observing and mapping the emission from the target galaxy and its CGM out to the IGM. The wealth of knowledge provided by absorption line studies does not explain any of this yet, but it is the first step in verifying or adjusting existing models of the cosmic web. These models are essential to understand what physical properties we can, or should, look for with our instrument.

Knowledge from Absorption Line Studies & Models

We are working under the assumption that galaxies have a two-mode gas accretion process—cold mode accretion and hot mode accretion—and we seek to determine whether these modes actually exist. Cold flow accretion refers to gas that enters the galaxy along the filamentary structure of the cosmic web, without going through a virial shock (Kleiner et al., 2016). Hot mode accretion, on the other hand, refers to shock-heated gas that falls in gravitationally.

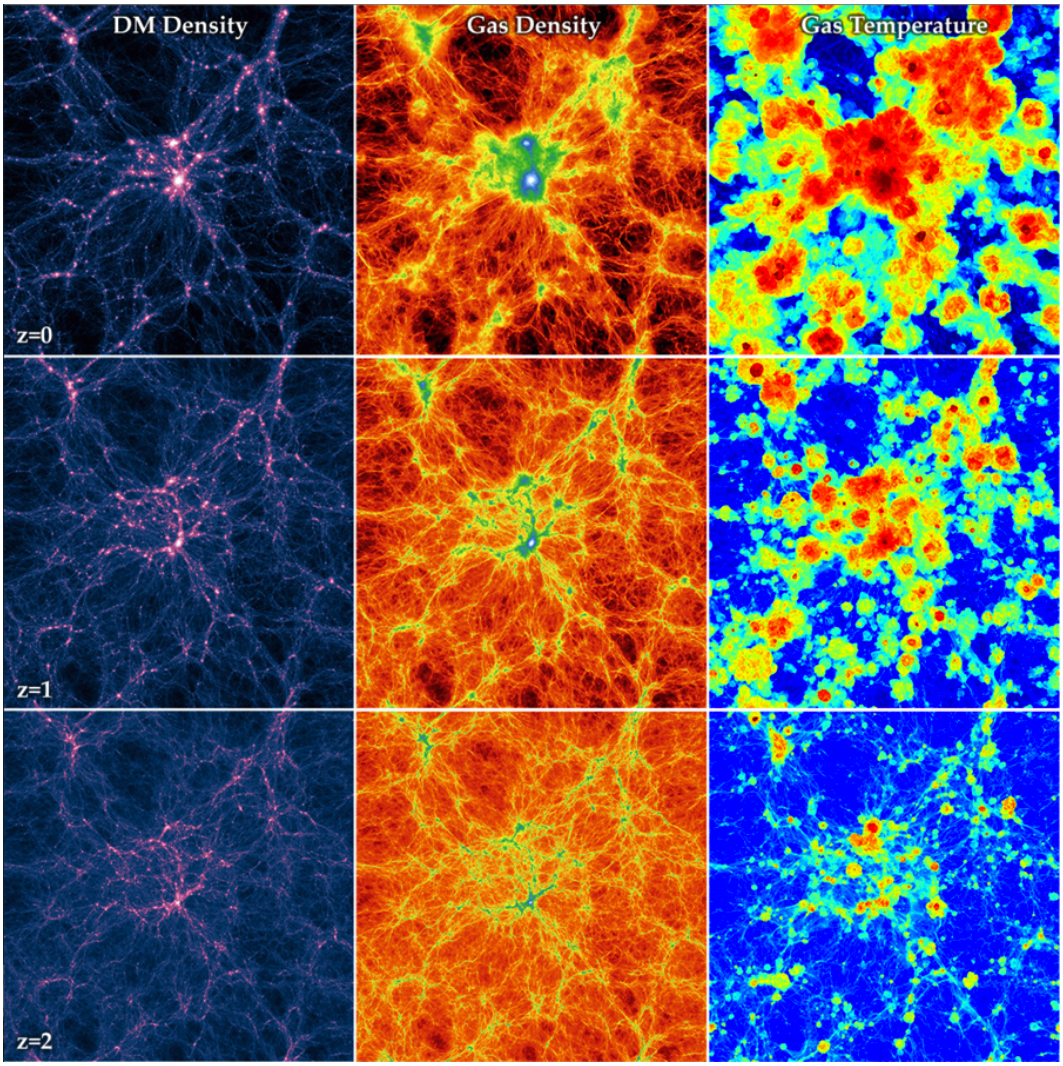

Figure 1.5: Redshift evolution of a whole box slice from $z = 2$ to $z = 0$, showing three projections: dark matter density, gas density, and gas temperature. (Illustris Project, 2017)

Simulations suggest that gas initially exists outside the virial radius of a dark matter halo and host galaxy. Over cosmic time, this gas accretes onto the halo and supplies the baryons that form new stars (Marinacci et al., 2010). For a large halo ($>10^{12} M_{\odot}$), this gas can become shock heated and will take a long time to cool and settle. The transition from the hot to the cold accretion mode is not sharp, however, and these halos will experience some cold mode accretion as well. For a less massive halo, the gas stays cool and can stream into the halo and galaxy (Dekel et al., 2009). There have been many estimates of the efficiency of gas accretion onto a halo via this mode, and this mechanism has also been cited in fueling ongoing star formation.

Real, directly-detected emission-line data that provides a three-dimensional view of the galaxy halo must be fed back into these simulations to refine the models and to verify that these modes are real (D. C. Martin, Chang, Matt Matuszewski, Morrissey, Rahman, Moore, Steidel, and Y. Matsuda, 2014; D. C. Martin, Chang, Matt Matuszewski, Morrissey, Rahman, Moore, and Steidel, 2014).

1.3 Detection of Emission Lines with FIREBall-2

Questions that can only be answered by emission line studies

The Faint Intergalactic-medium Redshifted Emission Balloon (FIREBall-2) has been designed to collect this 3D observational evidence. We will look for Lyman-α emission from the IGM and CGM. Cosmological simulations of large-scale structures suggest that the majority of the baryons at $z < 2$ reside in the Warm-Hot Intergalactic Medium (WHIM) at $T = 10^5$–10^7 K with moderate overdensities. In particular, we must distinguish between different morphologies of the WHIM.

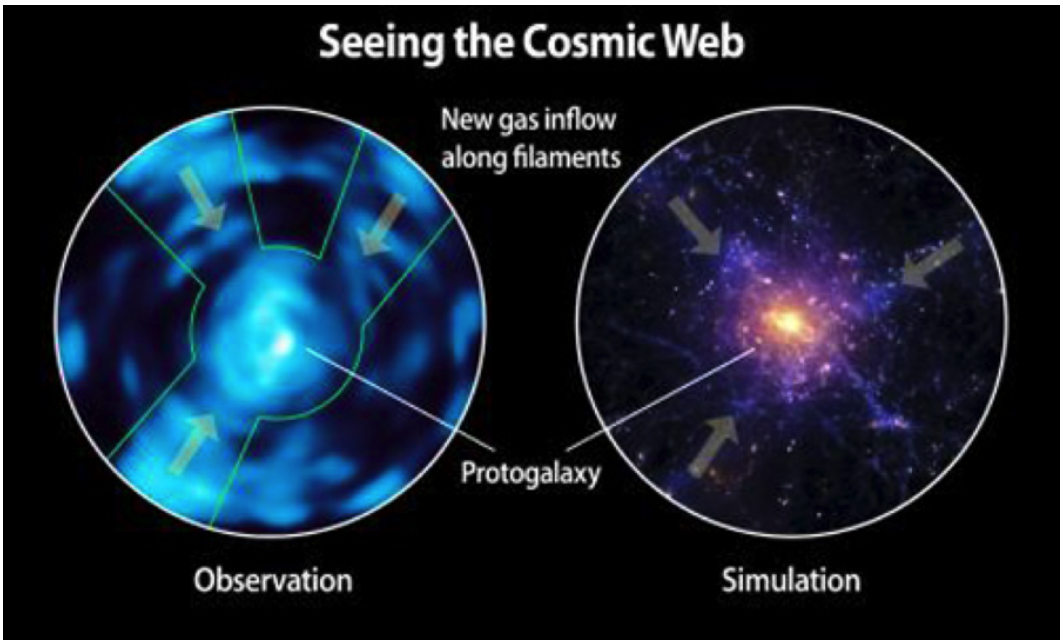

Figure 1.6: Direct detection map of the CGM captured by the Cosmic Web Imager, with a corresponding simulation on the right (D. C. Martin, Mateusz Matuszewski, et al., 2015)

Aside from Lyman-α, we can also observe metal lines (Oxygen, OVI 1033, and Carbon, CIV 1549) as well as several other strong metal line species in the CGM and WHIM (Jason Tumlinson et al., 2011). This gas cools radiatively via Lyman-α and metal line emission, and the spectrum of their emission changes with temperature.

The high ionization states will indicate hot gas and the low ionization species will identify colder gas. Then, by looking at the relative velocity shifts of these lines with respect to the systemic velocity, we can identify inflowing and outflowing gas.

The diffuse gas of the IGM is extremely dim and difficult to observe directly, yet it comprises nearly 75% of baryons. Thus, the vast majority of normal matter is essentially hidden from view. We aim to capture it by examining the brightest part of this dim matter, the CGM.

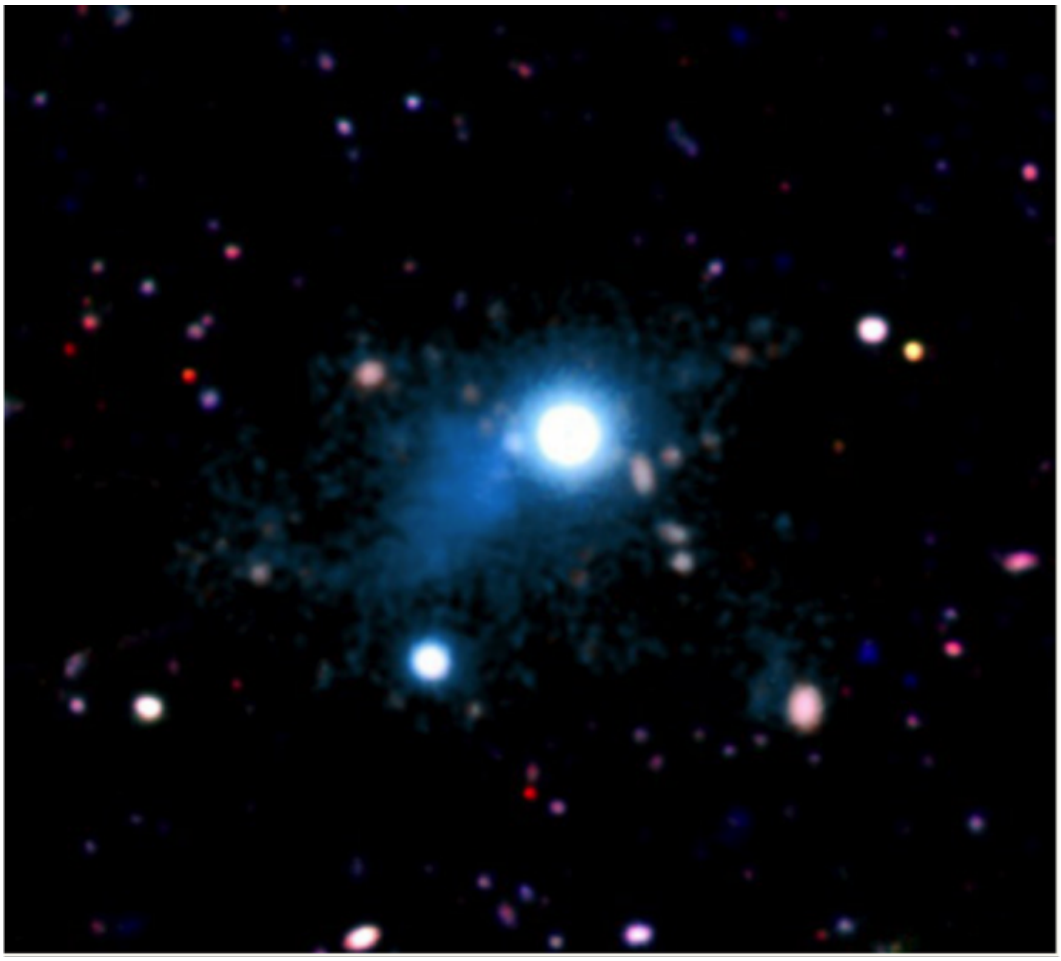

Figure 1.7: This nebula (cyan) was discovered around the bright quasar UM287 (at the center of the image) and extends across 2 million light-years. This deep image was obtained at the W. M. Keck Observatory (Credit: S. Cantalupo, UCSC)

Figure 1.7 shows a nebula discovered around bright quasar (UM287) spanning 2 million light-years. This quasar gives off strong radiation, which makes the surrounding intergalactic gas glow. The glowing gas reveals the morphology and physical properties of a cosmic web filament. The estimated amount of gas in this

nebula is at least ten times larger than computer simulations predicted. Recently, the Cosmic Web Imager (CWI), our very own Palomar instrument, has detected filamentary emission in the Circum-QSO Medium around a different quasar (QSO 1549+1919) and has found evidence for a cold inflow to the QSO (D. C. Martin, Chang, Matt Matuszewski, Morrissey, Rahman, Moore, and Steidel, 2014). For the upcoming FIREBall-2 flight, we will likewise choose one science field around a quasar, where the intense radiation might cause the extended hydrogen gas between galaxies to fluoresce, perhaps revealing a similar vast nebula of diffuse gas.

Emission line study of the circumgalactic medium & its host galaxy

The signal we are looking for is expected to be on the order of a few thousand photons cm^{-2} s^{-1} sr^{-1}. While we model the detection with simulations involving slits, there is uncertainty in the expected brightness of these lines. We will attempt to measure the intensity of escaping Lyman-α from the target galaxy as shown in Figure 1.8, which may give us some indication of the physical processes taking place within the galaxy and halo, its morphology, metallicity, and star formation rate (SFR) intensity. To do this, we need to determine the spectral signature of the host galaxy along with the emission from the surrounding CGM. With the extended emission spectrum, we should find some indication of winds by looking for the metal lines CIV and OVI (Frank et al., 2012). I will develop a camera that is sensitive at 205 nm, where Lyman-α at z = 0.7 is visible. In addition, CIV at z = 0.3 and OVI at z = 1.0 also fall into this band. Overall, the signal is very weak and we will require a high throughput spectrograph, good out-of-band light rejection, and a high efficiency, photon counting detector. Since the UV sky is fainter than the visible sky with a conventional detector, the measured noise will be dominated by the detector noise. With a standard imager, this noise would drown out the signal and thus we need a device with low noise and high quantum efficiency.

For our observing strategy, we have identified three distinct criteria to select galaxies that should host the brightest CGM regions: galaxy star formation rate, specific star formation rate, and galaxy/halo mass. With these in mind we selected two observing targets for the flight. Field 1 will use a single slit position for the entire exposure and Field 2 will use dithered microslits (as shown in panel d of Figure 1.8). The data will be dithered by ~40 arcsec in 5-arcsec steps to obtain a 2D spectroscopic "image" of Lyman-α from galaxies and their associated CGM. This objective necessitates a few other instrument requirements: we need very good pointing to stay on target over the integration periods and we need a new multi-object UV spectrograph to

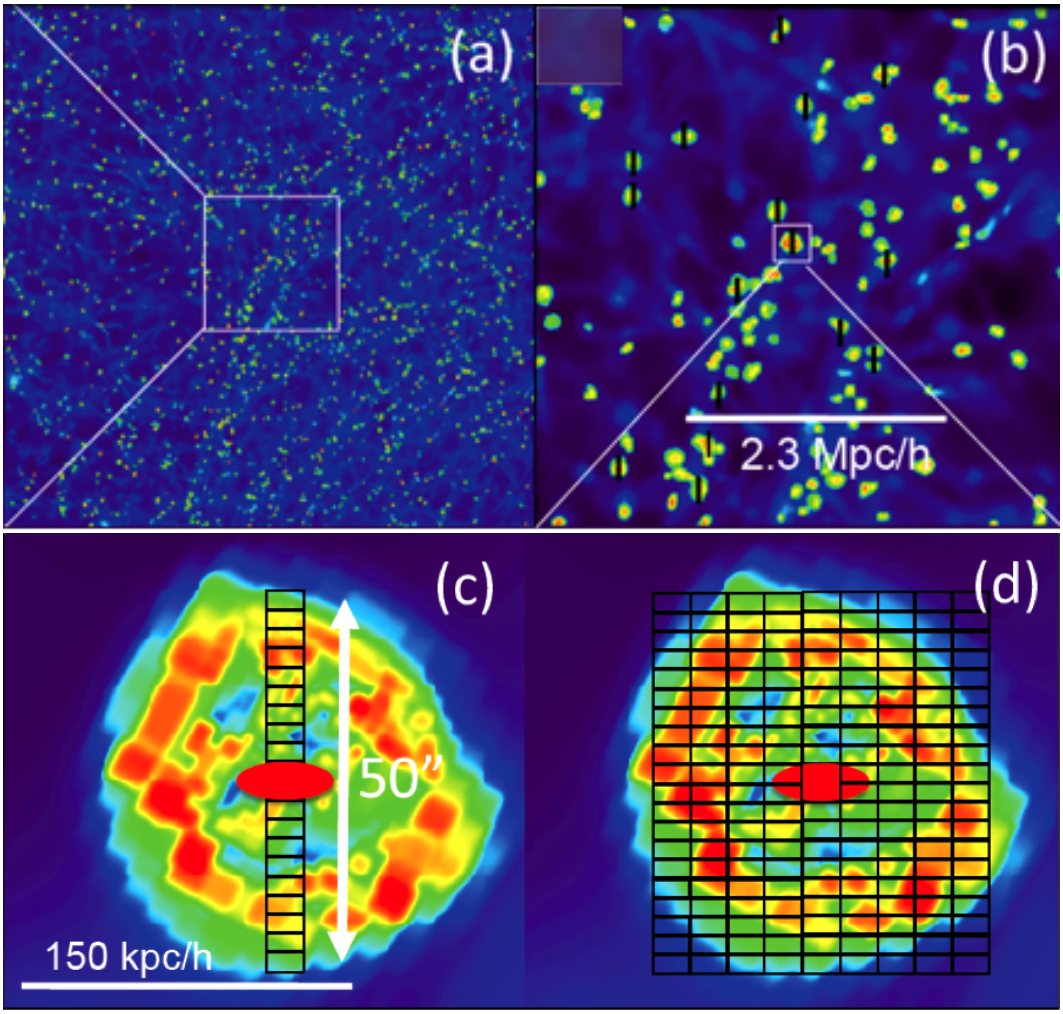

Figure 1.8: BINGO simulations of Lyman-α. We will choose several large-scale structures near a redshift of 0.7 by using ground-based redshift survey data with guidance from simulations like these. Because we will incorporate a multi-slit wide field spectrograph, we have estimated that there will be about 50 galaxies within the right redshift and luminosity range visible in each (11 by 35 arcmin^2) field, shown in image b. Panel c and d show zoomed-in looks at one of the 50 galaxies in the field demonstrating the two modes we will use to investigate the CGM. Credit: O'Shea, B.W., Introducing Enzo, an AMR Cosmology Application. ArXiv astrophysics e-prints, 2004.

distinguish between the light coming from the galaxy and the light coming from the CGM. We implemented several other improvements and the next chapter will discuss how the FIREBall-2 instrument builds on its predecessor, FIREBall-1. The table in Figure 1.9 describes the noise targets and requirements for the FIREBall-2 camera and cooling system, which flow down from the scientific goals described in this chapter. Each requirement will be discussed in the following chapters.

Noise Source:	Instrument Adjustments:	Noise:
Cherenkov Radiation Noise:		223 cts/cm2 s
high E particle directionality	**detector faces downward**	
dielectric medium distance	**larger vacuum tank: remove the window in front of the EMCCD**	negligible
blue dependence of photons emitted	**AR coatings reduce EMCCD transmission below 200nm (150nm)**	
Cosmic Web:		
faint UV emission signal	**photon counting detector: Silicon EMCCD with low noise & high QE**	
faint UV sky	**sky-limited**	0.01 ct/hr
Silicon EMCCD:		
UV absorption only in the surface layer	**DD improves EMCCD QE to about 30%**	
highly reflective	**AR coatings: (together with DD) brings QE to about 80% around 200nm**	
Detector Noise:		
Read Noise [RN]	**eliminated by using EMCCD**	
Clock Induced Charge [CIC]	**High Read-out speed & multi-level clocks Nüvü read-out electronics**	0.002 events/pix/frame
Dark Current [DC]	**Detector Temperature: -110˚C Exposure Time = 1000s**	0.001 events/pix/frame
Detector Temperature:		
balloon instrument cooling	**Mechanical Cryocooler**	
radiative & conductive heat load	**good Thermal Link performance**	
image degradation from frozen water layer	**High Vacuum [E-6 Torr] = lower water vapor pressure**	
High Vacuum:		
keep vacuum tank contaminants low	**low outgassing material only**	
vacuum leaks: make good O-ring seals	**smooth surface areas**	
minimize dewar ports	**PCB with internal wired feed-through**	
residual H2O, NI, O & H	**sorption pump: activated charcoal (cooled) & desiccant: zeolite**	
Mechanical Cryocooler:		
vibrational noise	**mech. mounting and coldtip decoupling (flexible copper straps)**	
excess internal heat	**mounting and wind tunnel + air fin solution**	
Thermal Link:		
external heat load	**minimize conducting material cross-sections (copper planes, wires, etc.)**	
external heat load	**minimize radiative load: gold plate copper will lower absorptivity**	
cooling geometries	**CCD pins are the only cross-sections that conduct heat**	

Figure 1.9: Noise summary for the FIREBall-2 scientific camera and cooling system. The cherenkov noise is discussed in detail in Chapter 2.2 and the noise value in this table reflects the original design of the instrument. I anticipate that the remaining radiation will be negligible due to specific instrument design changes. The detector noise sources have been reduced to the indicated values for an engineering grade device as well as several flight candidates. We are currently testing delta-doped devices to select a flight EMCCD with the best noise and imaging characteristics.

Chapter 2

FIREBALL-2 INSTRUMENT

2.1 Instrument Overview: FIREBall-1 vs. FIREBall-2

Gondola

FIREBall-1 was a 1-m telescope with a UV fiber-fed spectrograph (Tuttle et al., 2008). It used a GALEX (D. C. Martin, Fanson, et al., 2005) legacy NUV detector with modified pointing system from FOCA (Milliard, Donas, and Laget, 1991). The gondola stability system together with the 1.2-m siderostat allows for high pointing accuracy. FIREBall-1 was unable to detect emission from the CGM in three fields (Bruno Milliard et al., 2010), with limiting sensitivity of 20,000 LU (LU = photons cm^{-2} s^{-1} sr^{-1}). FIREBall-2 has been designed to overcome the limitations of FIREBall-1 and achieves significantly higher sensitivity. As shown in Section 2.1.

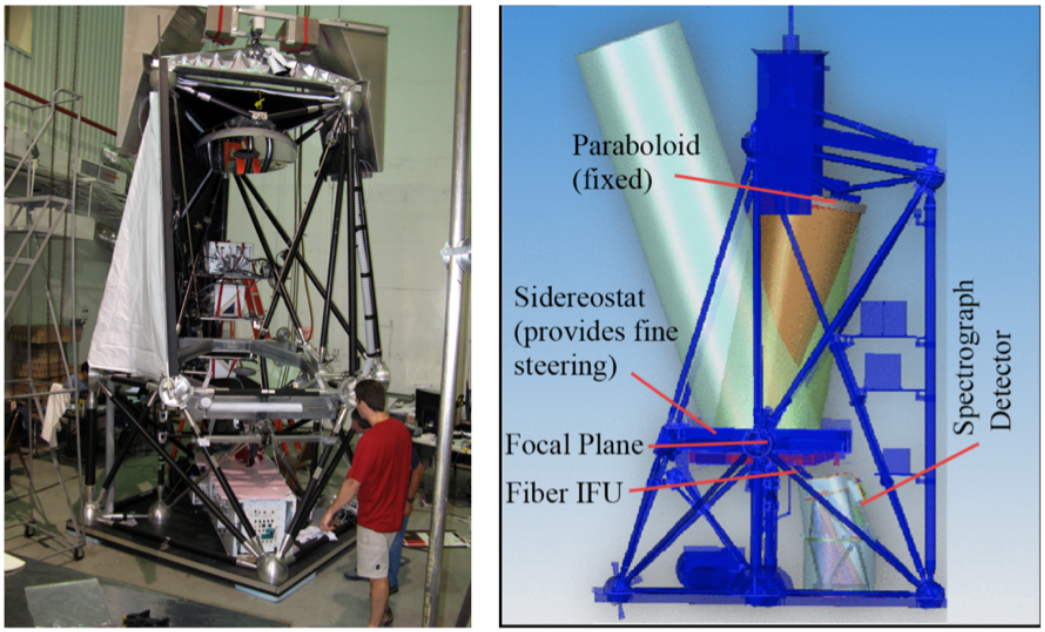

Figure 2.1: Interior of the FIREBall-1 gondola and its components. For FIREBall-2, the Fiber IFU and spectrograph tank have been removed and replaced with a new spectrograph.

FIREBall-2 builds on the proven FIREBall-1 payload and uses the same basic structure. This includes the balloon and carbon fiber gondola with its pivot mechanism,

which provides the coarse cross-elevation. Light enters the 1-m telescope, which consists of a 1.2-m plano siderostat that sits on a dual axis pivot assembly and reflects the light onto the fixed 1-m, f/2.5 parabolic primary mirror. The siderostat controls the elevation and provides stability over 45–70 degrees. The paraboloid reflects the light back down through an aperture in the siderostat and focuses it onto a 40-mm-diameter field lens. The fused silica field lens is the only window into the vacuum tank that holds the spectrograph and UV detector, for Cherenkov radiation considerations (see Chapter 2.2).

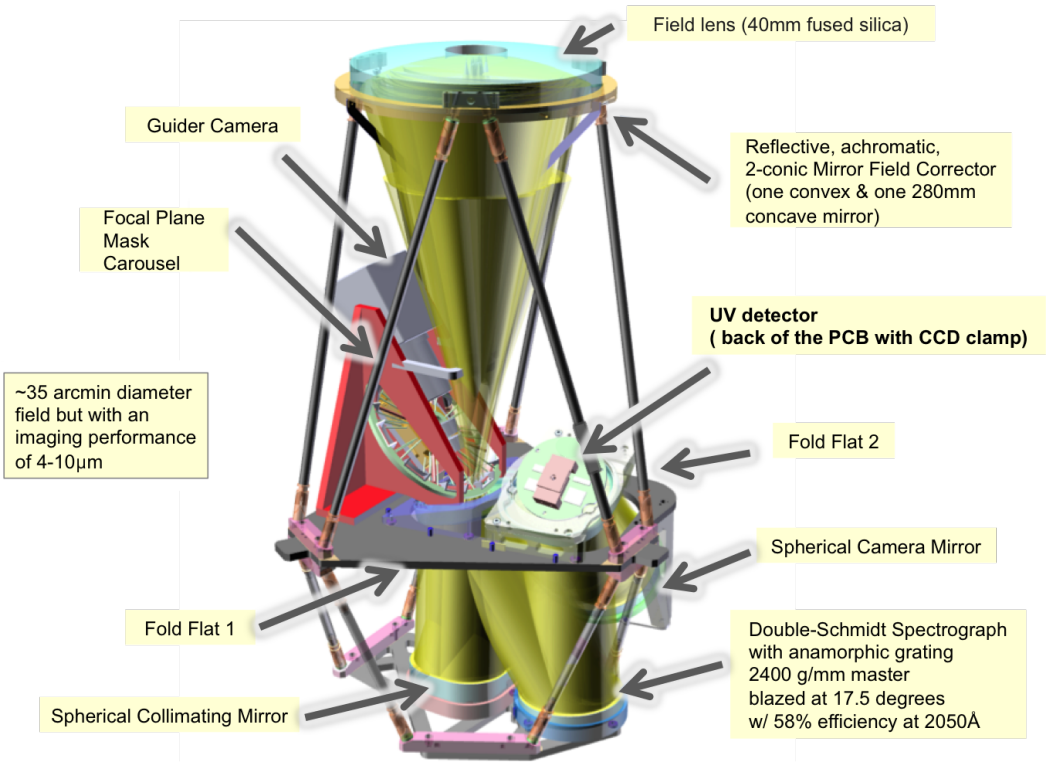

Figure 2.2: UV wide-field multi-slit spectrograph with an all-reflective design. It has a ~35 arcmin diameter field, but with an imaging performance of 4-10 μm. The camera presented here is on the right, facing downward and to the back.

Vacuum Tank

Inside the vacuum tank, the light is focused at the Mask/Guider System (MGS), which has a mask carousel with 9 slit-masks. These NiCoForm masks are 12 mm wide, 26 mm long and 40 microns thick. They are attached to mask holders, which sit on a circular carousel that can be rotated into the light beam for specific spectroscopic observations. The masks are laser cut with about 50 slits per mask, including masks for each science field, one star cluster, one long slit, one opaque,

and one open position. Another, thicker mask is used to reflect part of the light into the guider system, which can be seen on top of the carousel. The guider system is a PCO edge5.5 CMOS camera with two petzval-like lenses that focus light onto the camera, which is operated at a frame rate of about 30fps. Real-time guidance on stars, down to about 15th magnitude, is done in x and y as well as long-time guidance for rotation on a 1- to 5-minute baseline. The guidance software was originally developed by Matt Matuszewski for FIREBall-1 (M. K. Matuszewski, 2011). It performed flawlessly and has been adapted for this guider system to add the required long-time guidance for rotation as well as adding "lost in space" capability using astrometry.net (Lang, Hogg, Mierle, et al., 2009). Figure 2.2 shows the model design for the spectrograph, Mask/Guider system, and detector.

UV Wide-Field Multi Spectrograph

Once the light has been split by one of the masks at the focal plane, it enters the Double-Schmidt spectrograph with an anamorphic grating. The spectrograph is an all-reflective design.

The Focal Plane Mask is followed by a spherical collimator mirror, a fold flat, the reflective grating, a second fold flat, and the spherical camera mirror that directs the light into the UV camera. The camera presented here has been designed, manufactured, and built by the author. The Schmidt 1 Spectrograph and Folding Flat Mirror 1 have UV notch filters to absorb out-of-band light and reflect the FIREBall-2 bandpass. The other mirrors have an Al/MgF_2 coating to maximize the quantum efficiency within the FIREBall-2 band. The requirement was 85% reflectance and we have achieved 92% reflectance in the UV band.

Specific Improvements over FIREBall-1

The silicon detector with a special anti-reflective coating and a delta-doped layer has a response 10 times better than that of Flight 1 and Flight 2 of the FIREBall-1 microchannel plate detectors. The Al_2O_3 coating has been optimized around 205 nm, which is the spectral window for our observations (see Section 2.2) (E. T. Hamden, Jewell, et al., 2014) and brings the quantum efficiency (QE) to above 50% at a detector temperature of about -110°C. The noise of the science camera will be significantly lower as well, because it uses an Electron Multiplying CCD (EMCCD) that removes the (usually high) read noise, leaving only dark current and clock induced charge (CIC) (see Section 3.1) as noise sources. This instrument also incorporates two shutters to ensure low light exposure for the detector, since it is

very sensitive to low light levels.

The instrument efficiency was also drastically improved by removing the fibers from the system, and we can now better resolve the CGM regions. With this wide-field, multi-slit spectrograph, we will fit more CGM targets into our instrument. And with the help of newly developed anti-reflection coatings (E. T. Hamden, Greer, et al., 2011) we can now more accurately target the desired wavelength range.

Parameter	FB-1	FB-2	
Detector QE	0.05	0.7	Gen-1 → Gen-2
Instrument efficiency	0.0036	0.06	Fiber → slit, grating, cotings, etc
Atmosphere transmission	0.25	0.5	Afternoon flight, band narrowing
Net efficiency	0.0009	0.03	# CGM detectable ~ efficiency
Fiber/Slit Angular Diameter/width	8 arcsec	4 arcsec	To resolve CGM regions
Background Cts/PSF/hr	0.83 ct/hr	0.17 ct/hr[§]	CCD noise = CIC + dark (+ RN)
Sky Cts/PSF/hr [between NO bands]	0.01 ct/hr	0.56 ct/hr	FB2 is sky-limited
Effective Field of View	8 arcmin^2	900 arcmin^2	Field corrector, slit mask
Wavelength (Δz) range	150Å [0.12]	75Å [0.06]	Emphasize high t_{atm} band
Number of CGM regions/field	1	50	Multi Slit Mask, not an IFU
# of detectable CGM regions	0.1	~50/field	

Figure 2.3: Overview of the change in efficiencies from FIREBall-1 to FIREBall-2.

An effort was made to keep the total weight of the instrument lower than on the last flight by choosing appropriate electronics and hardware, where possible. The total weight, including gondola and ballast, is currently below 4500 pounds. With this reduced weight we should be able to reach a higher altitude on this flight, which will result in higher transmission through the atmospheric window (see Section 2.2).

2.2 Instrument Requirements

Instrument Signal to Noise Performance

The low surface brightness of our targets drives the low noise requirement. From multiple models and early observations of the CGM and its filaments, we expect to observe Lyman-α in the UV window with a signal strength as follows: the CGM survey will observe about 50 bright CGM regions around galaxies at z = 0.3, z =

0.7, and z = 1.0, which we can detect directly at 1400 LU and statistically at 200 LU. For the circum-quasar medium (CQM) survey we expect direct detection of the glowing gas at 5000 LU.

$$1 \text{ LU} = 1 \text{ photon s}^{-1}\text{cm}^{-2}\text{sr}^{-1}. \tag{2.1}$$

This relates to surface brightness as follows:

$$1\text{LU} \approx 1.16x \left(\frac{\lambda}{4000}\right)^{-1} x10^{-22}\text{ergs s}^{-1}\text{cm}^{-2}\text{arcsec}^{-2}. \tag{2.2}$$

A signal to noise of 10 is required to detect and characterize central galaxy regions (point sources). This means that the noise of the entire instrument has to stay below 140 LU.

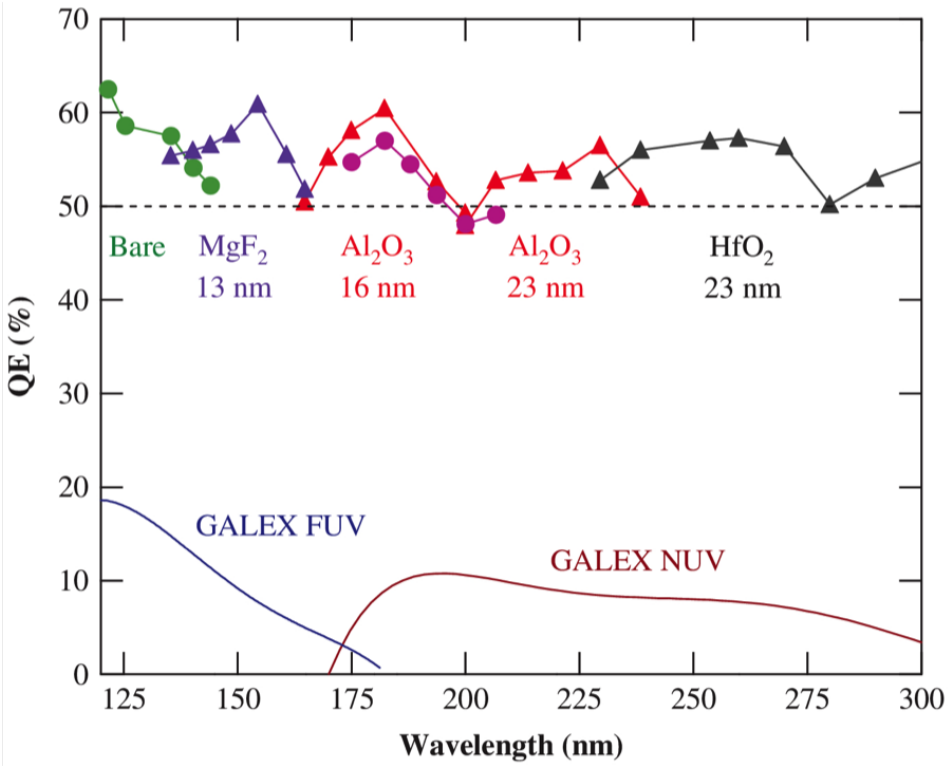

Figure 2.4: Response of the AR-coated, delta-doped arrays in the entire 100–300 nm region. Each region was coated with one simple coating. Further enhancements can be achieved by adding additional complexity in the coatings. In the third band, the response of AR-coated, delta-doped EMCCD (16 nm $Al_2O_3Al_2O_3$) is shown on the same plot as the conventional CCDs. (Nikzad, Michael E Hoenk, Grunthaner, et al., 1994)

The integration of an EMCCD is a major improvement over microchannel plates, which are traditionally used in space applications due to their low noise and high gain. However, they also have low UV quantum efficiency (~10%) and need high voltage power supplies. Due to recent efforts by Shouleh Nikzad and her group at JPL, the absolute quantum efficiency of delta-doped EMCCD arrays is over 50% in the near- to far-ultraviolet range for single photon counting applications as shown in Figure 2.4 (Nikzad, Michael E Hoenk, Grunthaner, et al., 1994). Specifically, the QE in our narrow band between 200 nm and 210 nm reaches a quantum efficiency of 56%, which we expect to be even higher for the flight device.

The read noise of this camera is negligible, and at a frame time of 1000 seconds the dark current has to be less than 0.003 events pixel^{-1} frame^{-1}. The CIC should be of the same magnitude, such that we remain sky-limited for our observations. Most of my camera noise effort has been focused on reducing the CIC and achieving low dark current noise (see Section 3.1).

The initial instrument design incorporated two CCD97-00 EMCCDs, rather than one CCD201-20. The CCCD97-00 has an imaging area of 8.192 mm by 16.384 mm and two sensors were needed to accommodate the required pixel field for the wavelength spread from the spectrograph. Due to the architecture of the chip, it was impossible for two chips to be butted together without losing imaging space at the center between the two chips. This would mean losing prime data with optimal center focus; to avoid this loss, we changed the design to incorporate a larger format CCD. The CCD201-20 has an imaging area of 13.3 by 26.6 mm and we were able to adjust the optics to focus all the data onto this sensor.

Cherenkov Radiation Analysis

Cherenkov Radiation Overview

Cherenkov radiation (Cherenkov, 1958) occurs when a charged particle passes through some dielectric medium at a speed (ν) faster than the speed of light (c) in that medium. Electrons in the material will interact with the particle and electromagnetic radiation is emitted due to a photonic shock wave. This photon shower is directional, emerging in a light cone (see Figure 2.5), which means that if there is a window close to the face of our detector, this radiation will spray uniformly over the CCD and appear as a background noise that we cannot identify and remove.

To address this issue, I conducted a study of potential contamination, or background noise, from Cherenkov radiation at our balloon altitude. To calculate the high-energy

particle count and directionality at flight altitude for a flight out of Fort Sumner, data was selected with a comparable Geomagnetic Cutoff Rigidity (θ_M). θ_M depends on the magnetosphere of Earth and therefore varies with location. Particles with less rigidity than the cutoff will be deflected before they reach a given point in the magnetosphere, whereas those with more than the cutoff will penetrate it. In our case this will manifest as a low energy drop-off in the cosmic ray flux for kinetic energies in the range from around 10 MeV to about 10 GeV. Only particles in an energy range that would affect the camera were considered. By combining this information with other high-altitude, high-energy particle data, I was able to estimate the resulting total UV photon flux on the detector face and make recommendations on how to reduce this noise.

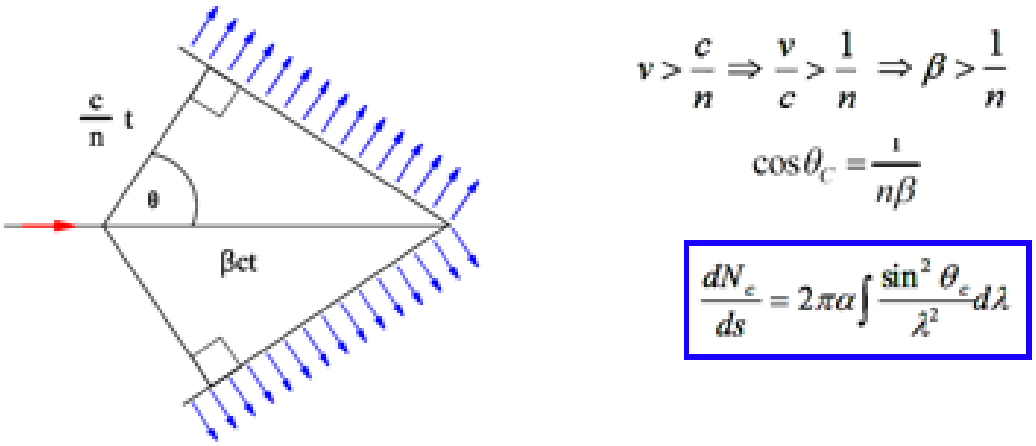

Figure 2.5: 2-D depiction: a charged particle (red arrow) enters the medium on the left with speed v and the resulting radiation (blue arrows) exits the medium at an angle θ. n is the index of refraction of the medium and c is the speed of light. The formula gives the number of photons emitted (blue arrows) per path-length at an angle θ (fine structure constant $\alpha = \frac{e^2}{\hbar} = \frac{1}{137}$ and wavelength λ, n=index of refraction of the medium, $\beta = \frac{v}{c}$, (Commons, 2017).

Parameters, Frank Tamm Formula, & Calculations

The only dielectric medium relevant to our instrument is the entrance window to the vacuum tank. It is made of fused silica and at the time of my calculations I used the dimensions of the original design, which were 5 cm by 7 cm by 5 mm. We have since changed the design to a round window with a diameter of 2 cm. The total surface area of the window, including the sides, is 47 cm^2. Fused silica has an index of refraction of $n = 1.47$ and a transmission window from 100 nm to about

5 μm. Silicon detectors have a wavelength cutoff at around 1000 nm, which leaves a standard detector behind a fused silica window with a wavelength range of about 0.1 μm < λ < 1 μm.

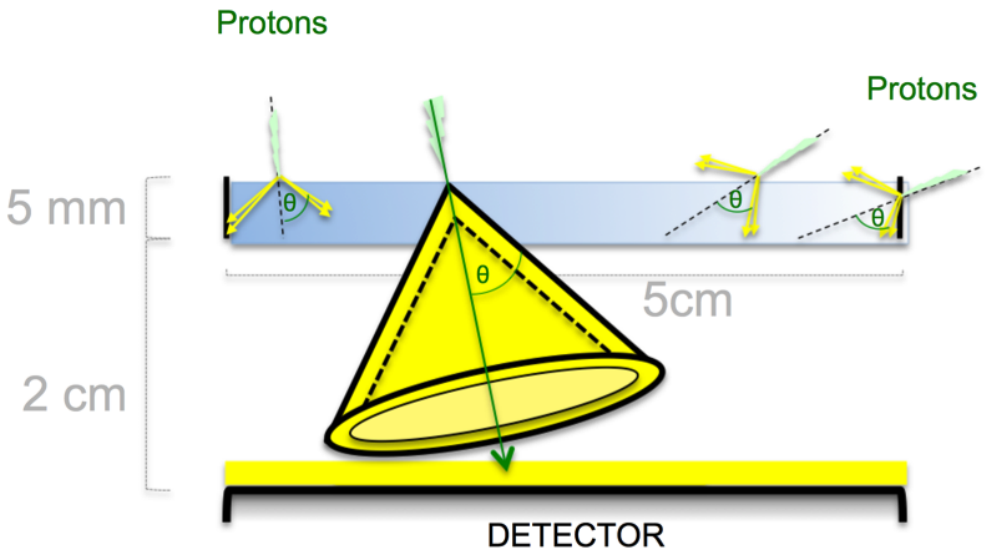

Figure 2.6: Cartoon of the window-detector geometry (yellow indicates the Cherenkov photons). Note that only half of the Cherenkov photons will enter the detector, due to the incoming angle of the protons.

The high-energy particles that we will encounter at our balloon altitude are predominantly protons with a proton rest mass (mc^2) of 0.938 GeV and the geomagnetic cutoff rigidity at Fort Sumner is 4.5 GV, which sets the lower bound for the proton energies that we will encounter during our balloon flight (Protons in near earth orbit AMS Collaboration).

$$\text{Rigidity} = \frac{M}{Z}\sqrt{E^2 + 2mE}. \tag{2.3}$$

Thus, protons with energies below 3.66 GeV are shielded by the Earth's geomagnetic field at Fort Sumner and will not penetrate to the lower atmosphere where we will conduct our experiment. This sets the energy range of the protons selected from the dataset taken by BESS (Balloon-borne Experiment with Superconducting Spectrometer).

To calculate the total photon flux from protons incident on our window, I used the Frank and Tamm formula (Kleinknecht, 1998), which gives the number of photons (N) in our wavelength (λ) interval for incoming protons of energy (E), for a medium with thickness l, including the fine-structure constant $\alpha = \frac{1}{137}$:

$$\frac{dN}{ds} = 2\pi\alpha \int_{\lambda_1}^{\lambda_2} \frac{\sin\theta^2}{\lambda^2} d\lambda \tag{2.4}$$

$$N = 2\pi\alpha l \left(\frac{1}{\lambda_1} - \frac{1}{\lambda_2}\right) (\sin\theta)^2. \tag{2.5}$$

With the geometric relation of the Cherenkov radiation from Figure 2.5 above we see that the Cherenkov angle θ can be derived as follows:

$$\cos\theta = \frac{(c/n)t}{\beta ct} = \frac{1}{\beta n} \tag{2.6}$$

so:

$$(\sin\theta)^2 = 1 - (\cos\theta)^2 = 1 - \frac{1}{(\beta n)^2}. \tag{2.7}$$

The proton energy is as follows:

$$E^2 = (pc)^2 + (m_0c^2)^2, \tag{2.8}$$

with $E << m_0c^2$:

$$E^2 \approx (pc)^2 \tag{2.9}$$

$$E^2 = \frac{\left(m_0c^2\right)^2}{1 - \frac{v^2}{c^2}}. \tag{2.10}$$

$\frac{v^2}{c^2}$ is well known as β. Rearranging, we find:

$$\frac{1}{\beta^2} = \frac{E^2}{E^2 - \left(m_0c^2\right)^2}. \tag{2.11}$$

Replacing these terms in the original equation yields:

$$N = 2\pi\alpha l \left(\frac{1}{\lambda_1} - \frac{1}{\lambda_2}\right)\left(1 - \frac{E^2}{\left(E^2 - \left(m_0c^2\right)^2\right)n^2}\right). \tag{2.12}$$

The derived formula yields the amount of radiation emitted by a medium with index of refraction n within a specific wavelengths range for an incoming particle of energy E. I determined the wavelength range by using parameters appropriate to the fused silica transmission window, which has a UV cutoff at around 100 nm, and the silicon detectors, which have a detection range out to about 1000 nm.

Choice of Particle Data

The spectra data that I selected from measurements of cosmic ray protons and helium from BESS were provided by Richard Mewaldt from Caltech's Cosmology Group Shikaze et al. (2007) and are shown in Figure 2.7.

Spectra measurements of cosmic ray protons taken with the Alpha Magnetic Spectrometer (AMS) during space shuttle flight STS-91 were also included. As Figure 2.8 shows, the AMS data consists mostly of lower energy secondary protons. For an in-depth discussion of primary and secondary particle fluxes at Fort Sumner, please refer to Abe et al. (2003). Primary cosmic rays enter our atmosphere from space, and as they travel through our atmosphere they interact with nuclei and produce secondary particles with lower energy (Shikaze et al., 2007). Only primary (high energy) protons have a downward direction, whereas secondary (lower energy) protons are traveling upward (see Figure 2.8).

These measurements were made in near-Earth orbit at 380 km from the surface, and this dataset is a good counterpart because the downward flux at this altitude is comparable to the balloon (BESS) data.

Data blocks b) and e) in Figure 2.8 give the particle count for the Geomagnetic Cutoff Latitude (θ_M) at Fort Sumner, where our balloon experiment will take place. Adding up the particle count in the correct θ_M range shows that BESS recorded about 66×10^{-3} cts GeV^{-1} cm^{-2} sr^{-1} s^{-1}, compared to AMS's count of about 76×10^{-3} cts GeV^{-1} cm^{-2} sr^{-1} s^{-1} downward and 1×10^{-3} cts GeV^{-1} cm^{-2} sr^{-1} s^{-1} upward. This ratio shows that the downward flux of primary protons is much higher than the upward flux of secondary protons.

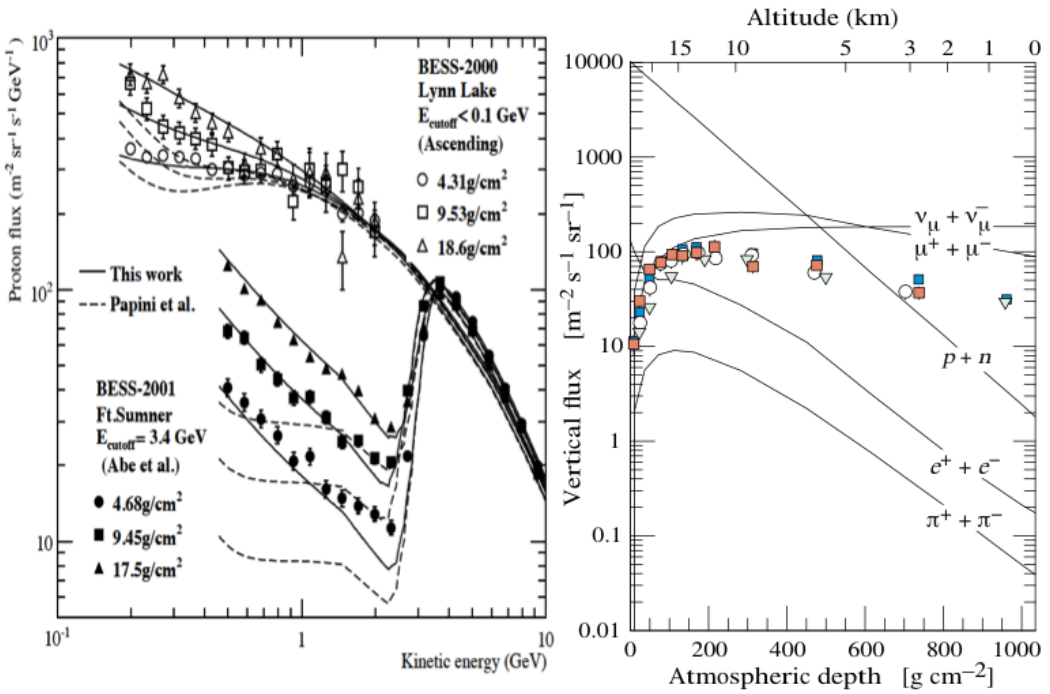

Figure 2.7: Left: Vertical fluxes of cosmic rays in the atmosphere with E > 1 GeV estimated from the nucleon flux from Shikaze et al. (2007). Right: Proton energy spectra at several atmospheric depths. Obtained spectra measured during the ascent in 2000 (open markers), and the descent at Ft. Sumner in 2001 (closed markers) are compared with calculated spectra before (dashed lines) and after (solid lines) modification of the recoil proton production energy spectrum.

These Cherenkov radiation results are compatible with data collected by Bruno Milliard, one of our collaborators at Laboratoire d'Astrophysique de Marseille (LAM) in France, as well as the quartz material results of a high energy group in Helsinki (quartz and fused silica have a similar index of refraction, Österberg (2008)). Milliard, informed by data taken with an upward facing detector from our previous FIREBall-1 flight, calculated a background noise figure of 400 cts GeV^{-1} cm^{-2} s^{-1} ± 200 cts GeV^{-1} cm^{-2} s^{-1} that he suspected was mainly Cherenkov radiation. My calculations fall within that range with a UV photon shower of about 200 cts GeV^{-1} cm^{-2} s^{-1} (upward facing detector), which is a noise background that we cannot afford.

Results and Instrument Design Recommendations

Clearly the upward flux of particles with energies that will produce Cherenkov radiation in the UV is much lower than the downward flux. Counts from the sides will be lower as well, because a larger volume of air (containing more molecules) has

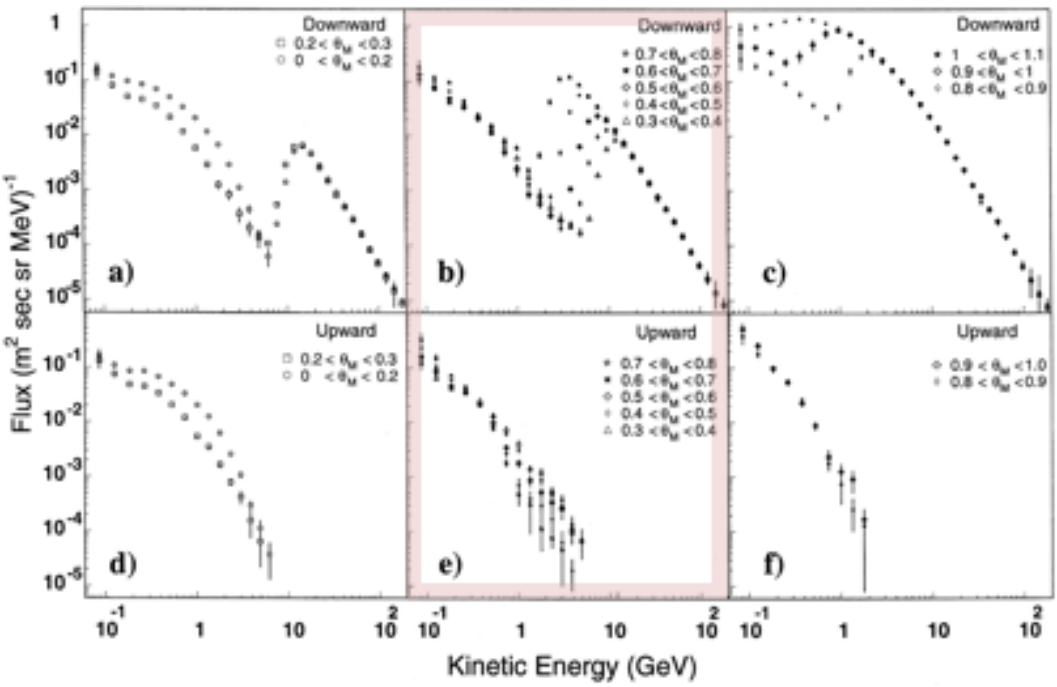

Figure 2.8: Vertical fluxes of cosmic rays in the atmosphere with E > 1 GeV estimated from the nucleon flux, according to θ_M (Alcazar et al., 2000).

been traversed and we are already looking at secondary lower energy particles from these directions. Therefore, I recommended a vacuum tank-spectrograph design, which accommodates a downward facing detector. A further reduction in Cherenkov noise can be achieved if we place the imaging detector far from the silica window, since the conical shape of the emitted radiation concentrates the flux of photons immediately behind the window.

We can further improve on the noise by simply blocking Cherenkov radiation photons from entering the detector. The number of photons increases exponentially in the UV, and if we limit the blue sensitivity of the chip, we greatly reduce the photon flux. A selection of sample wavelength cut-offs in the blue illustrate this drop: λ_1=100nm > 223 cts cm^{-2} s^{-1}; λ_1 = 150 nm > 140 cts cm^{-2} s^{-1}; λ_1 = 200 nm > 99 cts cm^{-2} s^{-1}. To address this concern, I suggested that an effort be made to develop anti-reflective (AR) coatings for the detector that specifically reduce the flux of Cherenkov noise radiation.

Based on these recommendations, our team made changes to the original spectrograph design. The flight instrument now features a detector location that is far from the silica window and light is focused into a downward-facing detector. JPL, in col-

laboration with Erika Hamden then of Columbia University, designed AR coatings that not only reduce the flux in the red, but also cut the incoming radiation below 200 nm by 80% (Jewell et al., 2015).

Gondola Altitude and UV Atmospheric Transmission Window

Another important aspect of the balloon payload is the overall weight, because it determines the maximum altitude of the gondola. We have spent a lot of effort reducing various noise sources in the instrument, but the atmosphere itself attenuates the UV signal by roughly 50%. As this is the peak of the UV transmission window in the stratosphere (see Figure 2.9), our FIREBall-2 instrument is designed to look at a narrow bandpass around 205 nm. Ozone (O_3) and oxygen gas (O_2) are the main absorbers in the UV, with a few nitric oxide bands in our observing bandpass (see Figure 2.9). Ideally our balloon will reach the top of the atmospheric ozone layer, where the transmission improves dramatically. We can gain about 5% in transmission for every kilometer of height change (M. K. Matuszewski, 2011) without significant increase in Cherenkov radiation. Most hardware has been designed and/or purchased with this requirement in mind and we estimate a total weight, including gondola and ballast, of less than 4500 lbs. By scheduling an afternoon launch, we can also increase the amount of helium that the balloon can retain due to loss of helium for zero pressure balloons (Space Sciences, 2017) launched early in the day (M. K. Matuszewski, 2011).

My specific recommendations for the instrument are:

- Downward-facing detector
- Light entrance window far removed from the detector
- Blue cut-off coating for the detector
- Gross weight below 5000 lbs, including all hardware and ballast
- Detector noise floor below 0.003 events pixel^{-1} frame^{-1}

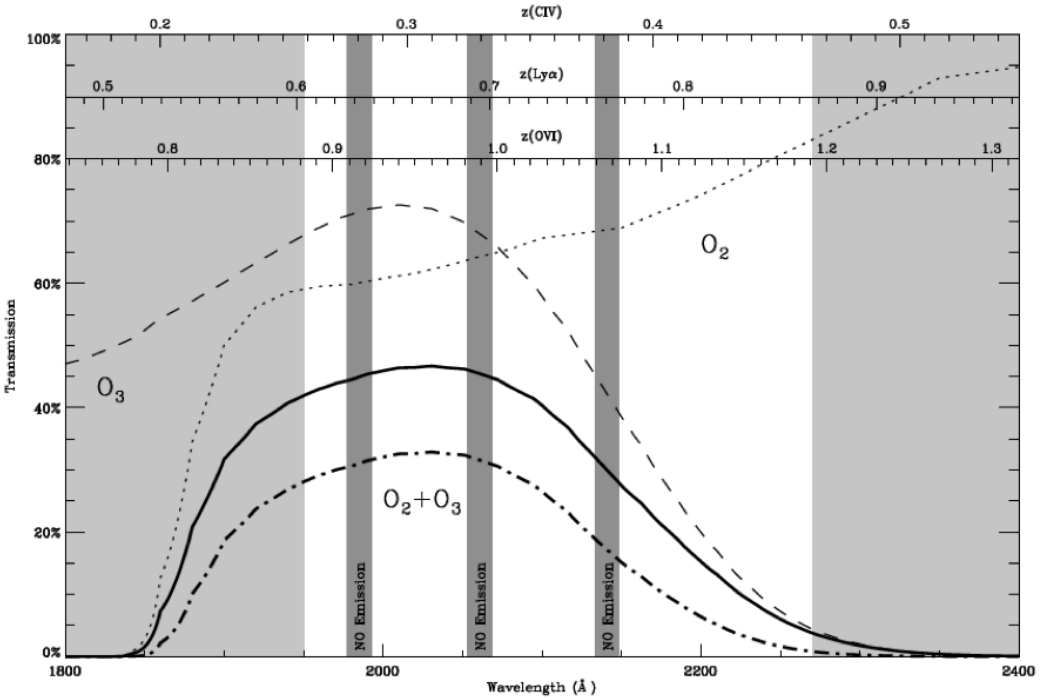

Figure 2.9: Atmospheric transmission curve for an altitude of 34 km, showing the Oxygen and Ozone transmission window for an elevation of 70° (solid curve) and 40° (dot-dashed curve). The light grey areas show the end of the transmission window for the FIREBall-1 instrument, the dark grey areas are nitric oxide airglow bands, and three red-shifted emission lines of interest are plotted on top of the image (M. K. Matuszewski, 2011).

Chapter 3

SCIENTIFIC CAMERA & COOLING SYSTEM

3.1 Requirements and New Instrumentation

New Sensor Technology: e2v CCD201-20

The faint signal of our science target in combination with the desired wavelength range (200–210 nm) drive the development of a new type of detector. We chose to use the low light level (L3) CCD technology from e2v (Technologies, 2017b), because it is specifically designed for photon-starved observations (Basden, Haniff, and Mackay, 2003), and amplification of the collected signal takes place on the chip while it is still in the charge domain. This allows for real-time operation of the EMCCD with sub-electron readout noise, permitting us to image the very dim filaments and CGM regions. e2v's CCD201-20 Back Illuminated 2-phase IMO series electron multiplying CCD sensor (Technologies, 2017a) was selected as the base sensor architecture to develop a high quantum efficiency, low noise, photon-counting UVCCD in collaboration with JPL's Micro Devices Lab (MDL). This new detector is called the next generation detector for UV spectroscopy (NEXUS).

Commercial e2v L3 CCD201-20 devices are customized at the wafer level with back illumination processes developed at JPL. This method includes Molecular Beam Epitaxy (MBE)-based superlattice and delta doping with Atomic Layer Deposition (ALD)-based custom anti-reflection (AR) coatings and integrated filters on EMCCDs (Nikzad, Michael E Hoenk, Greer, et al., 2012). The CCD is 13 mm by 26 mm and the actual pixel is 13 microns on a side. The device is a 1K by 2K EMCCD, as there are 1024 by 2048 pixels in the imaging area. The EMCCD has high voltage implants in the serial (horizontal) register, with a maximum voltage of 45 V (see Section 3.2). The Quantum Efficiency (QE) of the finished Aluminum Oxide AR coated, delta-doped device is between 50% & 80% at around 200 nm (depending on coating, refer to Jewell et al. (2015)), which is the UV photon energy range we need for our observations. The read noise is negligible, but the CIC will be about one thousandth of an electron per frame. By cooling the CCD to -110°C, we should be able to suppress the dark current to about the same low number for a 1000-second integration. The reason we have to integrate for this duration is to raise the signal level above the CIC floor for a 5.5σ cutoff. The extremely high

L3 CCD Design	
Parameter	Value
Format	13 x 26 mm^2
Pixels Size	13 x 13 microns
Pixel #	1024 x 2048
Maximum Voltage	45V
QE	50-70% (200 nm)
Read Noise	<1E -6 e^- $pixel^{-1}$ $frame^{-1}$
CIC Noise	1E -3 e^- $pixel^{-1}$ $frame^{-1}$
Dark Current (1000s)	<1E -3 e^- $pixel^{-1}$ $frame^{-1}$
Temperature	-110˚C
Frame rate	1 frame/(100 -1000 sec)
AR coating	Al_2O_3

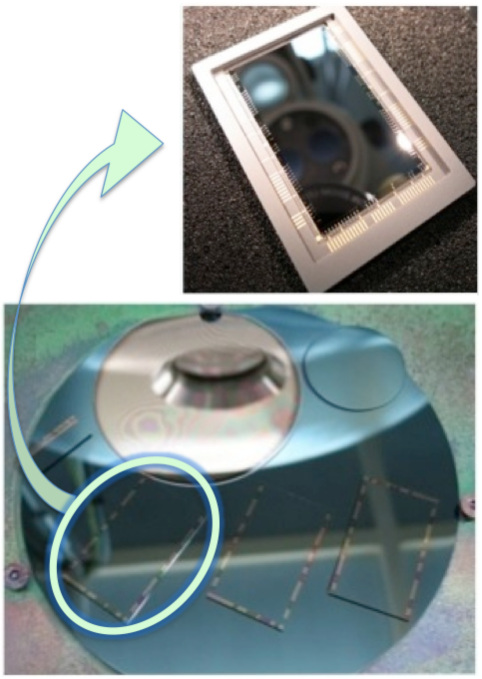

Figure 3.1: How a chip goes from wafer (bottom right) to packaged device (top right). This is a photograph of an anti-reflection (AR) coated DD-EMCCD, with the FIREBall-2 NEXUS parameters (left).

detector QE is a result of delta doping the EMCCD, which gets us to about 30% QE, and anti-reflection UV coatings on the detector. These coatings are enabled by the delta doping and will get the QE up to 80%. The delta doping process is the main contributor to the NEXUS (Nikzad, Hoenk, et al., 2013).

EMCCD stands for Electron Multiplying Charge-Coupled Device. An EMCCD camera operates by first transforming incident photons into photo-electrons (via the photoelectric effect) inside the bulk of the silicon body. The semiconducting surface is organized into an array of pixels, a grid of potential wells in which electrons can accumulate during an exposure. Negative charges are collected inside each pixel and then a series of voltages are applied to the wells. This forces the charge to travel vertically across the pixel grid, row by row, a process known as "Charge Coupling" (black arrows Figure 3.2). The electrons that arrive at the bottommost row of the CCD now travel horizontally through the serial, or horizontal, register to the right (curved black to red arrow, top row Figure 3.2). They do so again, pixel by pixel into the multiplication register, after the corner elements to the left (curved black to red arrow, bottom row Figure 3.2). Within this multiplication stage, the accumulated

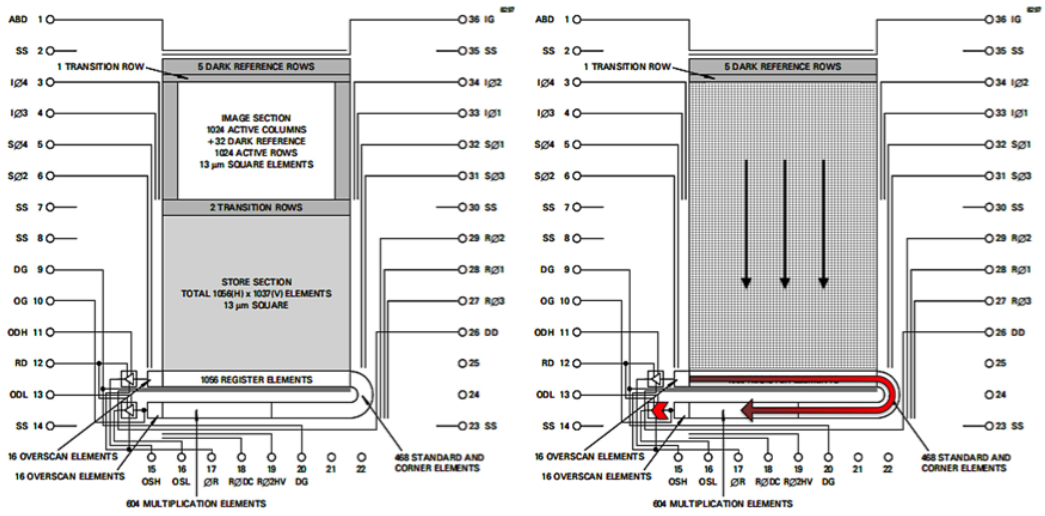

Figure 3.2: CCD201-20 schematic: Left: image area on top and covered storage area on the bottom (left). Once the delta doping and AR coating has been completed, the entire CCD will be used. Right: Cartoon demonstrating the EMCCD read-out scheme.

photoelectrons are amplified by impact ionization, which is an avalanche effect that generates several thousands of electrons per photoelectron. The multiplied charges finally reach an amplifier (short red arrow Figure 3.2), where they are converted into an electric impulse that is then digitized to form a high quality image. The charge package that reaches the output amplifier at the end of the multiplication chain is now much larger than the read-noise, which is why EMCCDs can achieve sub-electron read-out noise. However, the quality of output data is still affected by the read-out noise, as discussed in the next chapter.

The arrangement of different voltages is referred to as "clocking signals" and a waveform script defines in what order and at what level these voltages are applied to the CCD pins to move the charge efficiently from one pixel to the next. Most of my effort has been spent on optimizing the waveform script to reduce various noise sources and achieve maximum quantum efficiency in the detector.

EMCCD Noise Sources

During the movement of the charge packages from one pixel to the next, electrons can be generated, which is what we call CIC, or spurious charge (James R Janesick, 2001). Ironically, impact ionization is the source of the extra unwanted charge, which is the same mechanism an EMCCD camera uses to amplify signal in the multiplication register. There is always a small probability that during the charge

transfer between pixels some unwanted and later wanted extra charge is generated. This is the main noise source for EMCCDs, which require special read-out electronics to reduce the CIC levels. CIC is present in conventional CCDs as well, but it is drowned out by the much larger read noise, which is added to the signal by the output amplifier at the end of the read-out chain. EMCCDs are specifically designed for low light level observations by amplifying the incoming signal above the read-noise floor. To make quality measurements, the faint incoming signal cannot be saturated with CIC and must be amplified without degrading the final data, important factors for the read-out electronics. Another known noise source in silicon detectors is dark current, which can be lowered by cooling the device.

Before discussing the choice of controller and overall camera design, I will describe the main factors that govern detector noise and drive the controller requirements and options.

Excess Noise Factor

EMCCDs amplify the small signal that travels through the horizontal multiplication register by adjusting the voltage scheme of the second part of the serial register (see Figure 3.2). Serial register clock 2 pixels are replaced with high voltage clock (R∅2HV) pixels and a fixed low voltage (R∅DC) clock is added just before R∅2HV to create a large, repeatable voltage drop. CCD201-20 has 604 such elements and there is a low probability that impact ionization will create additional electrons, every time photoelectrons are moved from R∅DC to R∅2HV. Even though the probability is low for one event, the overall gain at the end of this multiplication chain is large and depends on the specific voltage drop used. However, the probabilistic nature of this signal multiplication makes EMCCD measurements stochastic and decreases the signal to noise ratio.

We can only know the mean gain, not the exact gain, for the charge of a given pixel. This introduces an error in the photometric measurement because the number of electrons that are measured at the end of the amplification process now carries an uncertainty. This is the excess noise factor (ENF), which shows up in the signal to noise ratio under the root in the denominator:

$$SNR = \frac{S}{\sqrt{ENF^2 S + \frac{\sigma_{real}^2}{G^2}}}, \tag{3.1}$$

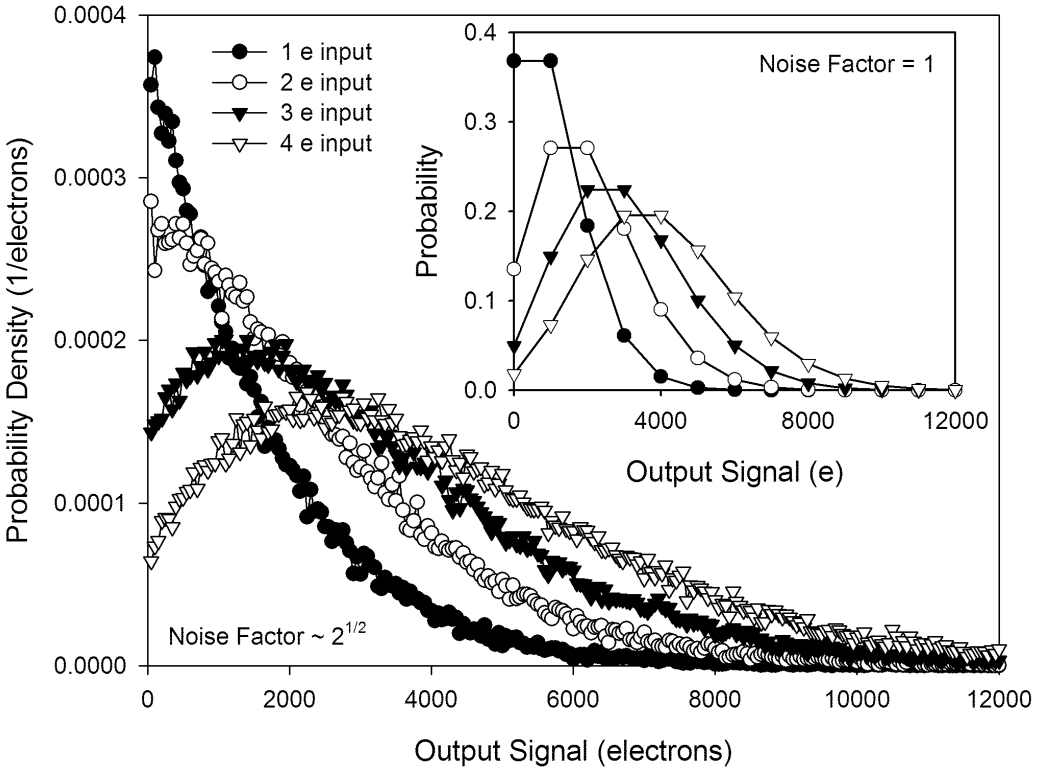

Figure 3.3: Simulated output signal distribution for a range of mean input signal levels (zero dark signal). Input follows a Poisson distribution. The signal passes through 536 multiplication stages providing a total mean gain of 1000. Inset shows the probability versus output assuming a noise factor of unity. Figure and caption taken from Robbins and Hadwen (2003).

where S is the number of electrons, G is the mean overall gain at the end of the multiplication register, and σ_{real} is the real read-out noise of the EMCCD. ENF^2 reaches a value of 2 with high gain and its effect on the SNR is the same as if the QE of the EMCCD were halved (Robbins and Hadwen, 2003).

Equation 3.1 shows that there are overlapping output probabilities for the different number of input electrons (at a given gain). As mentioned above, the ENF will degrade the signal by $\sqrt{2}$, but if we collect only one photon per pixel, the noise factor becomes 1 and we recover the standard SNR. (Robbins and Hadwen, 2003). Our observations will be extremely photon starved and we expect no more than one photon per pixel on average, with one event corresponding to one photon. This means that we are able to maintain the standard signal to noise of the silicon device, without the noise factor.

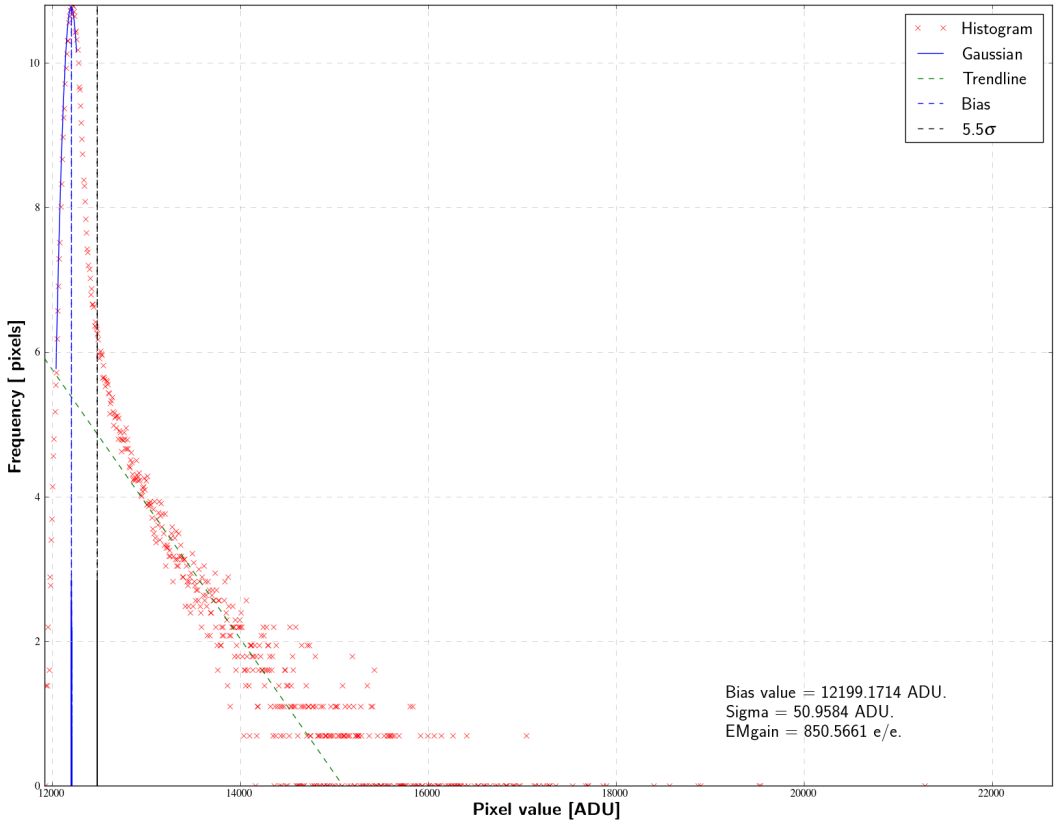

Figure 3.4: One of our histograms showing a 5.5σ threshold to distinguish between real and false events.

In order to evaluate the data collected with an EMCCD, we set a signal threshold; only pixels with counts above this threshold are considered to be an event (Daigle, Carignan, et al., 2009). The threshold is solely determined by the read-out noise of the device. We have to adjust the gain to find a balance between counting false events from the read-out noise, events lost due to low amplification, and coincident loss. Coincident loss occurs when more than one photon enters one imaging pixel. The final charge package of that two-photon pixel will still result in an event that we have defined as one photon, after the pixels are read out. Therefore we lose information about how many photons actually entered that single pixel. We will also lose events under the read-noise floor, if the charge from a given pixel was not amplified enough to get above the set threshold. Generally, a threshold of 5.5σ will exclude false events from the read-out noise without losing too many real events that were not amplified enough (Figure 3.4).

We work with an electron multiplication gain (G) of about 1000 e/e to amplify our faint signal above the read-noise floor. To achieve this gain in my camera system,

the HV clock is set to about 40 V, with a R∅DC level of 3 V, in order to amplify the single photon signal enough that it can be extracted from the read noise floor, while keeping the CIC low enough that it does not degrade our imaging quality.

Clock Induced Charge

Read noise has traditionally been the dominating noise factor for CCDs, which is why they have not been viable for low light level applications. This new technology, as described above, allows for very sensitive measurements, as long as the remaining noise is sufficiently suppressed. CIC (charge transfer generated charge) in EMCCDs can be significantly reduced through several methods, described below.

CIC reduction:

1. **Clock voltage wells:** Clock voltage wells must be reduced as much as possible without compromising the charge transfer efficiency (CTE). The voltage difference between the clock high level and the substrate bias level of the chip is important and should remain at a minimum.

2. **Clock speeds:** Parallel clocks will generally contribute more CIC, which can be mitigated by slowing them as much as possible. Slow rise and fall times with well overlapping clocks will reduce charge injection during transfer of the charge packages. We observed the best results with sinusoidal parallel clocks. Serial clocks, on the other hand, can be read out very fast, but this, together with small clock voltage wells, will negatively impact CTE. Thus, the shape of the serial clock overlap is even more important here and multi-level "trapezoidal" clocks will ensure smooth charge transfer between adjacent electrodes while preserving CTE.

3. **Non-inverted mode operation (NIMO):** The EMCCD 201-20 can be operated in inverted and non-inverted mode. For NIMO, the parallel clock low levels have to be less than -2 volts and the substrate should be at 0 volt. For optimized CIC performance, however, the difference between the SS voltage should not be greater than 4 volts with respect to the low levels of the parallel and serial clocks. This will keep the surface completely out of inversion, which has the negative effect that the surface dark current is no longer suppressed. See Section 3.1 below.

These requirements demonstrate how important it is to use a controller that will minimize CIC and also shape the clock waveforms for optimal CTE. We use a combination of all three methods to reduce CIC in our operation (Technologies, 2017c).

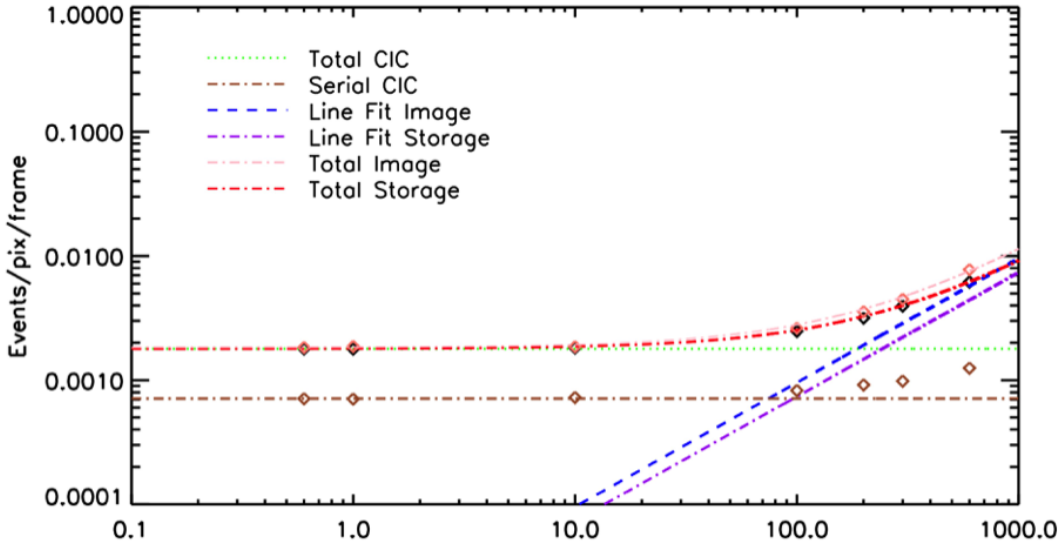

Figure 3.5: We plot the median of an image with increasing exposure time (x-axis) to evaluate our camera noise data. For short exposures CIC dominates, but for longer exposure times the dark current slowly increases and lifts off from the CIC bias noise level. We also take measurements of the serial register noise alone, so that we can subtract it from the total noise to regain the parallel CIC and identify where the CIC must be reduced further. The dark current noise for a 1000-second exposure is 0.001 events pixel^{-1} frame^{-1}.

Dark Current

Dark current is noise that stems from thermally generated charge in silicon devices (James R Janesick, 2001). It is unavoidable, but can be significantly reduced by cooling, which is why CCDs are preferable to CMOS devices. CCDs can be operated at cooler temperatures and perform with better quantum efficiency. Cooling the device too aggressively, however, will reduce electron mobility and degrade the CTE (Daigle, Gach, et al., 2008). The read-out speed of the serial register is very fast and it is difficult to efficiently move charge out of the EMCCD if the electrons are too sluggish. Initially, I had hoped to cool the device below -125°C for optimal dark current, but we have settled on -110°C for low enough dark current while maintaining good CTE. This sets a limit on the exposure time of up to 1000 seconds, since dark

current increases with time. But 1000 seconds is a long enough integration time to meet our science goal. Other considerations to reduce dark current are listed below.

Dark current reduction:

1. **Inverted mode operation (IMO):** This is the preferred mode to reduce dark current. The larger surface dark current is suppressed and only the bulk dark current remains. Surface dark current is the thermally generated charge at the silicon surface and bulk dark current is the thermally generated charge within the body of the silicon, which is a combination of diffusion current and depletion current. These two components have strong temperature dependences, and for operating temperatures below -30°C, depletion current is the main noise contributor (e2v).

2. **Non-inverted mode operation (NIMO):** At room temperature, the surface dark current is about two orders of magnitude larger than the bulk dark current. For a cooled device, however, the noise can be kept low for short integration times.

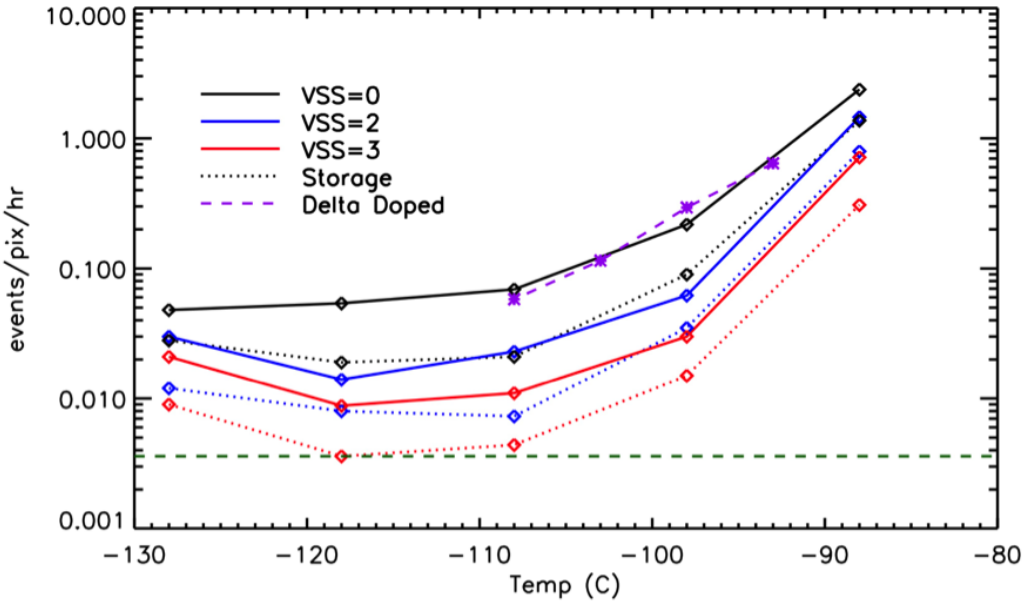

Figure 3.6: Dark current measurements at different substrate voltages. Results of the first delta-doped device are shown in purple at a substrate voltage of 0 V. Note the EMCCD is about 3° warmer than the diode that takes the temperature measurement at the back of the cold clamp.

After testing the camera in different configurations, we decided to operate the EMCCD in IMO. Unlike other applications, we will need a combination of fast read-out speeds with long exposure times, and this poses several challenges. We have to reduce CIC, which is very sensitive to speed, but we also need a very cold operating temperature to keep the dark current low for long exposure times. We can reduce CIC to a point, but the remaining difficulty comes from poor CTE at fast read-out speeds and low temperature, which makes charge transfer even less efficient. Thus, we have to find an optimal middle ground between low enough dark current, set by the EMCCD temperature, and CIC at that temperature. Ultimately, the long exposure time of about 1000 seconds is the key factor driving the search for the best operating conditions. By adjusting the multiplication gain, well depth, substrate voltage, and various parameters mentioned above, we are expecting to achieve a noise level of 0.002 events pixel^{-1} frame^{-1}. This sets a noise limit of about 0.001 events pixel^{-1} frame^{-1} for dark current as well as CIC (E. T. Hamden, Lingner, et al., 2015).

Camera testing with engineering grade devices have yielded a CIC just below 0.002 events pixel^{-1} frame^{-1} at -110°C and a dark current of about 0.01 events pixel^{-1} hour^{-1}. We are currently testing delta-doped devices to select a flight EMCCD with the best noise and imaging characteristics.

Cosmic Rays and Smearing

Cosmic rays that enter a CCD will saturate individual pixels, due to their high energy. They generate a lot of charge inside the silicon body of the CCD, which is captured by individual pixels. When the full well of a pixel is reached, excess electrons will spill into adjacent pixels, a phenomenon known as blooming (James R Janesick, 2001). Pixels that are nearly filled with charge can experience blooming as well. Even though cosmic rays can be identified and separated from the true signal by their signature—straight tracks across the pixel array—they are very undesirable. The pixels filled with charge can no longer capture the intended signal, and for EMCCDs the area with saturated pixels, or pixels with a high charge count, increases due to the multiplication process. Charge continues to be spread out horizontally into the trailing pixels in the multiplication register and the effect is naturally exacerbated for higher multiplication gain. The charge is no longer contained within one pixel, but spread across a tail that spans several pixels, and whatever charge was contained there previously has now been obscured and cannot be measured.

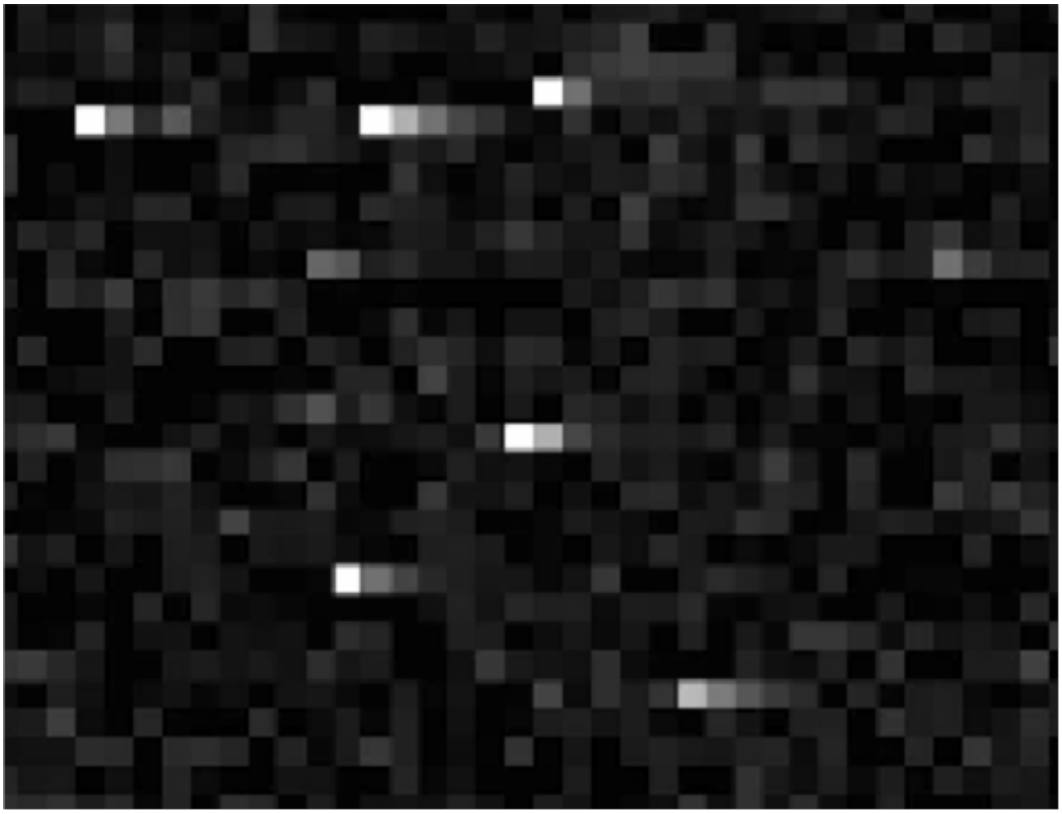

Figure 3.7: Raw image of a zero-second exposure at -110°C showing smearing of events. Bright pixels have several trailing pixels with counts significantly above the read-noise.

Good charge transfer efficiency is closely tied to the extent that the well of an individual pixel is filled. Cosmic rays are an unavoidable obstacle, but charge that was not transferred properly and amplified in the multiplication register also degrades image quality (Daigle, Djazovski, et al., 2012). This is particularly noticeable for fast read-out speeds and low potential wells at cold temperatures (below -95°C), which are our operation parameters. We observe a noticeable increase in smearing for increasingly lower temperatures and higher gain. When the charge that an individual pixel has collected is too large, it will not transfer properly and the leftover charge can smear across several pixels that follow in the multiplication register.

To account for the original charge in a given pixel, we implement a de-smearing process in our image processing. We correct for this horizontal tail of charge by performing statistical analysis on data taken from the same dataset with the same parameters (exposure time, pixel in question, as well as neighboring pixels). If the

charge count in the tail falls above the median plus a threshold for that region, the charge in that pixel is replaced with a random Gaussian read-noise level and the excess charge that was within these pixels is de-smeared back into the leading pixel. For a more thorough discussion please refer to E. T. Hamden, Lingner, et al. (2015).

3.2 New Readout Electronics: Nüvü Controller

EMCCD Controller Requirements

The noise inside an EMCCD camera can be greatly reduced with the right electronics. The precise shapes of the clocks will help with good charge transfer and specific combinations of clock and substrate voltages can reduce CIC while maintaining good CTE. It is also very important to achieve high electron multiplying (EM) gain while maintaining low CIC. This reduces the number of lost events, since higher gain will raise more events above the read noise floor (see the previous discussion in Section 3.1).

Figure 3.8: Lab camera continuity testing; PC board with CCD 201-20, connected to the read-out electronics from Nüvü.

Precise waveform shaping is critical to achieve low noise, especially at fast read-out speeds. The read-out speed of our device is 10 MHz, which leaves little time for the serial clocks to step up their voltages from the low to the high rail. It is therefore important for the electronics to put out precise voltages without overshooting or dragging. Giving the clock transfer multiple voltage levels to transfer the charge in the serial register will ensure that the charge in adjacent pixels is clocked smoothly

across the horizontal register, without leaving charge in the previous pixel. Charge in a given pixel can nearly fill the well toward the end of the gain register of the EMCCD, and it is difficult to completely transfer the charge packages from one pixel to the next, which adversely affects CTE (Murray et al., 2013). We chose a controller with these considerations in mind.

CCD Controller for Counting Photons (CCCP)

As we were looking for the right electronics to operate an EMCCD, a new controller by Nüvü Camēras came on the market. For his graduate work, Olivier Daigle had successfully designed and built specific read-out electronics for EMCCDs from e2v and was now making them commercially available. One of his first customers, we purchased the second-generation Nüvü CCD Controller for Counting Photons (CCCP) from his company. This controller generates electric signals that produce less spurious charge when electrons are moved around the EMCCD and it has an HV clock that can put out a high voltage at fast read-out speeds (10 MHz) without compromising other clock and bias signals in the system.

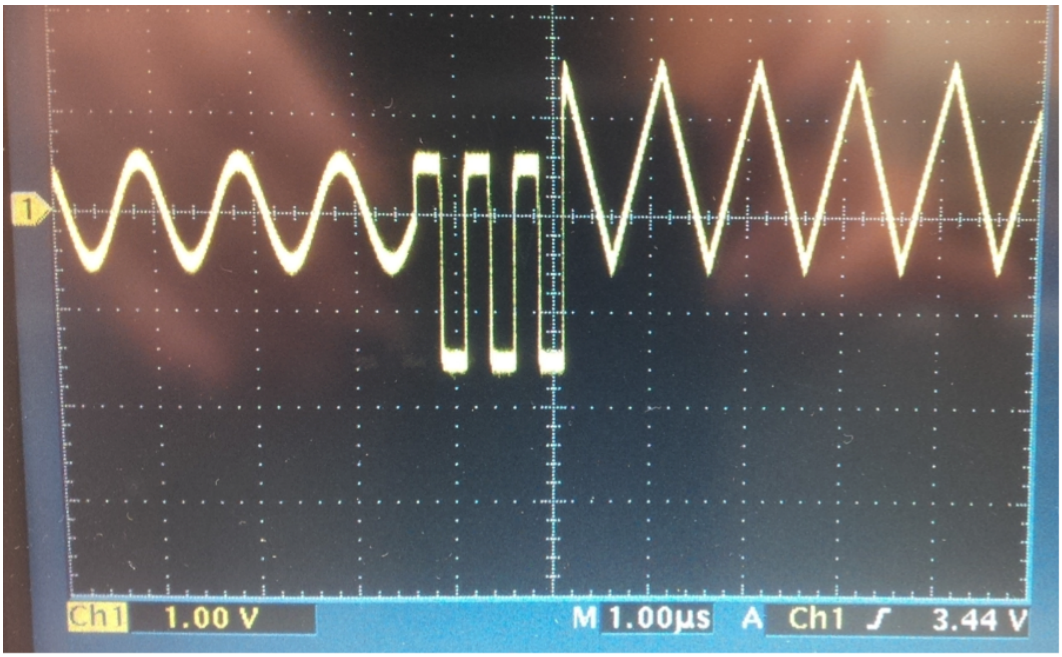

Figure 3.9: Sample of the sinusoidal, square, and triangular waveform output of the Nüvü controller.

The main difference between conventional controllers and this new device is the clock generation. The controller has four clock modes; it can produce square, triangular, sinusoidal and arbitrary clocks. Conventional controllers can generate

clocks with 2 to 3 steps in voltage levels, but CCCP drives the clock amplifiers with 12-bit DACs that refresh at a rate of 100 MHz. This yields 212 clock levels that can be sent to the clock inputs on the CCD every 10 ns. The controller generates 13 of these arbitrary clocks as well as the high voltage clock for the EM stage on the EMCCD, which is a sinusoidal resonant system with a switching precision of 1 ns.

The controller also provides 16 biases, some of which are used internally, and it incorporates an internal correlated double sampler (CDS) that is coupled to a 16-bit high-speed ADC. The digitized images are transmitted to a computer via a digital communication interface (Daigle, Gach, et al., 2008; Daigle, Carignan, et al., 2009). Initially, an EDT frame grabber card was used, but because it was incompatible with our system, I switched to a MATROX card. As MATROX discontinued updates for the Scientific Linux OS system, I incorporated a Gigabit-Pleora interface, and we have now switched to the Ubuntu operating system and are happily using the EDT card again.

The initial design of our flight instrument specified two EMCCD 97-00, to be operated side by side. We later discovered that we were unable to bring the devices close enough together and that we would lose valuable data at the central focus. By switching to the larger format EMCCD 201-20, we were able to recover the centrally focused data and still fit the desired wavelength range onto the chip. Because the controller was purchased before this adjustment, we had it customized to run two detectors, which is why it uses more power than the standard controller. We decided to purchase the third generation controller for the flight. Not only does it require less power, it is also radiation and vacuum safe and includes other hardware and software improvements. Notably, the waveform shaping can be done on a 5-ns scale, an improvement over the 10-ns steps of the version 2 controller.

With waveform shaping on a 10-ns scale and multi-level clocking, smooth charge transfer is ensured across the pixels in the serial register. By giving the clock transfer multiple levels to transfer the charge in the serial register, CCCP can achieve a CTE of 0.9999897 ± 0.0000008 (Daigle, Djazovski, et al., 2012).

CCCP Operation and EM Gain

To meet our science goals, it is important to attain sub-electron read-out noise, and the CCD controller for counting photons (CCCP controller) has that capability. This can be achieved without adding noise in form of the ENF factor, when the device is operated in binary mode, where binary here refers to signal in a pixel being counted

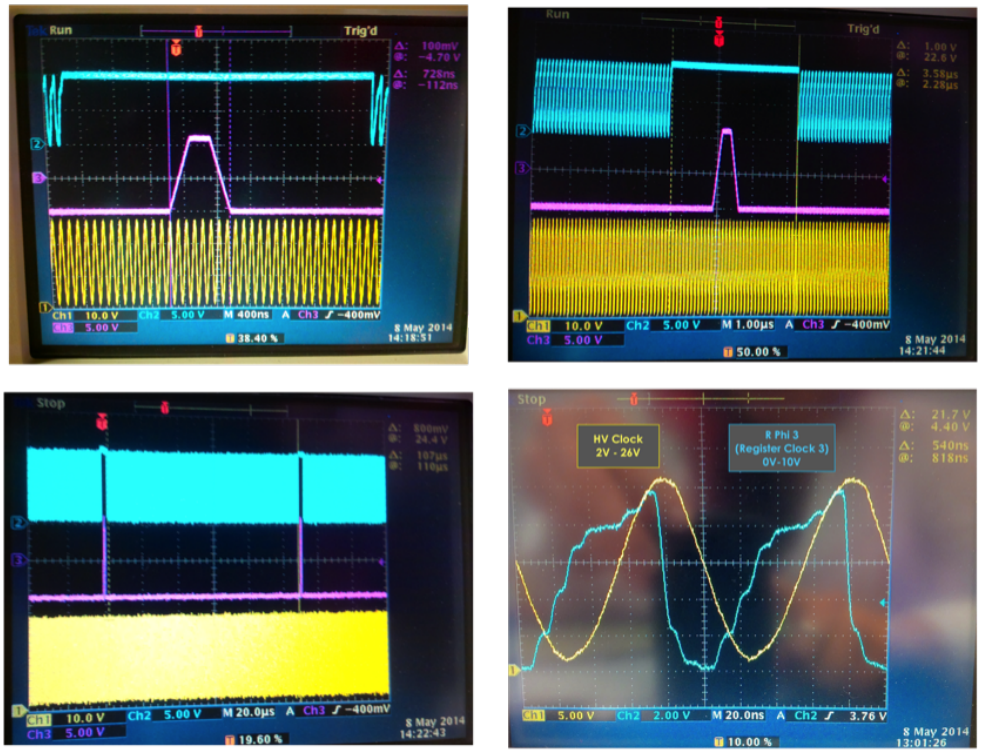

Figure 3.10: The blue trace is an arbitrary serial register clock, the magenta trace is a square image clock (with rise and fall time slopes for charge transfer) and the yellow trace is the sinusoidal HV clock. Here are four different zoom levels from one of my first read-out schemes. Top left: One row of the image area is clocked into the serial register via the magenta image clock(s), the blue serial clock(s) are stopped, so that the charge can accumulate in the pixels of the serial register. Top right: as soon as the charge has been moved from the bottom most row of the image area into the serial register, the blue clock(s) start moving the charge from the conventional read-out register into the multiplication register and out of the chip. Bottom left: After the last pixel with charge is transferred out of the device and output amplifier, the register clocks stop and the next row of charge is transferred from the image area into the serial register. These steps repeat, along with cleaning routines, until all rows of the EMCCD have been transferred and the next integration can begin. Bottom right: The HV clock replaces register clock 2 in the multiplication register and overlap with register clock 3 is shown. The voltage resolution of the image and register clocks is 12 bits.

as one photon per pixel per frame, if the signal is above a set threshold. If instead the signal is below that threshold, that pixel is assumed to contain no photon. To capture the incoming flux accurately, the signal strength of the observation must be considered when choosing the integration time. The signal amplification in the

multiplication gain register must be carefully set as well, to make sure that the charge is amplified enough to get above the threshold, but not so much that the signal will spill into adjacent pixels and generate horizontal CTE problems. The effective read-out noise (σ_{eff}) becomes the system's read-out noise (σ_{real}) from the output amplifier divided by the total multiplication gain (G) of the readout chain:

$$\sigma_{eff} = \frac{\sigma_{real}}{G}. \tag{3.2}$$

The total gain (G), called EM gain, is the result of the cumulative gain (g) from the 604 elements in the multiplication chain. The voltage drop set by the HV clock determines this value. Even though the probability that one of these stages will amplify the charge is low, the overall gain is high. We set the high voltage clock to about 40 V, which averages an EM gain of about 1000 e/e, while the read noise lies somewhere around $60e^{-}$, depending on the particular chip. Figure 3.4 shows the histogram of a photon counting device with a RN of around $60e^{-}$.

The detection probability of a photoelectron is set by the gain (G) divided by the system's read-out noise (σ_{real}). The detection threshold is generally set to 5.5σ, and a higher gain will result in a higher probability of amplifying a photoelectron above the threshold. However, the gain cannot simply be set as high as possible because CIC is proportional to gain and a compromise has to be found between photon detection, horizontal well depth, and background noise (Daigle, Carignan, et al., 2009).

Future Flight Electronics

We have purchased the third generation CCCP controller from Nüvü Camēras, which we will use for the second FIREBall-2 flight. This controller will need a new interface board, because it comes with different connectors and it also has a slightly different clock and bias output. We had this controller customized to test a slow read-out option and its effect on CIC. The HV clock is generated by a finely tuned resonant system inside the controller. It can resonate at two different frequencies, but they cannot be more than a factor of two apart. We wanted to test CIC with a read-out frequency of about 1 MHz, which is why we needed a second resonant system. Our new controller comes with two HV clock outputs and I made a PC board that allows us to choose between the two clocks during the testing phase. We are currently bringing the controller online and should have results within the next few months.

3.3 New Cooling Mechanism: Sunpower CryoTel MT

Cryocooler Criteria

Because CIC only happens once during read-out, it will have the same value regardless of integration time. Dark current, on the other hand, increases with integration time (due to thermal electron motion) and can therefore be greatly reduced if we slow down the electrons. Thus cooling the device to very low temperatures is crucial for dark current noise performance. At cryogenic temperatures, water and a few other molecules in the air will freeze onto the face of the sensor, which is why the sensor must be operated in a vacuum. To maintain high vacuum during the balloon flight, I incorporate a charcoal getter, which must be cooled as well to trap small molecules by adsorption. This addition increases the heat load that must be removed from the camera system. The vacuum enclosure, getter, and cooling design will be discussed in Chapter 4.

In the lab, low temperatures can be easily achieved with liquid nitrogen, but for a balloon flight other forms of cooling are preferred. Nitrogen can be utilized, but because air pressure is lower at higher altitudes, it will have to be pressurized and used in its solid state. Another option is a Pulse Tube Cooler, but their performance depends on orientation (convective effects, Ross and Johnson (2004)), and they are not a good match for this experiment because the cooler will be mounted to the tilting vacuum tank. After looking into various mechanical cooling systems, I found that a mechanical cryocooler from Sunpower met all of my criteria. There are multiple requirements for the cooler, which are mostly interconnected and linked to weight limitations on the balloon:

1. **Low temperature of the coldtip:** The EMCCD has to be operated at cryogenic temperatures to reduce dark current, and the charcoal getter performs best below -160°C.

2. **Low input power (efficiency):** The power on our balloon comes from heavy batteries, and the payload power budget must therefore remain as low as possible (see Section 2.1). The vibrational noise of the cooler will also increase with higher input power.

3. **High heat lift:** An important property of cryocoolers is the actual heat lift—that is, the heat load the cooler can remove while maintaining a set temperature—that it can achieve.

4. **Low vibration & active vibration cancellation electronics:** Vibrational noise from a cryocooler will degrade the point source imaging accuracy. The fundamental drive frequency of 60 Hz is generally tuned passively, but higher harmonics need active cancellation systems.

5. **Overall size and weight:** The instrument will be operated at different elevation angles and it rotates, leaving little available space for any hardware that has to be mounted externally to the vacuum tank. Also, if the cooler is too massive, heavy counterweights will have to be added to the tank for stability and smooth rotation, which increases the overall weight of the instrument.

6. **Vacuum safe components:** The instrument will fly above an altitude of 37 km, where the atmospheric pressure is about 3 Torr. If the electrical components are not vacuum safe, the device will have to be flown in a pressure box. Depending on the size of the cooler electronics, this can also add significant weight and volume.

The CryoTel coolers from Sunpower met these criteria and I decided to purchase the smallest version for lab testing. The CryoTel MT can be operated from 30-W to 80-W cooling input power and provides a heat lift up to 14 W.

Sunpower CryoTel MT Overview

The CryoTel family of coolers are Stirling cryocoolers (Sunpower, 2017c), which convert electricity into moving heat from the coldtip to the heat rejector. The cryocooler's coldfinger is a hollow cylinder, which is filled with helium and houses the piston that drives the heat engine. Sunpower's design features a frictionless gas bearing, which provides a few key benefits over other cryocoolers currently on the market (Lane, 2005). These engines are mechanically simpler than kinematic machines and less expensive to build. CryoTel coolers have a long lifetime (more than 10 years) and need less maintenance, because critical components are not touching and these free-piston machines do not need mechanical bearings, lubrication, or seals like kinematic Stirling engines. The compact linear alternators have a simple mechanical configuration, which keeps the design small, yet more efficient (up to 40% thermal efficiency) and allows for greater output power per mass. The power output can be easily adjusted, which is difficult to do with kinematic designs, and the electronic controller can modulate the free piston machines for load-matching capability.

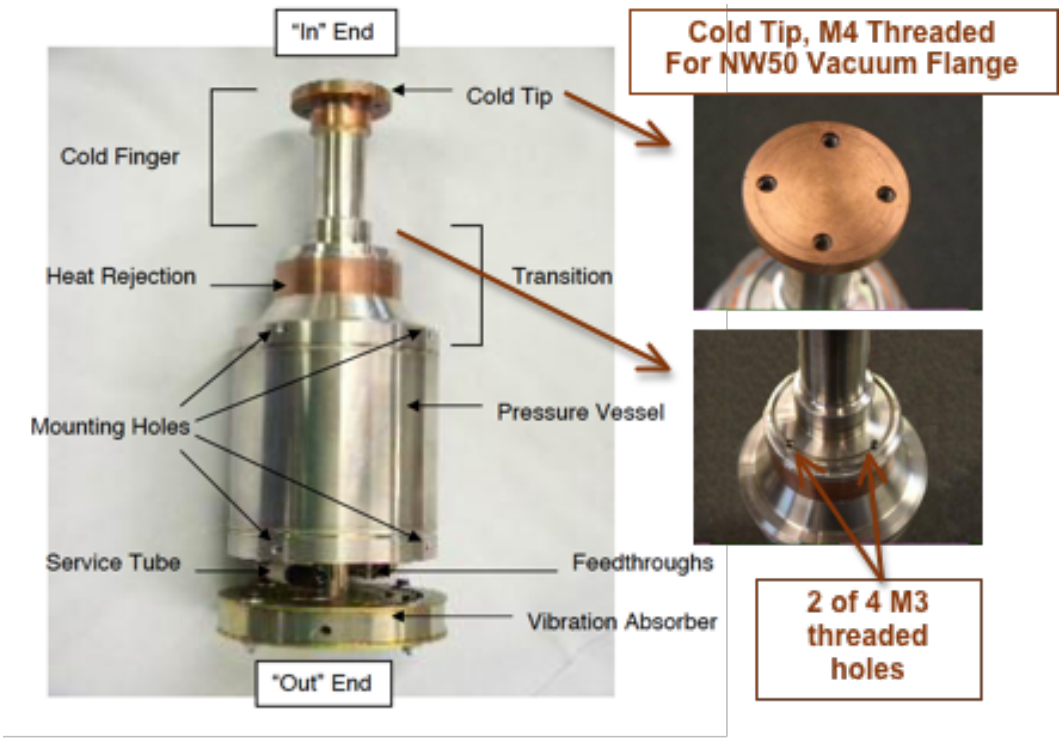

Figure 3.11: Details of the Sunpower CryoTel MT cryocooler. (Sunpower, 2017a)

The MT cooler features the coldtip at the cold end of the Stirling engine and the heat rejector at the hot end. The heat rejector resides just outside the vacuum vessel where heat has to be removed efficiently to prevent any damage to the cooler and maintain good performance (see Figure 3.12). The displacer and piston that move the hot and cold gas between the hot and cold ends inside the cylinder are sized precisely to fit. The piston has a very small clearance to the walls of the cylinder and a higher coefficient of thermal expansion, so if the heat rejector exceeds 80°C it will expand faster than the cylinder, resulting in a seize. This is the only way to break this otherwise very robust cryocooler.

The green curve in Figure 3.12 corresponds to a cooling input power of 30 watts (which requires around 48-W of electrical input power) and the orange curve shows the cooler performance for an input cooling power of 80-W (which requires around 126-W of electrical input power). The curves to the left of the graph indicate efficiency, while the more linear curves illustrate lift, or the heat load it can overcome. I operated the cooler at constant input power in the lab, rather than temperature controlling it, to determine the power budget for the cooler during the course of the balloon flight. After some testing in the lab, I found that optimal performance

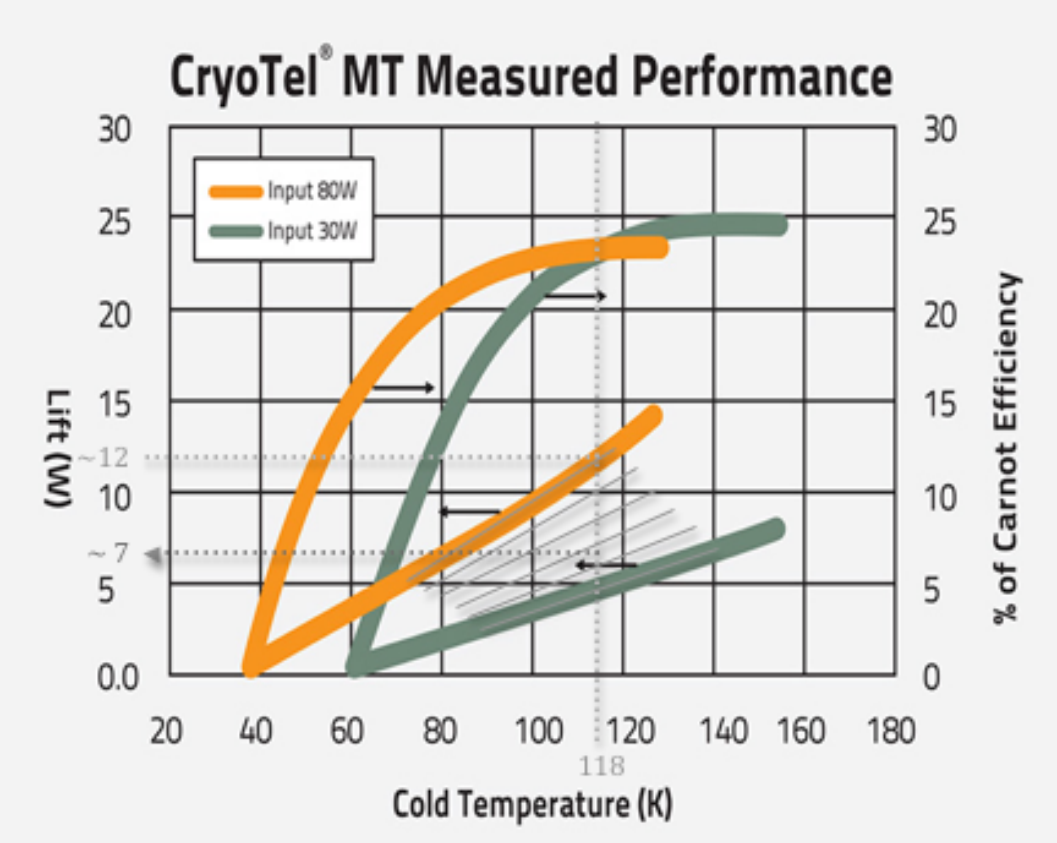

Figure 3.12: Sunpower Cryocooler Performance graph, for their smallest cooler. The graph shows the lowest (green) and highest (orange) input power (cooling power, not the electrical input power). The grey curves are an extrapolation in steps of 10 watts. (Sunpower, 2017b)

is achieved by using a constant input power of 45 watts, which leaves the coldtip temperature at -160°C, and the cooler achieves a lift of about 7 watts.

Cryocooler Hardware

The cooler weighs 2.1 kg and is about 24 cm long with a body diameter of 7.3 cm. It comes with a controller that has a serial (RS-232) connection and the unit can be operated from a computer with a simple hyperterminal interface. When the cooler was first purchased, the active damping electronics had not yet been developed and the controller was a simple electronics card. I provided the housing and wired everything myself, but the new Gen 2 controller comes with a finished electronics box. We later purchased active damping electronics with much more functionality, and they are almost twice as large in size. We also exchanged the passive damper at the back of the cooler with the slightly longer active damper, and

this new system required some rewiring. The cooler is supplied with a resistance temperature detector (RTD) sensor for the cryocooler tip, which feeds temperature information back to the controller. Depending on the users' needs, this mechanical cooler works in constant input cooling power mode, which can be adjusted in steps of 1 watt from 30 watts to 80 watts. Alternatively, it can be operated in temperature control mode for gradual lowering of the coldtip temperature to a set temperature. The video cable is a simple two-prong Mate-N-Lock connector and it is easily adapted to work with any power source. A 24-V power supply is required to run the cooler and controller, but the power supply cannot have a progressive current regulation, or current limiter, because the cooler will draw a lot of power at start-up and this spike in current has to be readily supplied; otherwise the cooler will shut down and leave the piston in the wrong position, as I discovered independently. After start-up, the power regulates itself, depending on the cooling scheme employed.

Figure 3.13: Sunpower Cryotel MT hardware.

Cryocooler Operation

The Sunpower cryocooler is a free-piston Stirling engine. The current coils that drive the piston are in the large-diameter cylinder of the cooler and the helium gas with the piston and displacer are inside the small-diameter cylinder; these are the cold

and hot ends of the heat engine. When the heat rejector does not sufficiently get rid of the heat that it generates and warms up to 80°C, the piston housing will expand too much and the piston will jam inside, subsequently breaking the mechanics. Sunpower suggests three cooling solutions: water jackets, copper fins, or a simple heat exchanger that will conductively transfer the heat. They can provide the fitted water cooling jacket or copper fins as removable or permanent attachments.

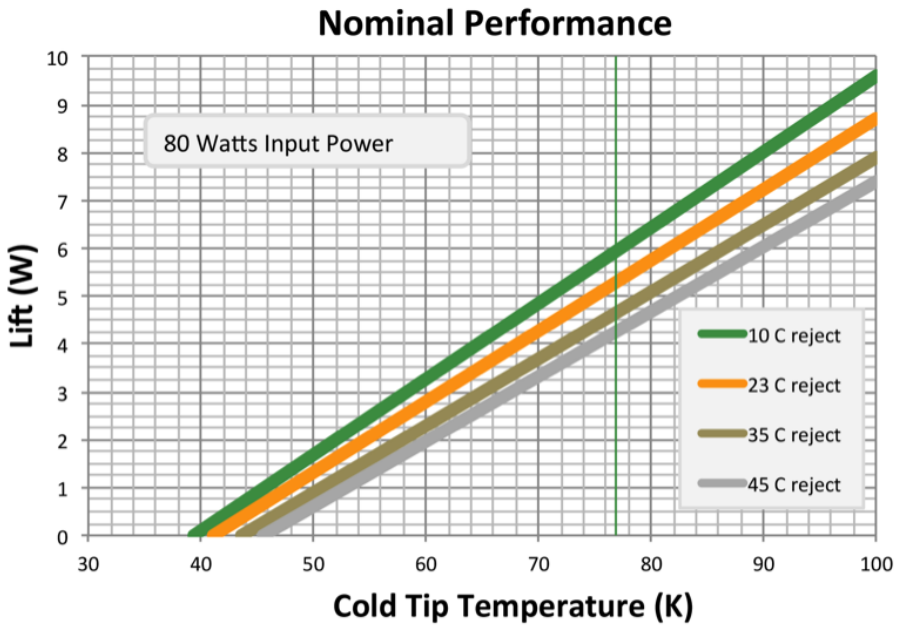

Figure 3.14: The reject temperature (ambient temperature) will determine how much lift the cooler can achieve for a given coldtip temperature. (Sunpower, 2017b)

I decided to attach copper fins to the heat rejection collar for lab operation. The unit comes with threaded holes for the fins and they can be easily detached (see Figure 3.13). I also built a small wind tunnel to pull cold air from the "In" End of the cooler through the fins (see Figures 3.11 and 3.15). By forcing cold air through the fins I can keep the reject temperature at around 30°C. This will not work for the flight, and we will conductively remove heat with copper blades to keep the reject temperature at about 45°C. The copper blades will also support the cooler while reducing the translational vibration from the cooler.

The rejection collar (hot end) is just outside of the vacuum enclosure and the coldfinger and tip (cool end) are on the inside. The cryocooler is mounted to the vacuum shroud with an NW50 flange (see Figure 3.15) that Sunpower welded to the cryocooler at my request.

The distance from the heat rejection collar to the coldtip is about 10 cm, so the cryocooler has to be attached close to the dewar wall (see Figures 3.15 and 4.1). I built a test camera system in the lab that could use both cooling solutions, the cryocooler and liquid nitrogen as a backup. The vacuum chamber utilized a charcoal (carbon) getter that needs to be cooled even more than the EMCCD and I incorporated a PC board for the electronic interface. All of these decisions put additional constraints on the lab camera design and drive the geometry of the lab camera system.

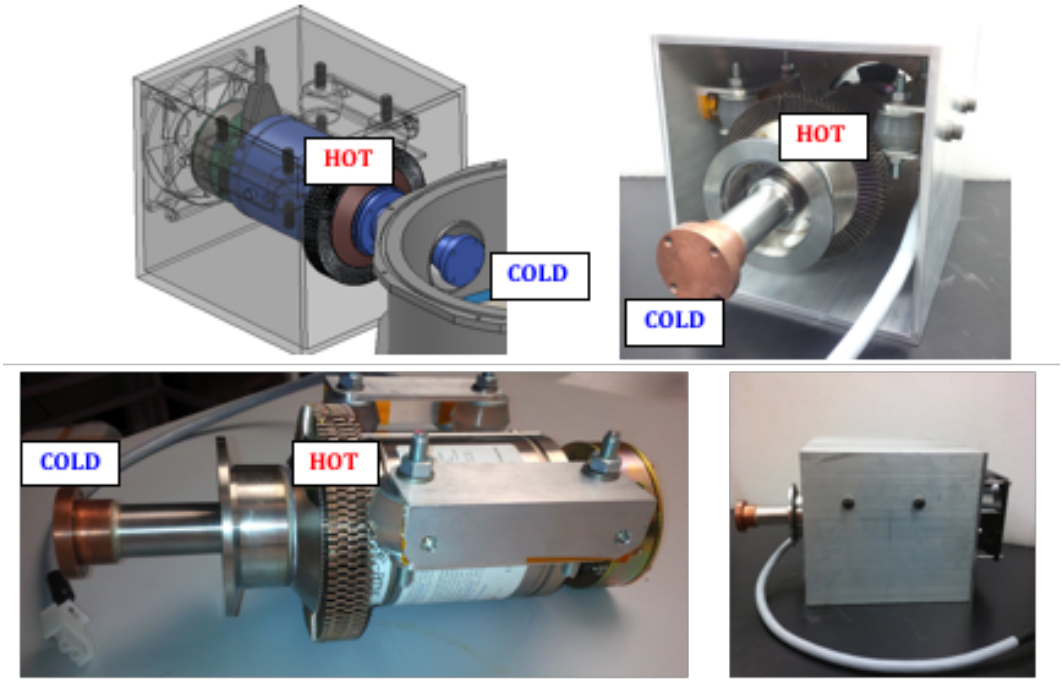

Figure 3.15: Mounting scheme of the cooler. Top left: the SolidWorks model for the wind tunnel with silicone vibration damping sandwich mounts that hold the cooler in place and fan that pulls the air from the cold "In" end of the cooler. Bottom left: mounting brackets attached to the sandwich mounts and cooling fins just outside the NW50 flange. Bottom Right: Aluminum wind tunnel with mounted cryocooler and fan.

Chapter 4

CAMERA DESIGN & IMPLEMENTATION

4.1 Mechanical Design, Construction, and Integration

Dewar

The starting point for my camera design is the e2v EMCCD, which requires cryogenic cooling and hence a vacuum chamber. From my group, I inherited an Infrared Laboratories ND-8 nitrogen dewar (Laboratories, 2017). It has a single cryogenic vessel with a cold work surface (cold plate) and it is surrounded by a vapor-cooled radiation shield. The dewar was not readily usable; it was losing heat much faster than it should and required too much nitrogen. I dissembled the dewar and found that the metal posts holding the vessel at the center of the dewar with G10 plates and keeping it thermally isolated had been slightly bent, such that one side of the vessel was touching the warm outside walls of the dewar. This created a direct conductive link to room temperature and was responsible for the high rate of nitrogen loss. I was able to adjust the posts without causing any damage and put everything back together successfully. The dewar was then fully functional and I proceeded to bake it out to quickly remove the bulk of contaminants.

The existing hardware for the vacuum chamber, which is located above the cold plate of the dewar, was sufficient to begin some initial vacuum chamber testing. Keeping the dewar under vacuum and continually cooling and heating it also helped to outgas the remaining molecules that were trapped on the surfaces inside the dewar.

I was able to do this with liquid nitrogen, while evaluating my cooling options for the flight camera and finalizing the search for the right mechanical cooler. I decided to design the cold-transfer link for my EMCCD in a way that would allow me to continue using this cooling mechanism as a backup, in case the mechanical cryocooler turned out to be unusable, too noisy, or not powerful enough. I also wanted the old heavy gold dewar to add some weight to the system, since the balloon vessel is very heavy and absorbs a lot of the vibration from the cryocooler. While I was outgassing the dewar in the lab I had to come up with a new vacuum chamber design to integrate the EMCCD with the Nüvü controller and Sunpower cryocooler.

Vacuum Chamber Hardware

Designing the new hardware required CAD software, and I learned how to program in SolidWorks, which is the standard design software for mechanical engineering. I was able to import existing models of the cryocooler and some connectors, and I made models of the ND-8 dewar and Nüvü controller myself. Shown in Figure 4.1, the vacuum chamber design had to accommodate the EMCCD at the center of the dewar, and the plan was to hold the sensor in place with a large diameter PCB. This would require two vacuum chamber shrouds with the circular PCB wedged between the two parts. The middle shroud should directly attach to the gold dewar and have the cryocooler port and one hermetically sealed connector for electronics. The top cover would contain the light entrance window and some hole patterns for optics attachments. The PCB should be pressed against both parts with O-rings to achieve a tight hermetic seal. I also wanted to place the EMCCD at the correct distance to the window for an existing CS-mount lens. The height between the window in the top cover and the face of the EMCCD should therefore be the flange focal distance of 12.50 mm, so that all optics can be attached directly to the dewar window to easily focus an image onto the EMCCD. Because the main focus of my camera testing takes place in the dark, I needed to take into account a baffle in front of the window. I machined an O-ring groove onto the face of the flange for a soft O-ring to eliminate any window light leak.

It is important to minimize the O-ring cross sections and apertures in a vacuum chamber. The diffusivity of the O-ring material will degrade the low pressure inside the chamber over time, but more importantly any aperture has the risk of creating a vacuum leak (Seals, 2001). These leaks are difficult to find and fix. Therefore, I kept the dewar ports to a minimum and used only one feed-through for all electronics. The electronics feed-through would support all cables for monitoring and controlling the EMCCD temperature within the dewar. All EMCCD signals, however, are wired through the center layer of the PCB and are completely separated from each other and other signals. The second port in the middle shroud is for the cryocooler attachment, which is just a standard NW50 flange.

Vacuum Chamber Mechanical Design

The design of the middle shroud matches the geometry of the gold dewar. I kept the octogonal shroud-dewar interface and reduced the height of the shroud to fit the NW50 flange, without interfering with the PCB that sits on top of this shroud. The aluminum middle shroud is machined out of one piece of aluminum, and thus

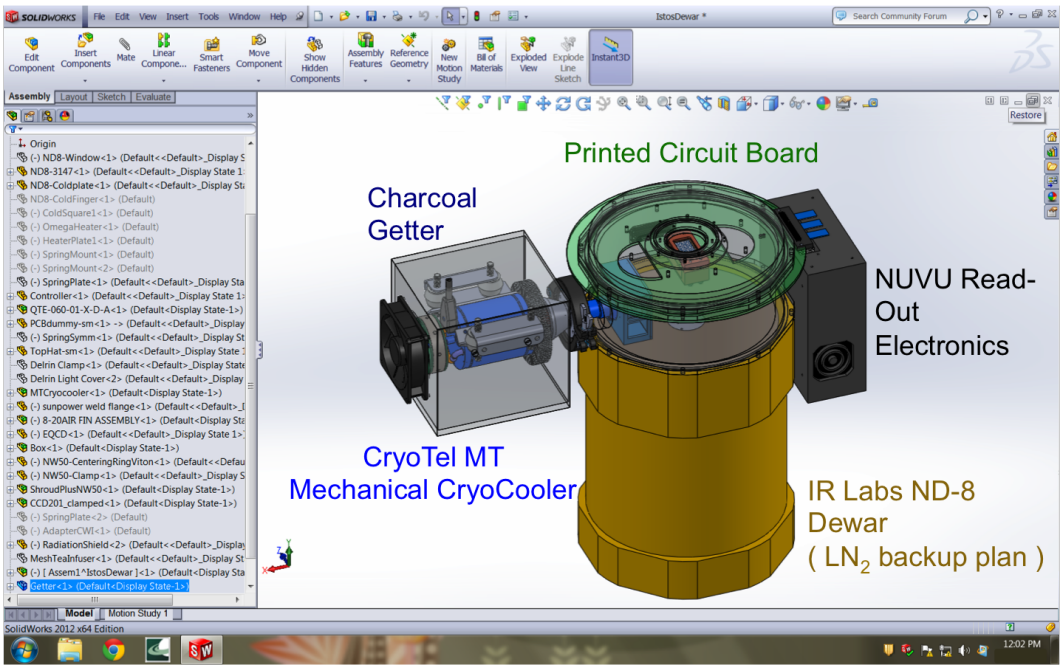

Figure 4.1: A look at the mechanical design software SolidWorks that I used to design the Lab & Flight Camera. This is the final result for my lab setup. The hole pattern on the top of the shroud accommodates the testing of the EMCCD at Palomar with CWI.

the taller the shroud, the more expensive it is to make. Another reason to reduce the shroud height is to accommodate the coldfinger alternative that will attach to the dewar cold plate, because its height is proportional to the height of the vacuum shroud. The second port on the shroud is for the heater electronics and I decided to use a D-subminiature (DB-25) connector. The DB-25 hermetic connector has an O-ring seal on the back side and is tightly screwed into the middle shroud. I also potted the seal with a NUSIL compound, to further minimize any vacuum leaks. I prepared this sealant in a vacuum oven to prevent any outgassing over time and to eliminate any air bubbles inside that could be slowly sucked into the high vacuum. I also fastened the connector with vented screws, again to avoid outgassing air pockets, because this will degrade the vacuum pressure and mimics a vacuum leak, which can be difficult to diagnose.

The electrical connector feeds power into the vacuum to drive the heater that regulates the EMCCD temperature, which is why the connector on the shroud is a male (pins) DB-25 connector and the mating connector on the harness is female (sockets). The connector is also wired for four temperature diodes that monitor the temperature inside the vacuum chamber at different locations along the coldlink. These monitor

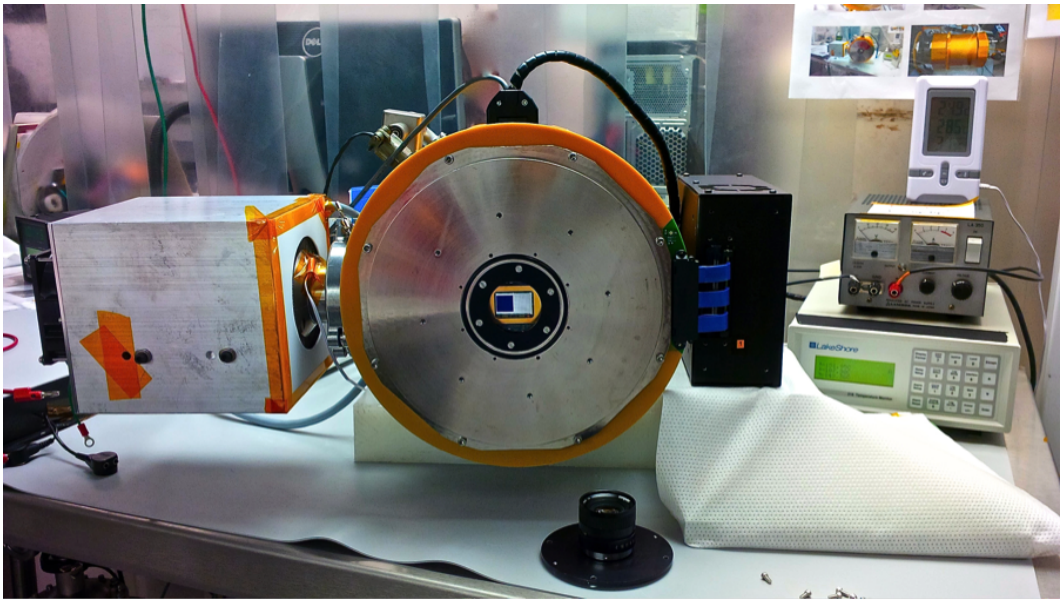

Figure 4.2: Front view of the test camera system. From left to right: cryocooler inside the wind tunnel with temperature diode on the heat rejection collar visible. Dewar face with the EMCCD and gold clamp behind the window. Nüvü Controller attached to the PCB with a Samtec EQCD high-speed coax cable assembly. Lakeshore temperature controller to monitor the coldlink chain inside the dewar from cryocooler coldtip to EMCCD; including a feedback loop connected to the power supply for the heater to regulate the sensor's temperature. CS-mounting plate and optics in front of the dewar.

and record the temperature flow, allowing me to verify the actual heat load while also keeping an eye on the temperature gradient. Because there is no way to view the inside of the chamber once closed, it is important to ensure that no surfaces inside the vacuum tank come into contact with each other.

The diode at the back of the EMCCD is used for the heater feed-back loop. The lab setup has only one 10-W heater at the end of the cold clamp, just behind the EMCCD, but the flight system will have two such heaters for redundancy. When I first started testing the lab camera setup, I used a dummy CCD with an identical Alumina package (as the real EMCCD) and recorded the temperature directly on the face of the CCD package. The diode on the back of the sensor is 3 degrees colder than the face of the sensor, which is why the control temperature has to be set 3 degrees lower than the target temperature. The EMCCD heater (which can be seen in Figures 4.3 and 4.8) is connected to a power source, which is connected to a simple relay system on the Lakeshore temperature controller. The system's alarm setting controls the power output. Two alarms can be set up per input diode: a

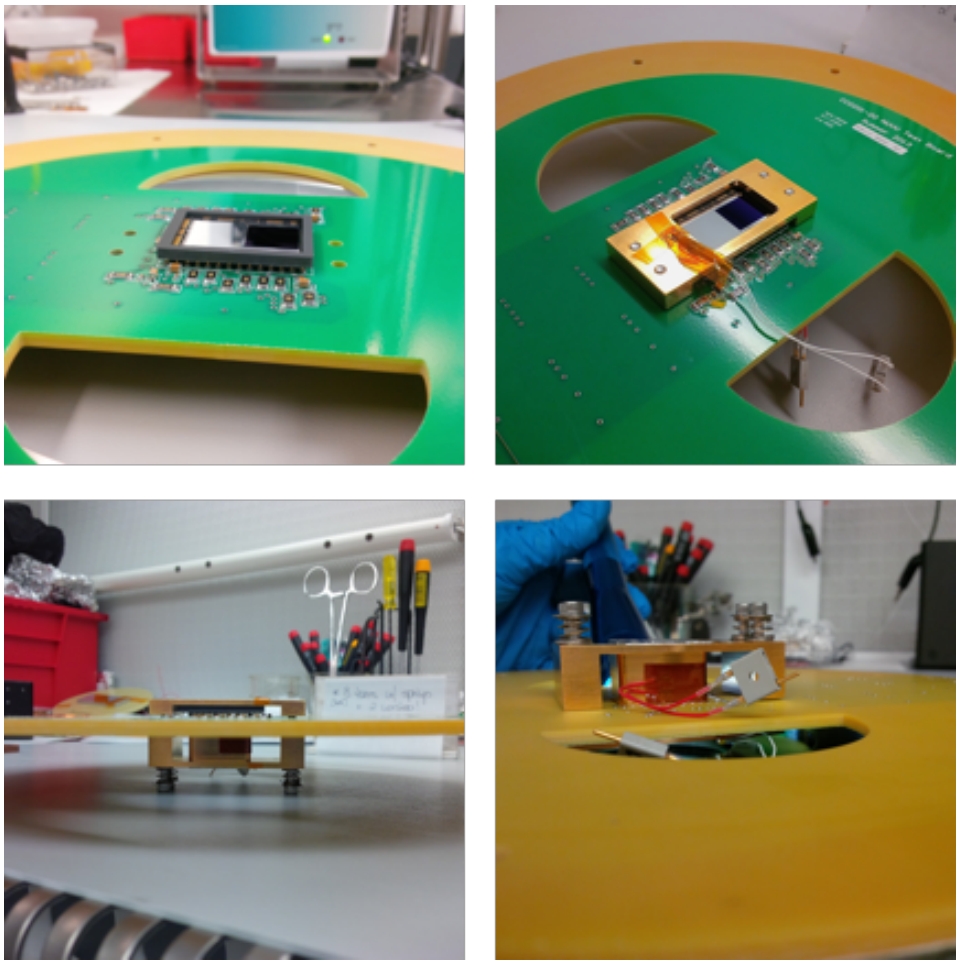

Figure 4.3: Close up look of the gilded copper clamp that holds the EMCCD in place. The cold clamp does not touch the PCB, only the backside and the front frame of the sensor's Alumina package. The springs on the four screws that hold the clamp together allow for controlled repeatable mechanical pressure.

"high value" and a "low value." In my case only the "low value" alarm is necessary, because I cool the EMCCD from room temperature down to about -115°C and then use the heater to warm the sensor back up to reach my target temperature of -110°C, which corresponds to a "low value" of -113°C. The low alarm is activated if the temperature reading falls below the "low value." The alarm is deactivated when the reading rises above the "low value" plus the dead band, for which I have chosen a value of 0.5 degrees. This keeps the temperature stable and minimizes temperature swings around the target temperature. I set the power source outputs at 24 V with a current of about 0.2 A for a gentle warm up, so the EMCCD temperature does not overshoot the target temperature by a large margin.

The copper spring clamp that holds the EMCCD in place also doubles as the final heat link in the cold chain from the coldtip of the cryocooler. The copper conductively removes heat from the sensor's package on the back side as well as from the front frame surrounding the silicon chip. This keeps the temperature of the imaging surface uniform. All copper parts inside the vacuum are used as part of the cold chain and are gold plated to increase their emissivity as well as make them slightly more malleable for better contact. The conducting surfaces are tightly pressed together and the conductive cross section is further increased with indium.

My goal of achieving a high vacuum of 10^{-6} Torr requires tight seals at all hardware interfaces as well as a clean vacuum environment. I use Viton O-rings (Fluorocarbons), because they allow for adequate compression and are capable of making a reliable vacuum-tight seal. Nitrile (Buna-N) O-rings have an even lower permeability, but they have a higher outgassing rate than Viton (Chief Engineer, 2010). Increasing the compression of the O-ring can reduce the permeation rate. The compression set—the tendency of the material to stay compressed after removal from vacuum—is very small but not negligible, and I renew the O-rings about once per year, depending on how much time they spend under vacuum.

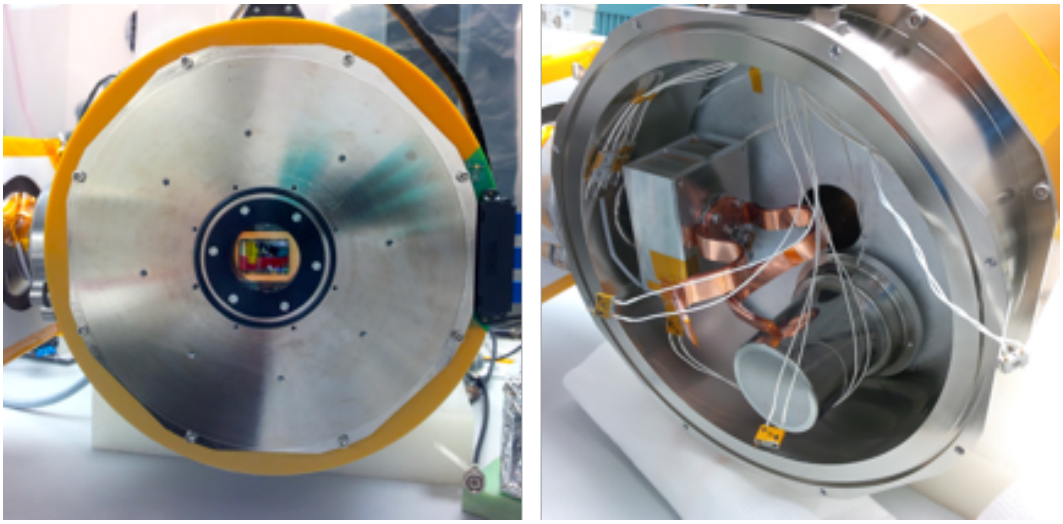

Figure 4.4: Front view of the test camera system with different hole patterns for various attachments. Left: The top cover with UV window, part of the gold clamp, and EMCCD. Right: View of the inside of the middle shroud with charcoal getter, cold ribbon, and Zeolite getter on top of the radiation shield of the dewar.

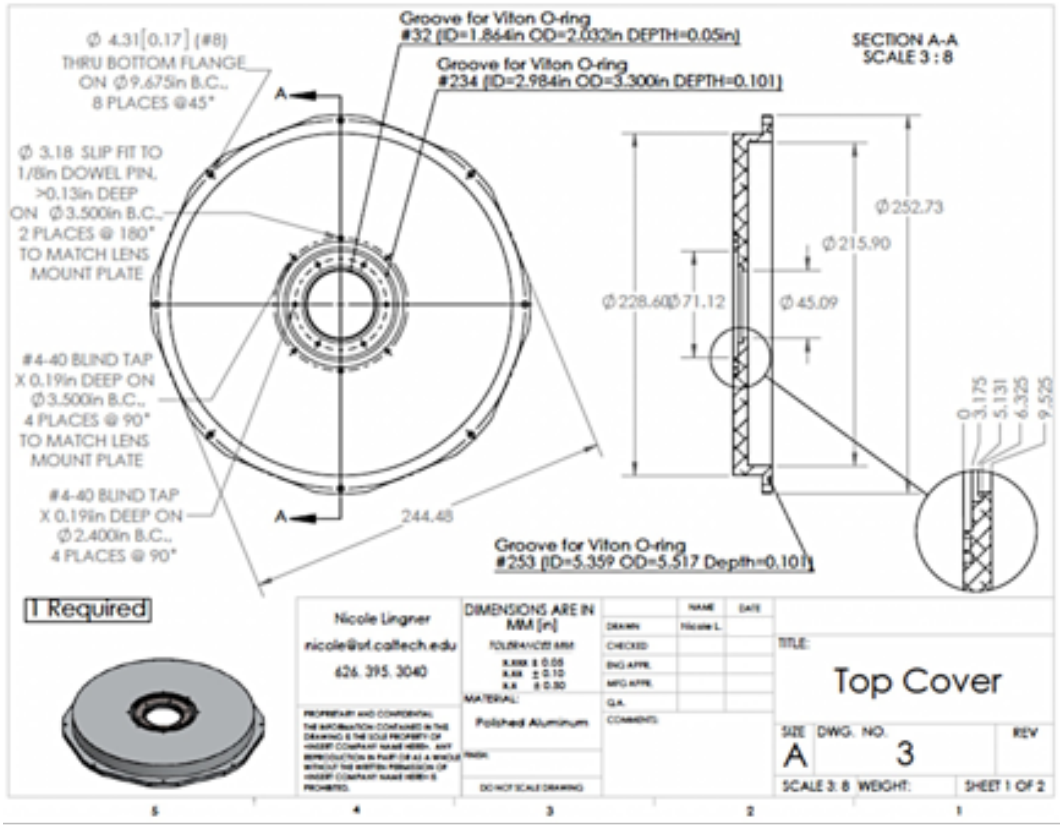

Figure 4.5: SolidWorks drawing of the vacuum chamber's front cover.

Vacuum Chamber Low Pressure Considerations

The lab setup consists of a turbo pump, connected to the dewar, that can run uninterrupted, even during testing. The vacuum pressure can be reduced to almost 10^{-5} Torr with the roughing pump, and a sorption pump can bring the vacuum down to 10^{-7} Torr. A sorption pump is a vacuum pump that creates a vacuum by adsorbing molecules on a very porous material like activated charcoal (carbon) and zeolite. The charcoal getter that came with the dewar did not have enough pumping power, but I was able to make a substantial vacuum pressure improvement by designing and implementing a new getter. I removed the existing 20mL charcoal getter and replaced it with a 120mL charcoal getter, which I designed in a way to keep the charcoal inside uniformly cold. Charcoal has to be surrounded by cold surfaces, because it primarily cools by radiation (Kindall, 1963), which is why the new design has mostly solid walls and four equally-spaced metal cylinders inside the volume. The colder the granular activated charcoal or carbon, the more contaminants it can absorb, and the pumping speed increases (Kindall, 1963). The new getter makes better cold contact with the coldtip of the cryocooler, which is very important

to reduce heat loss and hence power, but also to quickly reach a lower charcoal temperature.

Activated charcoal is carbon, in this case derived from coconut husk, that has been treated to create microporosity—many tiny small volume pores that increase the surface area of the carbon for adsorption. The size of these small pores is not uniform, allowing activated charcoal to adsorb many different molecules and remove various contaminants from the vacuum. Water is the main molecule that must be removed form the vacuum chamber, but the smaller molecules like CO_2, nitrogen, oxygen, argon, and hydrocarbons must be taken out to achieve a high vacuum.

To help remove water from the dewar, I also added a zeolite getter as a desiccant (it is not actively cooled to be used as another sorption pump). If the desiccant is relatively cool and dry, its surface vapor pressure remain low enough to attract moisture from the air, which has a high vapor pressure when it is damp. After the desiccant becomes wet and hot, its surface vapor pressure rises and it will give off water vapor to the surrounding air. Zeolite, in contrast to activated charcoal, is a molecular sieve, which means that it can only adsorb molecules of a specific size. There are a few different types of Zeolite (Weitkamp, 1989); I use Type 5A, which adsorbs molecules with a diameter less than 5 Å. This works well for water molecules, which are a little less than 3 Å in diameter, but it can also adsorb some hydrocarbons and other molecules. Zeolite never loses its ability to adsorb water, as long as it can be baked off after saturation. For the flight, we will have several baked and vacuum-sealed pouches of zeolite that can be changed out, replacing saturated zeolite in the vacuum tank.

With this getter system, I have accomplished a very low-pressure environment for my lab camera. The cooled charcoal captures any remaining outgassing and the Zeolite adsorbs water vapor that escapes from the charcoal during warm-up, which makes the vacuum chamber safer for the EMCCD. The flight vacuum tank will be more difficult to control because it houses multiple surfaces that can continue to outgas for a long time. We will have to start pumping on the flight chamber long before we start any camera testing, to remove as much of the initial vacuum contaminants as possible. The total coldlink inside the vacuum tank is much longer and wider as well, which will require the cooler to be run at higher input power. After initial testing, we decided to purchase the larger version of the cryocooler for the flight instrument, which provides more cooling and thus pumping power to keep the larger volume at acceptable vacuum levels during the flight. We will have to find

a balance between removing the heat generated by running the cooler at high input power and the required cooling power necessary to sufficiently cool the activated charcoal to achieve high vacuum.

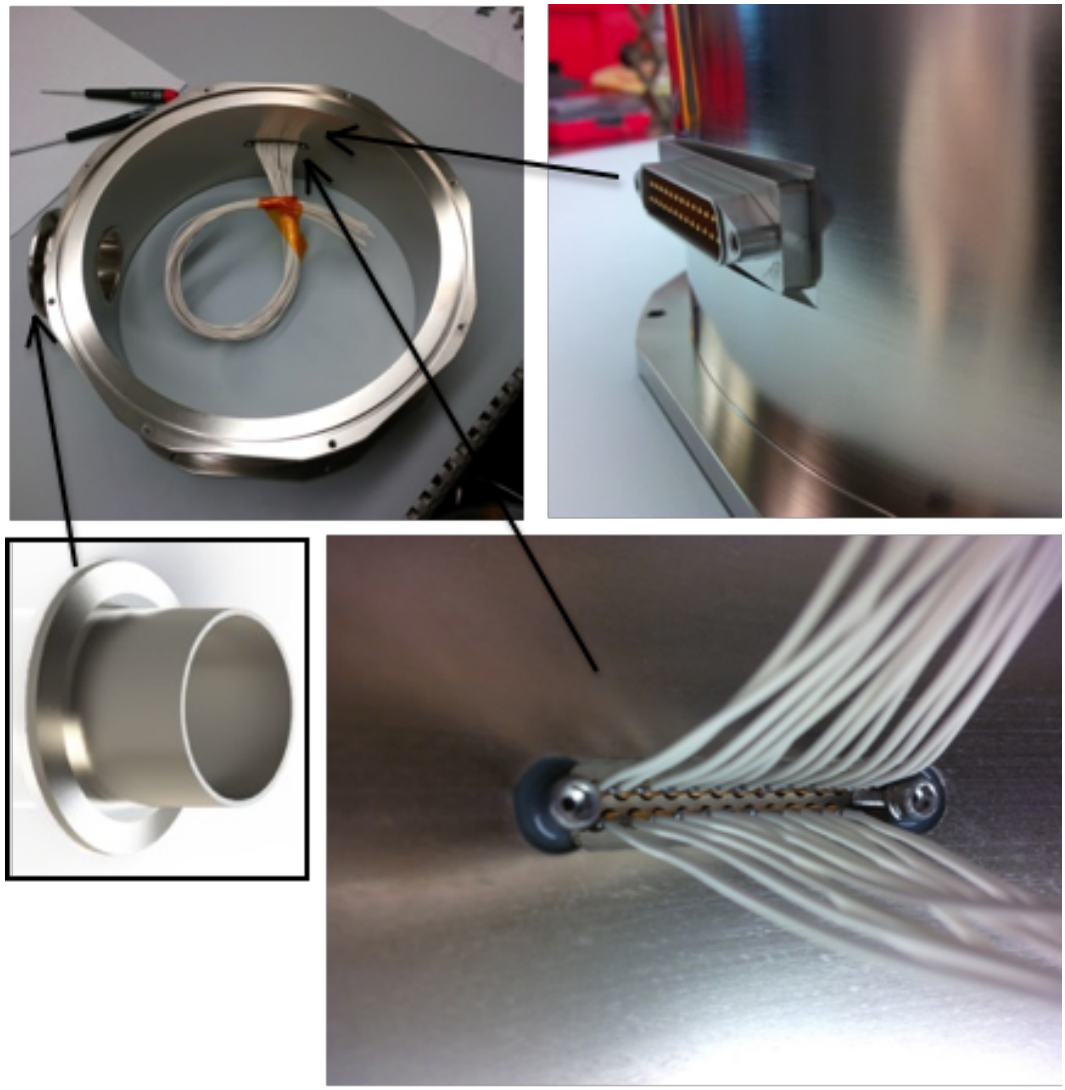

Figure 4.6: Components of the middle shroud. Hermetic seal connector from the outside (top right) and inside (bottom right). The other port in this shroud—a standard NW50 half nipple (bottom left)—is for the cryocooler attachment.

Mechanical Cryocooler Interface

I requested cooler fins and an NW50 flange for the cryocooler from Sunpower at the time of purchase. I then had an off-the-shelf NW50 half nipple welded into the shroud.

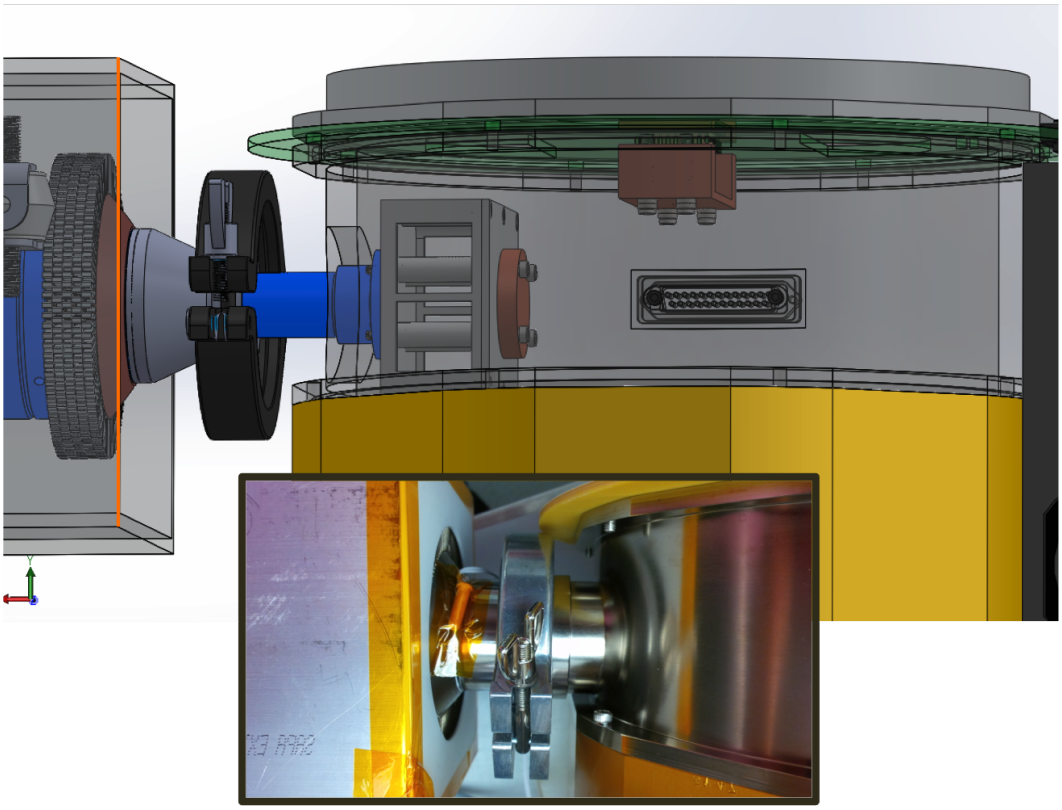

Figure 4.7: Rendering of the cooler-dewar-getter attachment without the half-nipple and inset of the hardware. The cooler is hard-mounted to the heavy dewar. The cryocooler cooling fins and charcoal getter are visible in the SolidWorks model, with the getter fastened to the coldtip of the cryocooler. The PC board supports the EMCCD and attached cooling copper clamp at the center of the dewar. The PCB transmits all camera signals to the EMCCD and the micro-D connector on the dewar wall drives the heaters and temperature diodes.

Thermal Design

The EMCCD needs to be operated at a range of low temperatures. This is easily achieved by adjusting the input power of the cryocooler in combination with a 10-W heater behind the cold clamp that holds the EMCCD in place. Since the cryocooler must cool the activated charcoal to a lower temperature than the EMCCD, the charcoal getter is directly attached to the coldtip. From there, multiple thin and narrow cold straps (36 straps that are 0.003 in thick, 0.5 in wide, and about 3 in long) conduct heat while isolating the EMCCD from the vibrational noise of the cryocooler tip. A copper clamp conducts heat from the back of the EMCCD to these cold straps. The CryoTel MT cryocooler can be operated at a range of input cooling powers anywhere between 30 watts and 80 watts, in steps of 1 watt.

In my system, 30 watts input-cooling power corresponds to an EMCCD temperature of -97°C and a charcoal getter temperature of -122°C. At 85 watts input-cooling power, the EMCCD reaches -147°C and the charcoal getter at its far end reaches -177°C.

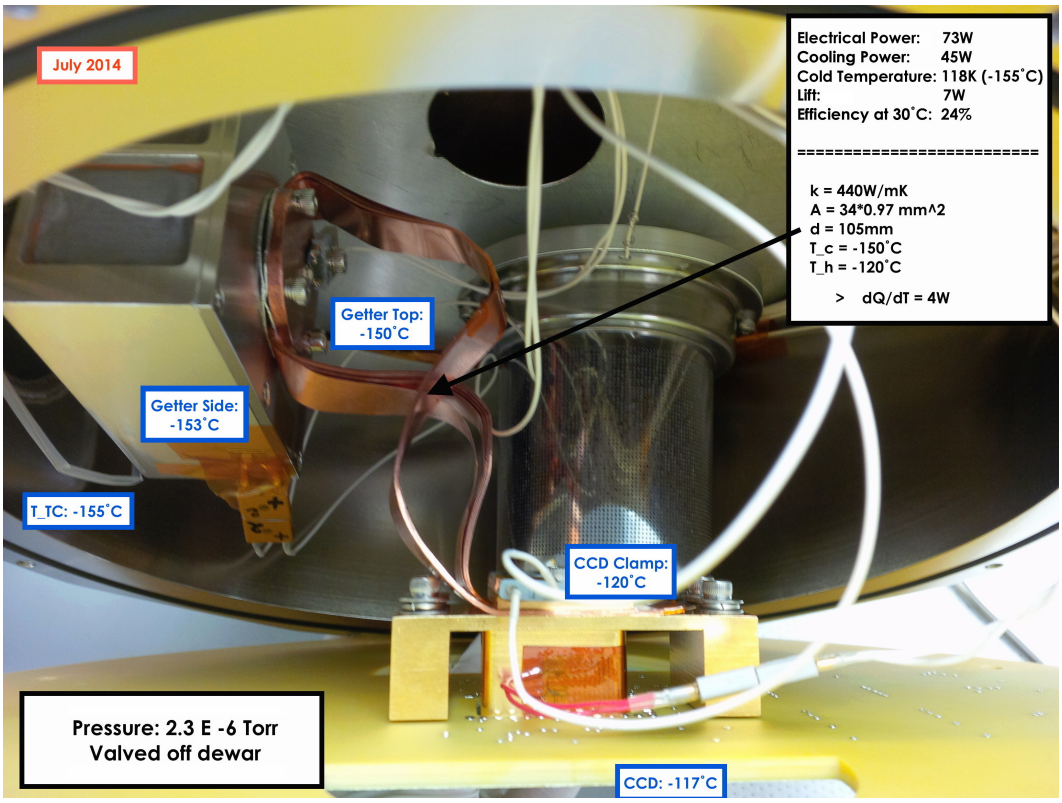

Figure 4.8: Full thermal link inside the vacuum chamber with temperature diodes. The new charcoal getter attaches directly to the coldtip of the cryocooler, just inside the chamber. A circular copper brace presses the cold ribbon onto the backside of the charcoal getter and is attached to the EMCCD copper spring at the other end. The 10-W Omega Kapton insulated flexible heater is placed on the copper spring, between the cold ribbon and the back of the EMCCD.

System Performance, Testing & Adjustments

To evaluate the vacuum of the dewar plus new vacuum chamber, the initial test camera setup had a dummy aluminum ring instead of the real PCB mounted between the two vacuum shrouds. This allowed me to get a baseline for the pressure that can be achieved with this new design, providing a means of determining whether the camera PCB is outgassing and/or too diffuse to maintain good vacuum. This was a valid concern, since FR-4 is a composite material made of woven fiberglass cloth and held together with an epoxy resin binder. All of the hardware that enters the

vacuum chamber, including the shroud and lid, was cleaned prior to assembly of the new system. The cleaning protocol is as follows, based on correspondence with John Klemic. First, all parts are washed for 20 minutes at 40°C in the ultrasonic cleaner with a soap solution. Everything is then rinsed in deionized (DI) water until the soapy residue is mostly removed, followed by another 20-minute ultrasonic wash in DI water. Finally, water is blown off of all of the parts with nitrogen and they can either be returned to the ultrasonic cleaner for an alcohol bath or a simple isopropyl rinse can be done. The system was pumped on for a few weeks and I did multiple nitrogen purges to clean all of the walls and other surfaces from any contaminants that had accumulated. The gold dewar itself had been open to the atmosphere for months and the inside walls are inaccessible for cleaning.

I mounted an EMCCD package at the center of the dummy aluminum PCB via a G10 plate to measure the temperature gradient in the system. This test showed that the 20mL charcoal getter did not have sufficient pumping power, and that the cold strap was not as flexible as I had envisioned. I designed and built a new charcoal getter that could hold 120mL of charcoal with four cylindrical poles in the center of the getter that uniformly cool this larger volume. I also made better cold contact with the new smaller cold straps by pressing them onto the back of the charcoal getter with four screws. These screws go through the center of the four posts inside the charcoal getter and screw directly into the coldtip of the cryocooler. They press the charcoal getter against the coldtip and the cold straps against the getter, with indium between all interfaces. This new getter performed very well at low temperatures, quickly removing water and other molecules from the chamber. The vacuum pressure is now directly correlated to the temperature of the activated charcoal. The two temperatures, one for the charcoal getter and the other at the EMCCD, can be easily regulated by adjusting the input cooling power and the cold straps that conduct the heat from the getter to the back of the EMCCD. The desired temperature will vary depending on the testing objective.

To explore the noise characteristics of a given chip, we performed dark current testing at temperatures much lower than the imaging temperature. Every EMCCD will perform slightly differently, depending, for example, on how it was processed, the characteristics of the wire bonding, and how the silicon itself grew, since not all devices come from the same wafer. For us, the different CCD201-20 devices will vary even more, because of the delta doping process and AR coatings. Thus every device must be individually tested for noise characteristics.

EMCCD Characterizing Performance

The first step in evaluating the performance of an EMCCD camera is to create a Photon Transfer Curve (PTC, Figure 4.9). Digital cameras are devices that take in light and put out digital data, which means that the noise associated with a particular camera can be easily determined. The only noise at the input of the camera is shot noise (SN), which is the fluctuation in the number of photons at any given illumination level. We can determine this noise exactly by finding the square root of the signal. The rest of the noise in the digital data comes from the camera electronics. The noise floor is set by the random noise of the output amplifier, which is called read noise (RN). In the case of an EMCCD, read noise becomes the noise level above which the signal has to be raised to reduce the total noise of the camera to sub-electron level.

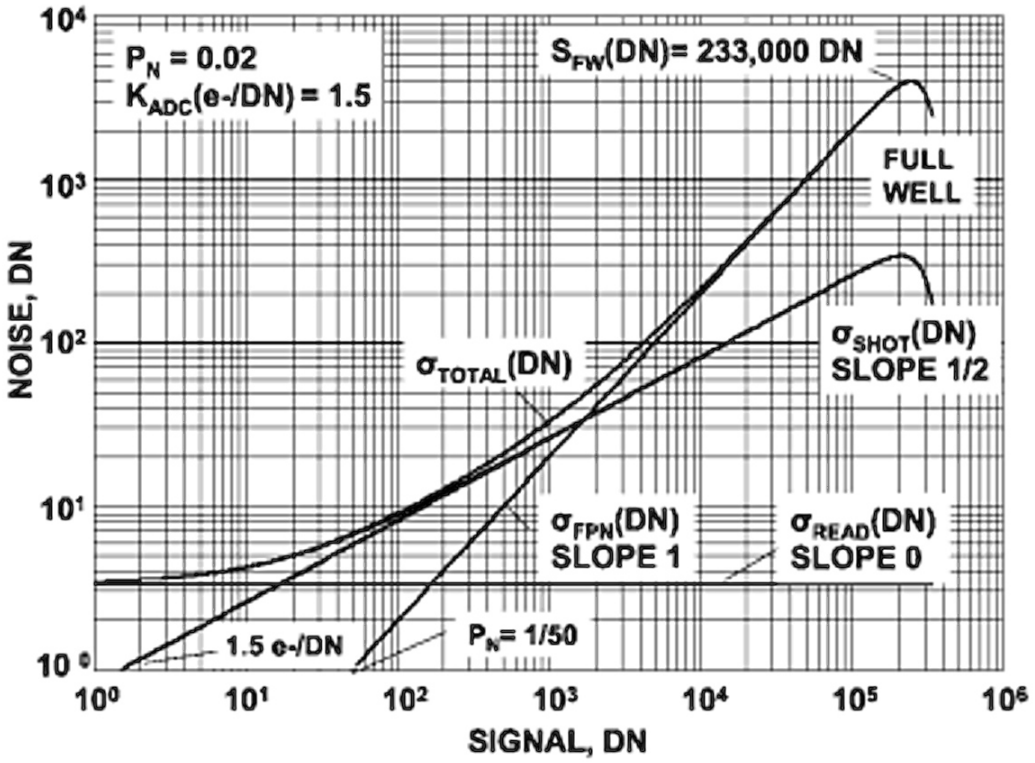

Figure 4.9: Photon Transfer Curve characterizing CCD performance. σ_{Total} data is plotted and three distinct slopes are added that correspond to Read Noise (RN), Shot Noise (SHOT) and Fixed Pattern Noise (FPN), as well a few other important parameters that are gained from this process. For far more detail than can be included in this writing, see James R Janesick (2001) and J. R. Janesick (2007).

Shot noise dominates at low light levels, but at high light levels, fixed pattern noise (FPN) takes over. FPN results from the photo-response non-uniformity (PRNU) of the pixels in a sensor array. Every pixel has a slightly different sensitivity, which is directly proportional to the input light. Pixels in an array can hold a fixed amount of charge, called full well capacity (FWC). Once this level is reached, blooming occurs, in which additional charge spills into adjacent pixels. When the charge in a pixel exceeds the saturation level, the signal degrades because blooming averages the signal between pixels and suppresses random noise. At this point, any increase in input signal will not increase the output signal; it will only continue to spill further across pixels and lower the noise. These clear noise characteristics make it easy to define the basic noise of a given EMCCD. To do so, a set of images is taken at different illumination levels and the results are compiled in one photon transfer graph (PTC).

To assess the dark current accurately, we first determine the conversion gain number (the number of electrons per digital number (DN)) of a given device, as well as basic noise statistics as described in the previous section. This can all be done with a photon transfer curve, as shown in Figure 4.9, which gives the read noise, full well capacity, sensitivity, dynamic range, gain, and linearity of the device. The PTC can be compiled from data that is taken by increasing the illumination of the device gradually. To find these parameters, the EMCCD has to be operated without signal amplification in the multiplication chain.

Our protocol for gathering this data consists of taking sets of three frames at every illumination level, taken in log-steps, with a set of 50 bias (dark) frames at the beginning and at the end of the illuminated images. The dark frames taken at the beginning and end show if there is any residual signal left after the device has been exposed to light for further analysis. We then make a "master" dark frame from the 0-second exposures taken at the beginning of the data series by finding the median signal in every pixel. We also median combine images of the same illumination and subtract the master dark frame from each image, leaving only the signal with an added $\sqrt{(2)}$ of noise. Next, the median signal is plotted against the median standard deviation, which includes the combined noise sources. The resulting graph should have three distinct slopes. At low integration times, the slope should be zero because in conventional CCDs the only noise at low light levels is the read noise, which can be determined by finding where the curve crosses the y-axis of the graph. The next section of the curve should have a slope of 1/2, corresponding to the shot noise,

followed by a slope of 1 for the fixed pattern noise. The curve will eventually turn over as it reaches full well capacity.

We explored a range of EMCCD temperatures to arrive at the best quality image that the EMCCD can produce for a faint source. Because photon counting only generates a small charge package inside the imaging array of the EMCCD, charge transfer becomes very important. Any leftover charge inside a pixel can appear as a signal and the small amount of charge transferred could be mistaken as noise, since we operate our EMCCD in binary mode. Our initial goal to operate the EMCCD at a target temperature of -125°C for low dark current with read-out speeds of 10 MHz for low CIC made good charge transfer impossible. Lowering the read-out speed increased the CIC by too much, but increasing the EMCCD temperature still yielded low enough dark current for shorter integration times. We found a good compromise between dark current and charge transfer at around -110°C. The vacuum pressure is directly correlated to the temperature of the charcoal getter, which should be around -160°C for a vacuum in the 10^{-6} Torr range. I also adjusted the cold straps between the back of the getter and the back of the EMCCD gold clamp. The overall size of the cold straps is important to find a balance between heat conduction and heat loss due to radiation in combination with vibration damping from the mechanical cryocooler. Originally I used four cold straps that were 0.01 in thick and 1.4 in wide and now I use 36 straps that are 0.003 in thick and 0.5 in wide, which makes for a system with better vibration damping between the EMCCD and coldtip (getter) while keeping the correct temperature at the EMCCD. Additionally, I used indium at all surface interfaces.

4.2 Electronic Design, Manufacture, and Integration

PCB

To Design the electronics for the camera I learned how to program in Altium Designer, which provides tools to create the schematic for the boards as well as PCB layout (see Figure 4.26).

Electronic Noise

An important consideration in my camera design was the choice of harness that feeds all EMCCD signals to and from the Nüvü controller. Ideally, the number of hermetic seal connectors in the signal chain is kept to a minimum to reduce additional electronic noise that can be introduced at this junction. A simple solution is the use

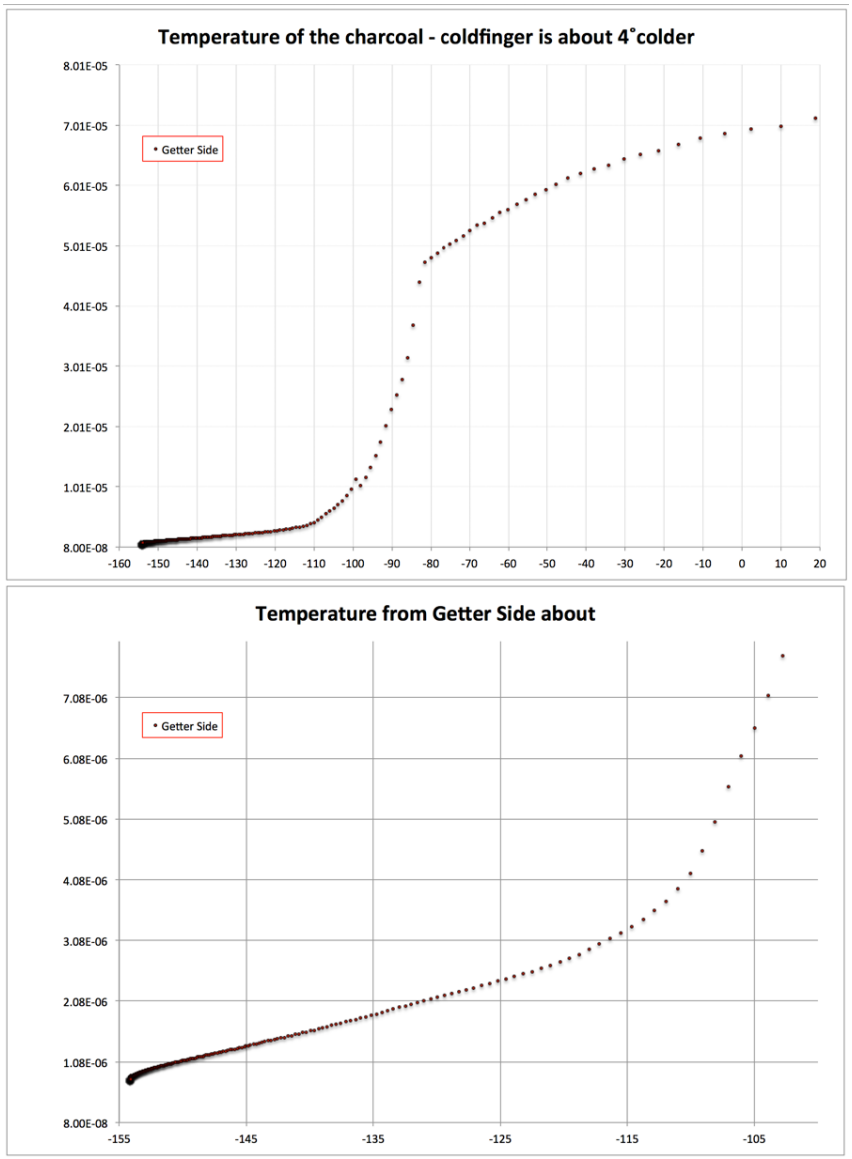

Figure 4.10: Pressure drop as the charcoal getter is cooled. The vacuum pressure drops exponentially between -85° and -105°C (top), and it continues to drop at lower temperatures (bottom). We start with a vacuum in the 10^{-7} Torr range at the beginning of a test run and the pressure rises into the mid-10^{-6} Torr range within a few days.

of a PCB that feeds the signals from the controller on the outside of the vacuum chamber to the EMCCD inside the vacuum. This setup assumes that the PCB can make a good vacuum seal with the aluminum shrouds of the vacuum tank, and the roughness of the PCB will determine how well the O-ring can hold the pressure inside the tank. The smoother the surface, the better the seal, and a surface finish of 128 micro-inches RMS is generally considered acceptable. After consulting with

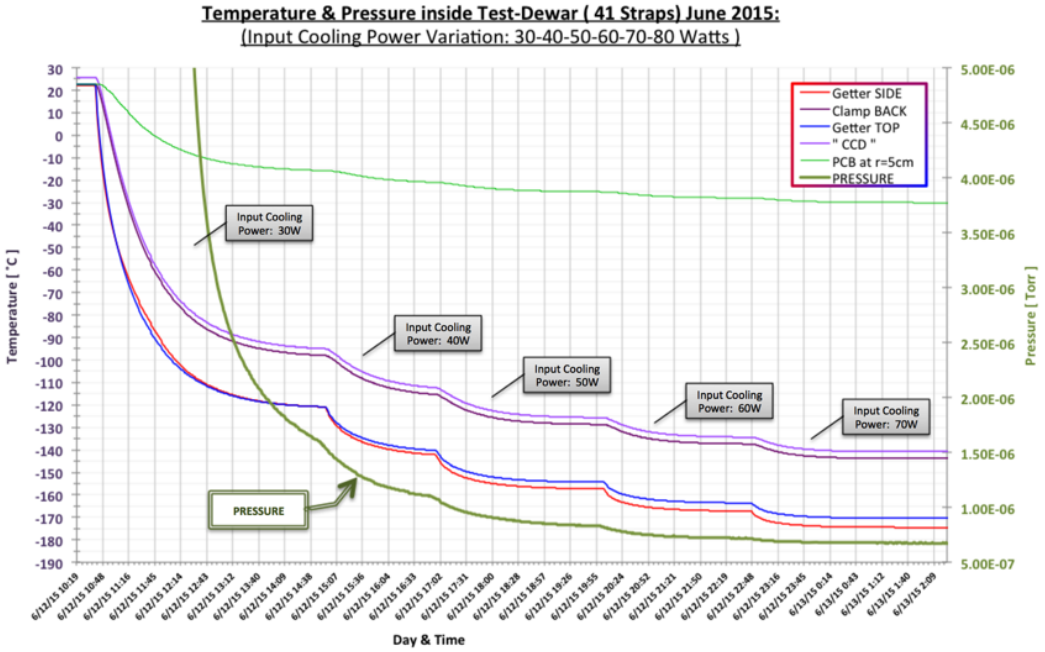

Figure 4.11: Overview of the temperatures and pressure inside the lab vacuum tank for increasingly lower temperatures of the cryocooler coldtip. Incrementally increasing the input cooling power lowers the curves and the vacuum pressure. With each cold step, the charcoal capacity increases and more contaminants are adsorbed. Also shown is the temperature of the PCB, which is nearly isolated from the cold chain. The diode measures the PCB temperature 5 cm from the EMCCD pins, which are the only conductive link in the electronic system.

PCB manufacturers and several groups who have used similar setups, I arrived at the following solution: for outgassing considerations, I used FR-4 material rather than polyimides, and instead of plating the PCB with gold at the sealing edge, I chose to leave the PCB bare. The finish of the solder mask is rougher than the bare FR-4, and I noticed that the gold plating came loose from the board for systems that were opened and closed repeatedly. With this design I am able to use only one connector at the controller and one at the PCB, with the rest of the signal chain internal to the PCB.

This PCB is very useful for several reasons. It is a rigid harness that preserves signal integrity; no matter what orientation the vacuum tank is in or how often I remove and replace the EMCCD, the signal wires inside the board will always experience the same environment. Because the inductance in each path will always be the same, we avoid adjusting the EMCCD waveforms to get the correct amplitude and signal phase to the pins of the chip. It provides internal wire feed-troughs in place

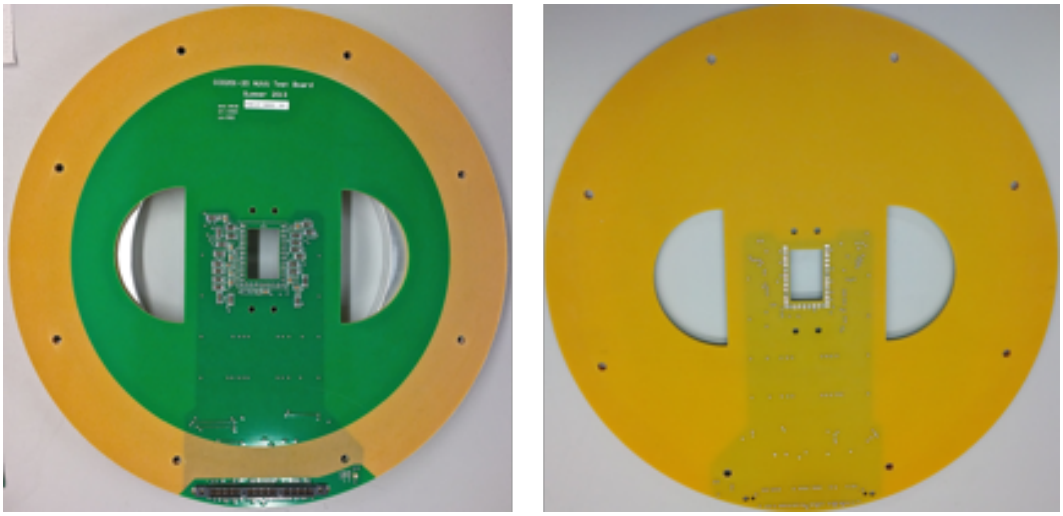

Figure 4.12: The signal tracks enter the vacuum dewar wall—indicated by the circle of holes within the board—with bare FR-4 visible. The front and back of the board is shown. The green material on the board is a solder mask, which protects the copper traces on the board from oxidation.

of another hermetic seal connector and the signal chain ends directly at the EMCCD pins. This structure also provides a rigid and reliable support for the detector, with a small enough thermal conductivity to keep the EMCCD at a low temperature. The center of the PCB supports the EMCCD, which is thermally isolated from the PCB by a small gap between the board and the gold clamp, which conducts the heat from the EMCCD. The only thermal contact is made between the pins of the EMCCD and the copper sockets inside the PCB. The board itself has a low thermal conductivity, but it will still get cold.

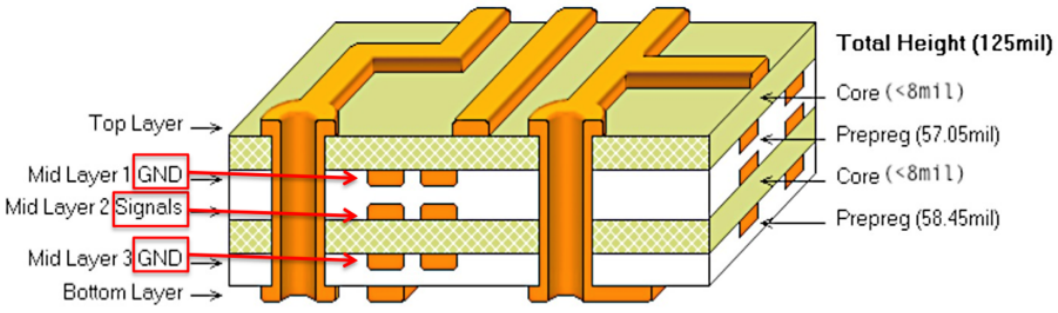

Figure 4.13: Layout of the board-layers inside the PCB, with specific thicknesses.

The signal tracks are on the middle inside layer of the PCB and are brought out to the EMCCD contacts inside the vacuum enclosure. In Figures 4.16 and 4.17 the RC filters can be seen close to the EMCCD sockets. These are low pass filters, which

give high frequency current a path of least resistance to the ground plane, while the low frequency current continues to the EMCCD pins. This will reduce any noise ripple left in the DC signals just before they get to the EMCCD. The 1 Ω resister together with a 0.1μF capacitor and a 10μF capacitor make a low pass filter with a cut-off frequency at about:

$$f_c = \frac{1}{2\pi RC} = \frac{1}{2\pi \times 10 \times (0.1 \cdot 10^{-6} + 10 \cdot 10^{-6})} = 1.6 \text{ KHz}. \tag{4.1}$$

It is good practice to put a serial resistance right at the output of the THS3091, which are the clock drivers of the Nüvü controller. This will shield the clock driver's feedback loop from the capacitance of the transmission line and CCD pin, because the THS3091 is a current feedback OP Amp. After consulting with Olivier Daigle, I used 10 ohm for horizontal clocks and 1 ohm for vertical clocks. I made sure that the resistors are rated for high power, which in this context is about 0.5 watt, as the current can peak at 350 mA.

I also added two grounding plates, above and below the signal traces, which are in the center of the board. Figure 4.13 is a generic output from the software, which is why the solid copper planes appear as traces.

Figure 4.12 shows the bottom of the PCB and the internal ground planes are visible as the darker shaded area of the board. To reduce the heat conducted to the CCD, I made these copper sheets as small as possible so that the cryocooler has to overcome a smaller heat load. To minimize the inductance of the transmission line, I placed grounding traces next to every signal wire, providing a low inductance path for the current at high frequencies. This will keep current from crossing over to another signal trace, because it will cross over to a ground path instead. I chose the FR-4 material because it has very low outgassing compared to other PCB material. However, it does have the highest outgassing rate of all the materials used in the dewar. Because it is composed of woven fiberglass cloth (with an epoxy resin binder), it soaks up water from air when exposed to atmosphere for more than two days or so. To overcome this, I bake the board in a vacuum oven for 3–4 hours, removing the bulk of the water, after which I can quickly get down to a good vacuum.

Because the EMCCD has such a low operating temperature, I looked into the thermal expansion properties of various materials at the center of the PCB. The temperature gradient for the EMCCD goes from room temperature to about -120°C. Even though the PCB has low thermal conductivity, it will cool down close to the sockets of the

chip, which is where the surface mount components are. The mismatch of thermal expansion for ceramic and FR-4 could cause the components to pop off of the board. To address this issue, I made the top layer of the PCB from FR-4 thermount (Tarzwell, 2010), a material with an adjustable thermal expansion coefficient. In my case, the coefficient of thermal expansion of the board is now matched to that of the ceramic surface mount components, and the camera can be temperature cycled without issue.

I adjusted the thickness of the various layers of the board for noise considerations. At low frequency, current will always find the shortest path, while at high frequency, the lowest inductance path will take precedence. The core between the signal plane and ground plane is 5 mils thick, a shorter distance than the gap between traces on the same surface, which are 8 mils apart. In this configuration, shown in Figure 4.13, the return current will jump to the ground plane rather than to the ground trace next to the signal path. I also used vias to tie the ground planes and traces together along the signal tracks for multiple return current paths. This extra step should eliminate any crosstalk between signals along the transmission lines. Crosstalk can, of course, still happen at the EMCCD pins, which are spaced very closely together, and I tried to locate as many grounding traces in the vicinity of the sensor as possible. It is difficult to tell what the true signal at the pins actually looks like, because the wires of our oscilloscope probes introduce noise themselves. Once they are removed, the signal improves, which I was able to document in images taken both with and without the probe (see Figure 4.20).

Thermal Considerations

The green material on the board in Figure 4.12 is a solder mask, which protects the copper traces on the board from oxidation, as well as any unintended electrical connection between two conductors, such as a bead of solder. The yellow material is the bare FR-4 board material and the darker shaded area indicates the location of the copper planes inside the board. FR-4 itself has very low thermal conductivity (k = 0.5 W/mK), but the copper traces and planes conduct heat very well (k = 400 W/mK). To reduce the cross-sectional area of the copper, I did not cover the entire board diameter with internal copper planes, instead they are only where they were needed to shield the signal traces.

The EMCCD sockets inside the board make contact with the cold EMCCD, and this heat is conducted to the copper planes by the ground connections (sockets).

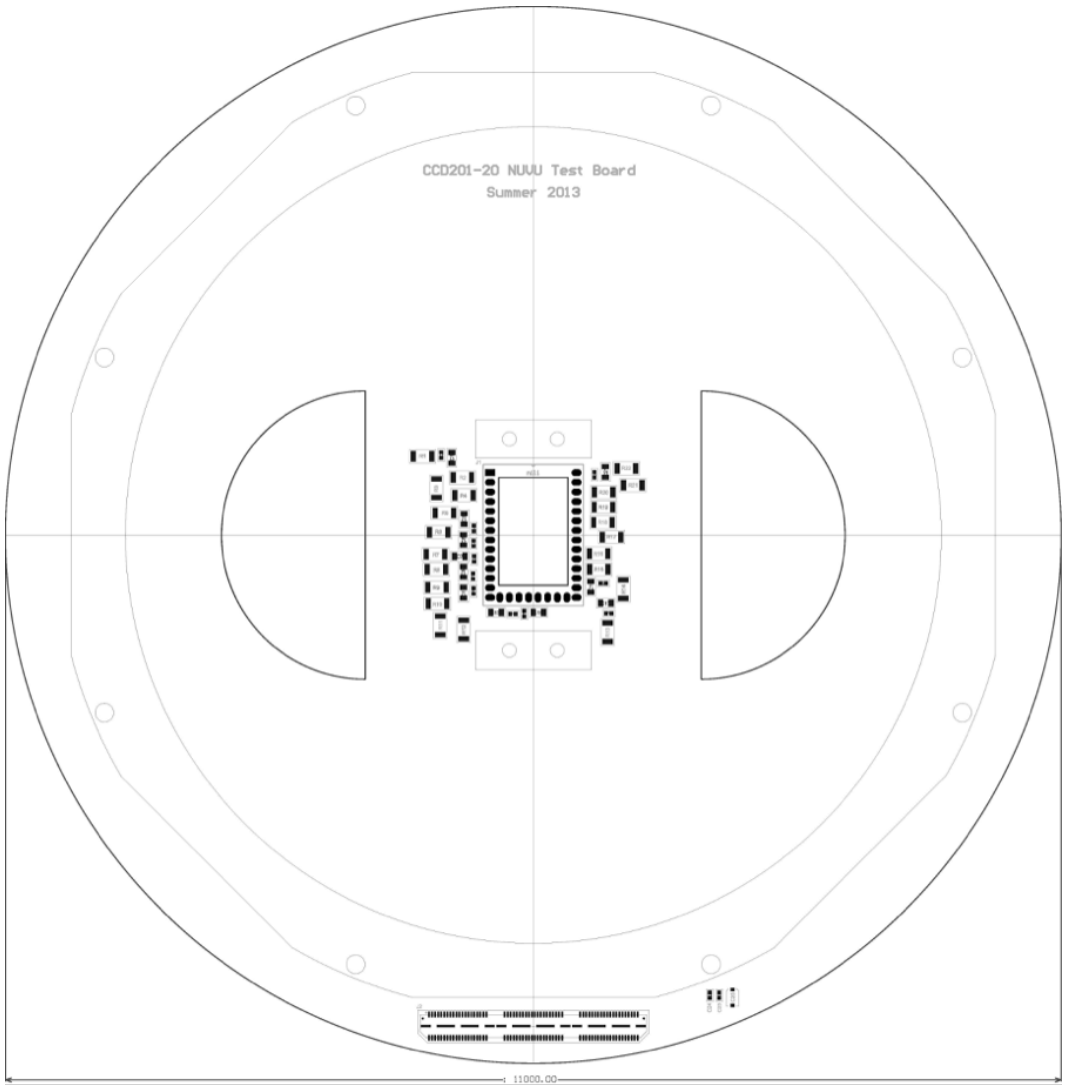

Figure 4.14: Lab PCB schematic. The vacuum shroud placement is visible in light grey, and the screw pattern and the footprints for the electrical components are stacked in close proximity to the EMCCD sockets. The SAMTEC connector footprint sits at the bottom of the board just outside the vacuum shroud, and the footprint for the HV clock capacitors is visible to the right of the SAMTEC footprint.

Thus, the FR-4 board itself cools down in the vicinity of the chip and also adds to the load of the cryocooler. Even though the board has low thermal conductivity, its large diameter will add to the heat that is conducted to the center. The rest of the heat load comes from radiation. The walls of the aluminum shroud are at room temperature and will re-radiate this heat inside the vacuum chamber. To lower the emissivity of the aluminum, the shroud can either be polished (emissivity coefficient ϵ = 0.039–0.057) or nickel plated (ϵ = 0.059), and my shroud and charcoal getter

are polished. The copper parts inside the chamber are gold plated to reduce their absorptivity, except for the copper ribbon. I exchanged the initial copper ribbon for one with a small surface area that is exposed to the radiation from the shroud.

4.3 Nüvü Controller

EMCCDs are a new kind of sensor that requires a novel technical approach in order to clock out frames at high read-out speeds, with multi-level clock-levels. The first controllers of this kind had just come on the market as we decided to use an EMCCD in our experiment. We purchased Nüvü's CCD Controller For Counting Photons (CCCP) and requested an adjusted controller that would support two EMCCDs with synched signals. This modified controller required some troubleshooting and we ultimately only needed one detector board, due to a change in design.

The controller came with a power cord and I purchased a SAMTEC EQCD high-speed coaxial cable to connect the controller to the PCB. There are two crucial design aspects for connecting the controller to the EMCCD: self-inductance and transmission line length. To minimize the inductance of the transmission line I tried to build coax signal traces by placing ground traces next to every signal transmission line with grounded copper planes above and below the signal plane inside the PCB. The core between the signal plane and ground plane is 5 mils thick, a shorter distance than the 8-mil gap between traces on the same surface, further helping with return current paths and signal crosstalk.

Our version 2 Nüvü controller comes with the option of using a real substrate voltage (VSS) or a virtual VSS. The controller has a real VSS signal output that is driven by a higher current OP-Amp, and the controller can be easily modified to use this option (by exchanging a couple of resistors). For virtual VSS operation, the substrate is tied to ground and all other voltage levels are adjusted to achieve the desired substrate voltage. Choosing this mode, I tied the EMCCD VSS bias to ground and linked all of the ground planes and ground return paths to these pins. This results in the smallest inductance return for the current.

A critical design aspect is the total transmission line length. The distance between the clock drivers and the pins on the EMCCD can negatively affect a clean signal due to propagation delay. Nüvü performed extensive tests and showed that the edges of the clock signals that arrive at the sensor pins with a delay of more than 2 ns will overshoot and consequently degrade the signal. To determine the allowed length of the signal chain, I assumed a speed of $0.6c$ for the signal inside the wire, which

gives a total length of 36 cm for a propagation delay of 2 ns. There is already a 100 ps delay between the clock driver and the connector on the inside of the Nüvü box and the PCB traces are about 70 mm long. That leaves a little less than 30 cm for the cable from the controller to the connector on the PCB. For the lab setup this is not a problem, because the controller sits next to the PCB on the work bench, but this restriction presented a challenge for the controller attachment and for the flight camera inside the vacuum tank.

Hardware Troubleshooting

Figure 4.15: Top and bottom left: dual channel CCCP, internal detector boards are not accessible, due to a permanent ribbon cable on the side. Top and bottom right: visit with Olivier Daigle at his lab in Montreal to test my camera system end-to-end (controller, computer, chip, and PCB).

On receiving the custom Nüvü controller, I immediately realized a problem. I was unable to configure the controller in my lab via software and I could not access the detector board inside the housing because of the way it is mounted and wired. I consulted with Olivier Daigle, the head engineer from Nüvü, who logged into my computer remotely to assess the situation. After weeks of unsuccessful remote trouble-shooting, I flew to Montreal to collaborate in person. We were not able to

make the controller work with my PCB and computer while I was there, but Nüvü shipped the fixed controller back a few weeks later.

The problem lay with the frame grabber card, the software that accepts the images from the read-out electronics and stores them on the computer for access. The EDT card that Nüvü had initially shipped with the controller did not accept images, and several other versions were also unsuccessful. In the end, we used a MATROX card instead, which required entirely different software to communicate with the Nüvü controller and computer. This change left the software system in a convoluted state, but it enabled all of the different software components to communicate with each other. The software system was later improved and bundled, and it can now be accessed from one simple graphical user interface. It became the mainstream software package that comes with Nüvü's version 3 controller.

Although the hardware/software issue had been resolved and shown to work for the camera system in Montreal, my camera still failed to take real images. I observed that the data that I was reading out did not change in pixel counts for different illumination levels, or when I grounded the input to the video line. The video output also did not respond to changes in the temperature or parameters in the software, like the gain. It seemed as though the video line was not connected, and by probing the resistance of the video pin, I found that the line was indeed open. As it turned out, the load resistor of the video line was missing.

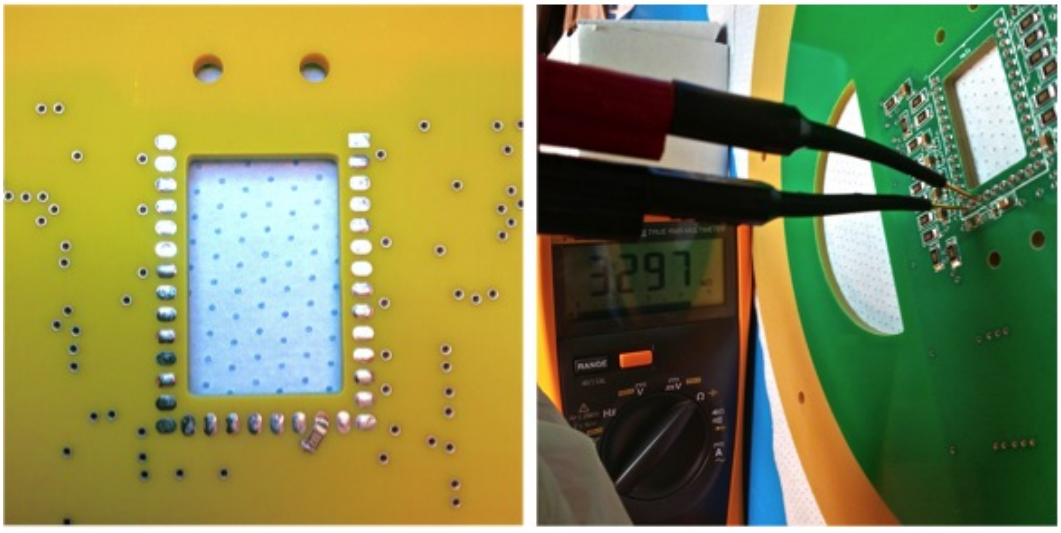

Figure 4.16: Placement of the 3.3K Ohm video load resistor on the back of the PC board

When we first ordered the controller, we had been testing e2v's CCD97-00 as our potential flight sensor. As the design of the flight instrument changed from two EMCCDs to only one, we decided to use e2v's CCD201-20. These two chips require different size charge resistors. Nüvü built the controller with a 5 KΩ resistor for the CCD 97 camera, but had removed it prior to shipping. To overcome this problem, I added the correct load resistor for the CCD201-20 to the back of my PCB (see Figure 4.16).

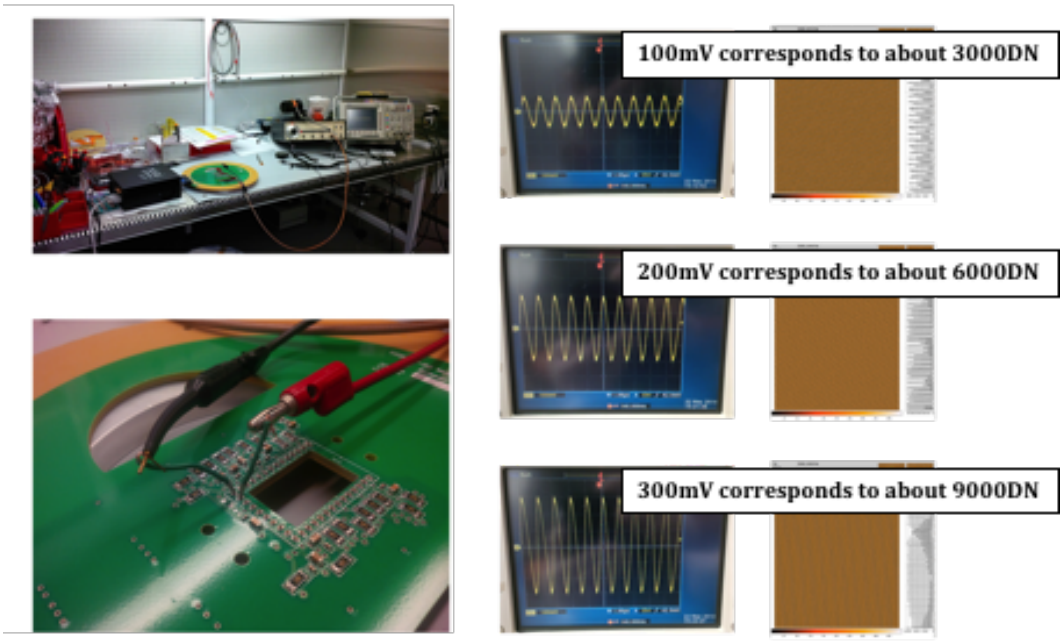

Figure 4.17: Testing the video channel with a waveform generator once the load resistor was placed on the back of the PC board. Top Left: The Nüvü controller is connected to the PCB on the workbench and a waveform generator (right) feeds the signal into the video channel of the board. Bottom Left: Black probe is connected to ground and the red lead feeds the signal into the video channel. Right: input of the sine signal with increasing amplitude and corresponding data output from the controller.

Once the 3.3 KΩ resistor was added, I performed tests to ensure that the video line was now correctly connected. I sent sine waves of three different amplitudes (100 mV, 200 mV, and 300 mV) into the video pin and took an image each time. I checked these images with IDL to compare the signal values, which correctly reflected the amplitude increases and confirmed that the controller was now finally talking to the video pin on the board. The read-out speed of the controller is 10 MHz, so I used a 1-MHz wave for sampling. The 100-mV amplitude wave corresponds to a max

pixel value of about 3000 DN, 200-mV corresponds to about 6000 DN, and 300-mV corresponds to about 9000 DN.

Programming and Optimization

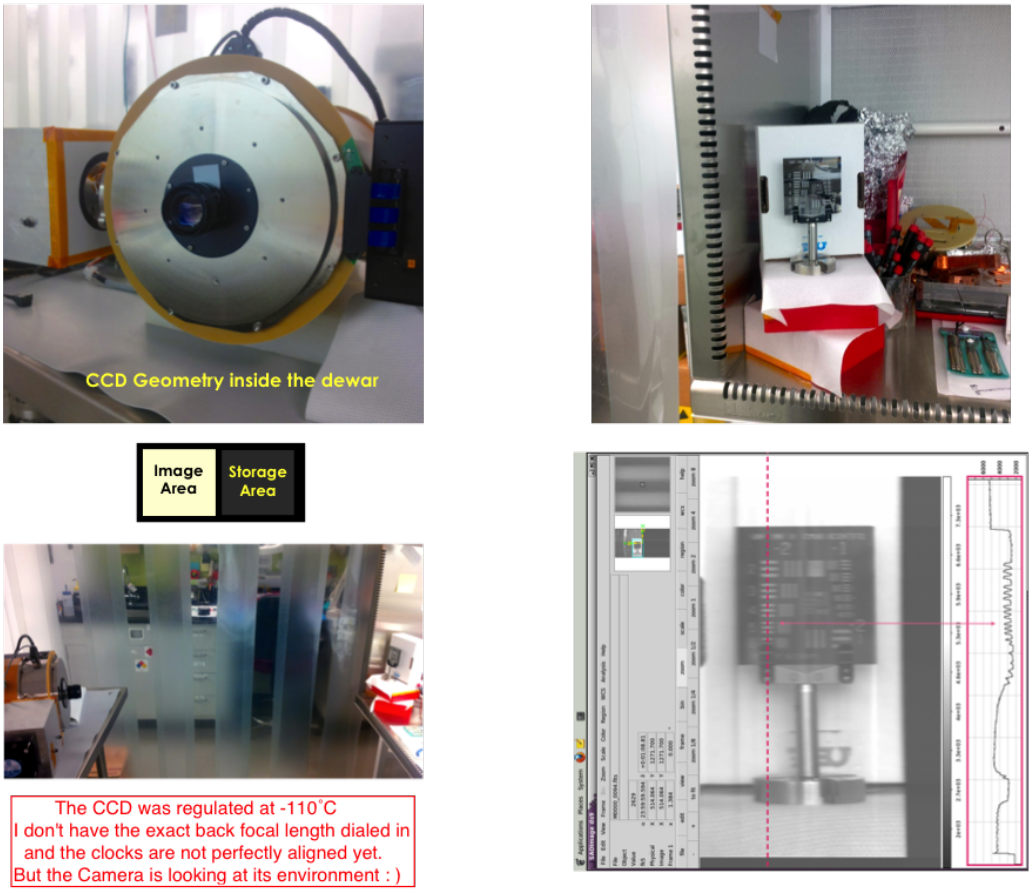

Figure 4.18: First light for the FIREBall-2 test camera in the lab. Top Left: View of the test setup and orientation of the chip behind the dewar window. Bottom Left: Camera and target geometry. Top Right: View of the target and surroundings. Bottom Right: First image of the EMCCD 201-20 shown in DS9 software. Now the real work begins as we adjust various parameters in the waveform script to optimize the camera and take images in electron multiplication mode.

The process of getting an EMCCD online starts with configuring the new sensor in analog mode and making it operational. Analog mode refers to setting the HV clock high rail to a low value, such that no charge amplification occurs in the multiplication chain of the read-out register. In this configuration, test images can be taken warm and at low illumination levels without saturation. Some experimentation with the different waveform parameters was required before the camera was able to function

properly. To create a digital image, I had to adjust pin voltages, clock overlaps, and the combination of timing diagrams necessary to move the charge from every pixel out of the sensor. With the first functional waveforms programmed, the new frame grabber card inside the computer, the resistor in place, and a new software package running, I was able to get the first image out of the camera. I placed a military target about half a meter from the camera, corresponding to the focal length of the lens attached to the dewar window. The sensor was cooled to -110°C and the image was integrated for one second. The resulting image shows that the waveform overlap requires further adjustment, but the system hardware finally worked properly.

The EMCCD is oriented horizontally inside the dewar with the shielded storage area of the sensor on the right (see Figure 4.18). Because the HV clock was not yet increased to amplify the signal, this is a regular CCD image. The next step was to slowly increase the HV clock and start amplifying the signal, to put the EMCCD in photon counting mode.

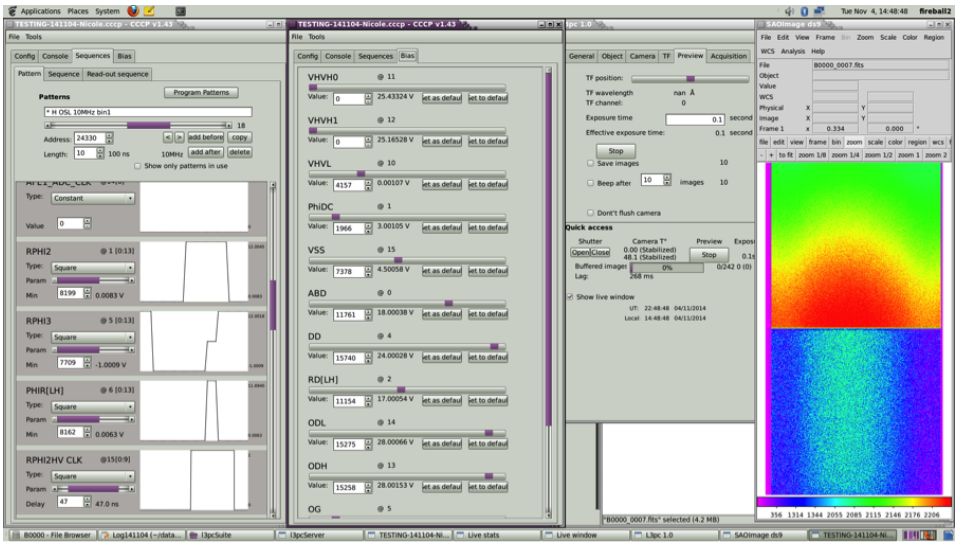

Figure 4.19: A screenshot of the software GUIs: serial register clocks are on the left, other bias levels, including the HV clock are in the middle, with the live image on the right for easy real time adjustments.

4.4 Camera Testing and Optimization

Waveform Programming

The controller software necessary to operate EMCCDs (Figure 4.19) is not as straightforward to program as conventional read-out electronics for CCDs. This complexity is inherent in the need for more intricate electronics that drive the fast-

changing clocks required to reduce the noise sources of an EMCCD. Any new controller that is built for EMCCDs will have more complex programming schemes in order to drive these electronics. For the CCCP, one difficulty arises from the very sensitive HV clock, which is driven by a resonance system that responds to any change in the circuit. The HV clock is a sine wave whose peak will shift when the inductance in the path changes, for example with different harness lengths or capacitance added to the transmission line (e.g., oscilloscope probes, seen in Figure 4.20). Another major difference from conventional controllers is the variety of read-out waveforms that can be programmed. Here the percent of overlap will vary, as will the duty cycles for the various shapes that the clocks can take on.

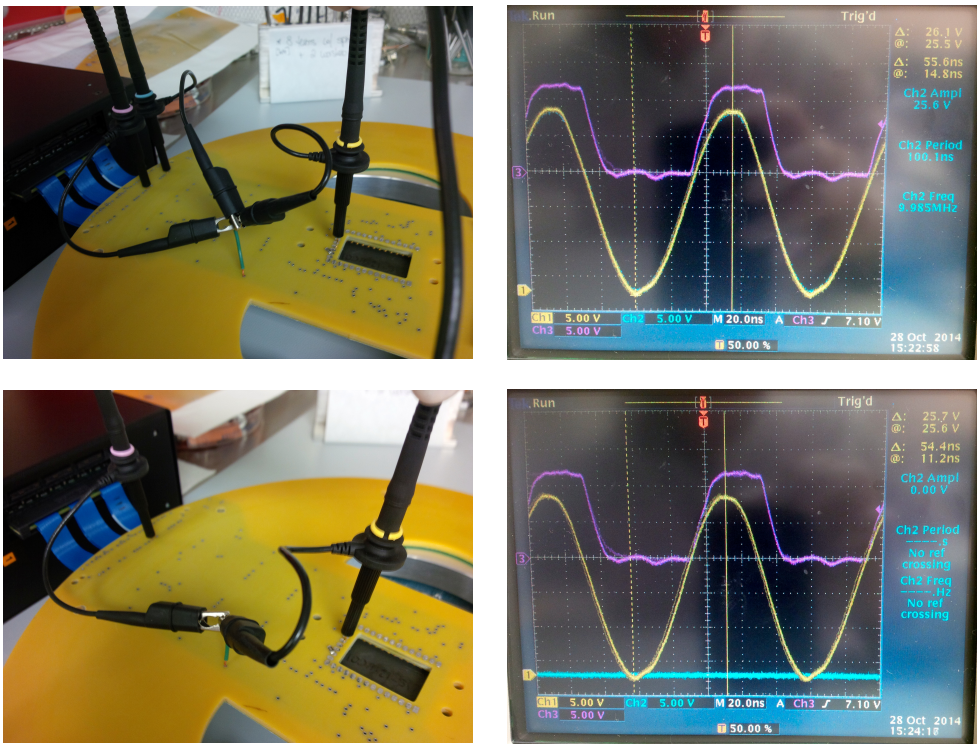

Figure 4.20: Testing the effect of added inductance from the probes on the resonance of the HV clock. The scope is triggered on serial register clock 2 and the HV wave amplitude peak shifts by about 4 ns when one of the two probes in line with the HV signal is removed.

I began by following e2v's read-out scheme from the CCD201-20 datasheet. I followed the pulse timing and clock overlaps for trapezoidal clock shapes. For better overlap during the fast charge transfer in the serial register, I adjusted the serial register clocks to follow a multi-level scheme instead. The first analog image

of the FIREBall-2 camera was taken with conventional trapezoidal clocks for the parallel transfer at about 200 KHz and multi-level clocks for the register clocks at 10 MHz. The first amplified image was challenging and required multiple hardware and software changes. I connected a set of probes to the test pins for the clocks on the PCB and investigated the overlap and relationships between the clocks (Right panel of Figure 4.22). The first change involved the capacitance in the circuit of the HV clock. My board design anticipated necessary small changes in capacitance with an adjustable capacitor, but I also added two more capacitors for a wider range of allowable values. The adjustable capacitor went from 6 to 33 pF and the two additional capacitors were 51 pf each. By removing one of the 51-pF capacitors and tweaking of the adjustable capacitor, I achieved resonance. By putting the HV clock on the scope together with serial register clocks 1, 2, and 3, I could verify the change in amplitude of the HV clock as I was adjusting it in the software and show that it was now working properly (Figures 4.21 and 4.22).

I moved the low rail of the HV clock to about 2 V and raised the high rail in slow increments to about 42 V. This was insufficient to initiate the multiplication process, and I had to also move the rise and fall time of the register clocks until they were completely different from the scheme specified in the e2V datasheet. I also adjusted the reset pulse timing and amplitude together with the timing of the sampling clock of the charge capacitor. This finally resulted in an amplified image. To see the amplification, I placed a cover with three small pinholes above the imaging area of the chip with very little ambient light in the room to illuminate the sensor. I took short images at three different high voltages and showed that the pinhole images on the EMCCD increased in brightness. With the hardware and HV clock now working, I then focused on eliminating various noise sources that were present in the data. My main effort went towards proper charge transfer in the parallel as well as serial registers, with many parameters to adjust for a better quality image. I consulted with Olivier Daigle on a proper clocking scheme for my EMCCD and implemented sinusoidal clocks for the parallel transfer of the EMCCD. I later improved the illumination scheme for more controlled experimentation and attached some optics to the front cover of the dewar window, so that I could attach a light diode inside a dark cylinder. I connected this diode to the trigger output (5V) of the Nüvü controller with a voltage-lowering resister in line. This setup provided a great repeatable illumination level that I used to improve and test the amplification of the EMCCD201-20.

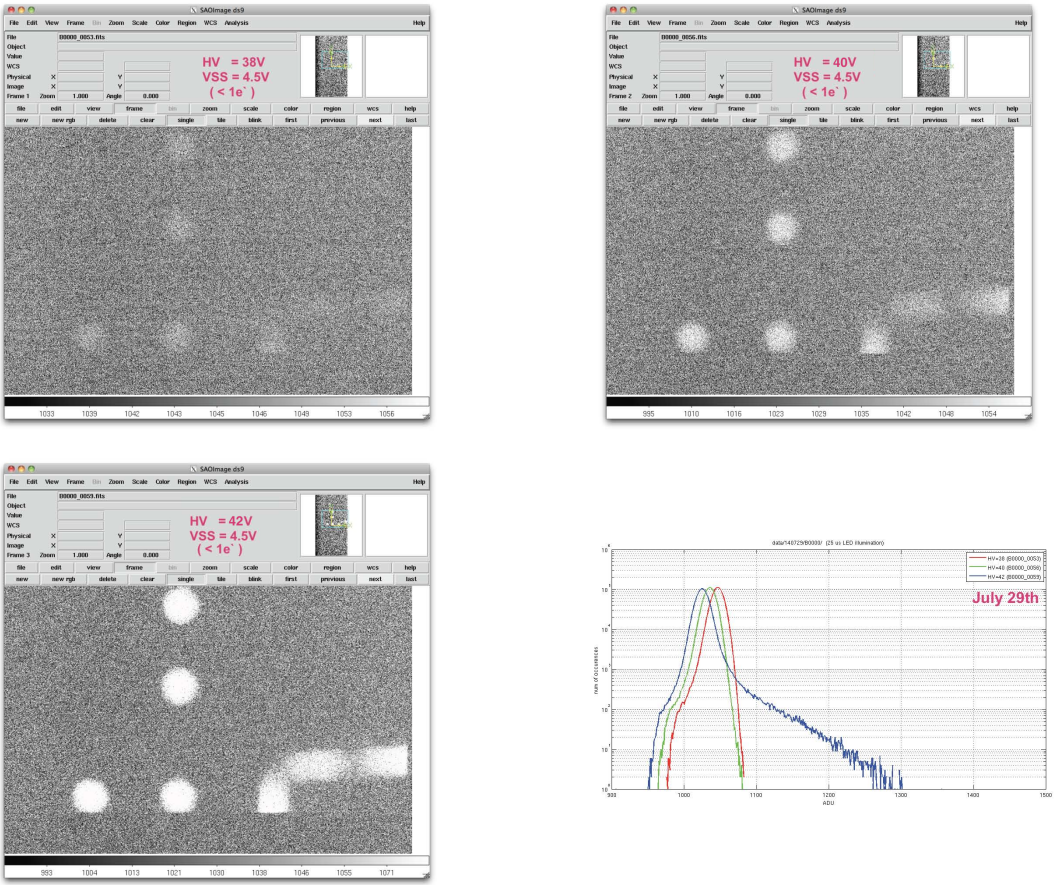

Figure 4.21: The first EMCCD images of a few faint pinhole spots that have gone though a multiplication register of increasing gain. Top left: HV clock high voltage 38 V (red trace on the graph) Top Right: HV clock high voltage 40 V (green trace on the graph) HV clock high voltage 42 V (blue trace on the graph). Bottom Right: EMCCD image histograms. The blue trace shows that some events have been amplified and moved out from under the read noise bell of the curve.

Waveform Shaping

For the next few months I continued to adjust the waveform script in various ways with varying results. I went back to trapezoidal shapes for the parallel transfer and adjusted the overlap as well as the shape and amplitude of the register clocks. During this time, I achieved a slow but steady improvement of the image quality and learned to work with EMCCDs and the Nüvü controller.

Help was received from Roger Smith, who has extensive knowledge of CCDs, and one of his engineers, as well as from two new postdoctoral scholars who joined our team and started to work on the camera with me. We were able to make great progress, ultimately developing a script that produced images with acceptable levels

of clock induced charge and dark current. Over the following few months, the final waveforms for the flight were developed; these could likely be further improved for months to come.

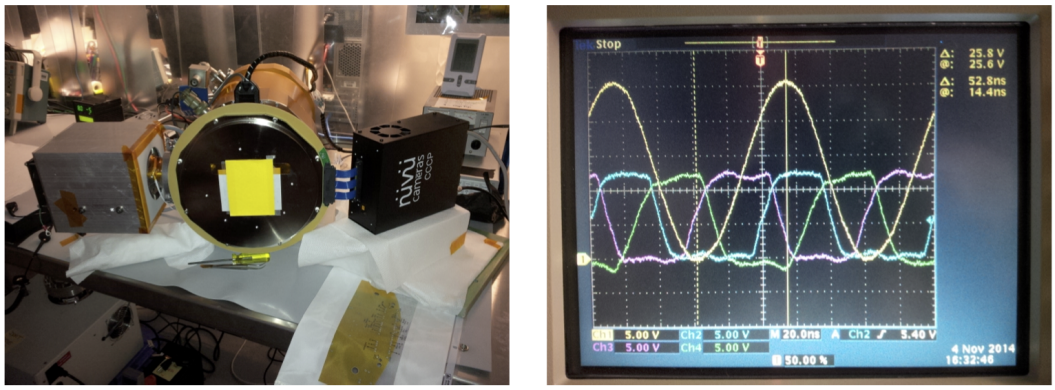

Figure 4.22: A look at the different waveforms on the scope. The HV clock will shift by about 4 ns when the probes are removed from the circuit. Most of the noise/crosstalk visible in the waves is added by the long probe wires.

4.5 Flight Camera Design and Integration

Flight Vacuum Tank

The flight vacuum tank will be more difficult to control than the lab dewar, because it has a larger volume and much larger surface area. It houses multiple bodies with surfaces that can continue to outgas for a long time. We will have to start pumping on the flight chamber long before we start any camera testing, to remove as much of the initial vacuum contaminants as possible. The total coldlink inside the vacuum tank is much longer and wider as well, requiring that the cooler be run at a higher cooling input power.

Flight Camera Hardware

The flight PCB is a smaller-diameter version of the test PCB. The main differences between the flight and test scientific camera are the PCB size, the controller cable length, and the coldlink and cold strap lengths. The PCB is attached to the back of a mirror mount and will not be used as a vacuum seal for the flight instrument. To feed the required cables into the vacuum chamber, I have designed two flanges that follow a standard DN 100 design at the vacuum interface. I worked with Douglas Electrical, who can provide hermetic seal solutions for electronic cables needed to operate the read-out electronics.

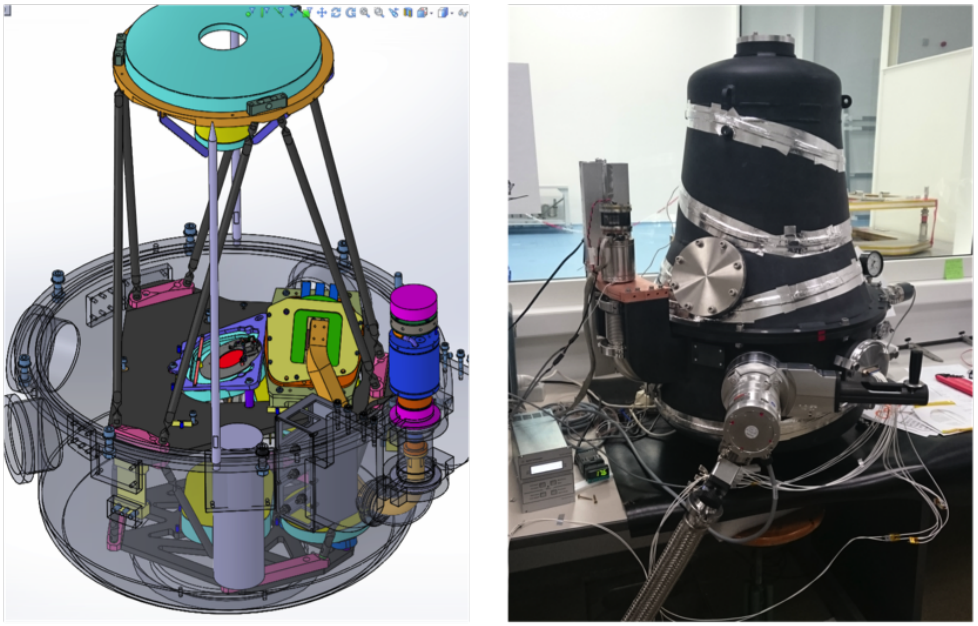

Figure 4.23: Left: SolidWorks model of the flight vacuum tank, including the cylindrical zeolite can and rectangular charcoal container mounted to the front of the tank wall, next to the blue and purple cryocooler. The green PCB with gold plated copper link is also visible inside the tank, just behind the charcoal getter. Right: an image of the tank during initial testing. The silver cryocooler is mounted on top of the copper blades on the outside of the vacuum tank to the left and the circular feed-through for the controller has been closed off with a flange for this cooling and vacuum pressure test.

Due to the limited space on the smaller diameter board, some changes to the wiring layout of the flight PCB design had to be made. Nevertheless, I was able to mimic the lab design and every signal trace is still surrounded by a ground trace for signal isolation. The flight board behaves similarly to the lab board, and we have been able to bring it online without any complications.

I also designed and built zeolite and charcoal getters for the flight vacuum tank, each with a larger volume than the lab versions to accommodate the larger vacuum tank. The flight charcoal getter attaches to the gold plated copper coldfinger. It is the first thermal load on the cryocooler and can be kept at a much lower temperature than the EMCCD, which attaches to the end of the coldfinger via a flexible cold ribbon. The flight coldfinger is significantly longer than the lab version, but it is a solid copper link with only one joint. The joint connection has a layer of indium for better surface contact and thermal conductivity.

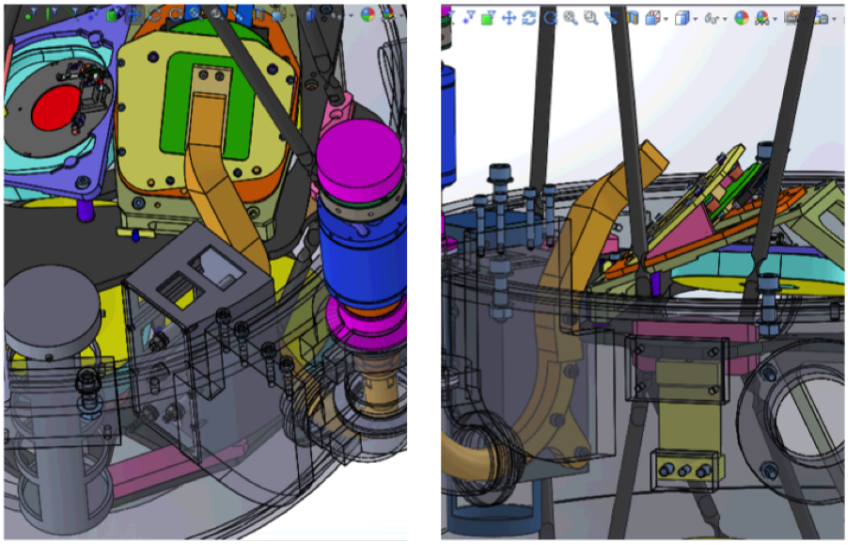

Figure 4.24: Left: close look at the mounting scheme for the zeolite can and charcoal container inside the tank. The back of the green PCB with the coldlink attached is also visible at the top of the image and the mechanical cryocooler is to the right. Right: A view from the back of the charcoal getter to show the coldlink attachment, with a side view of the mounted PCB.

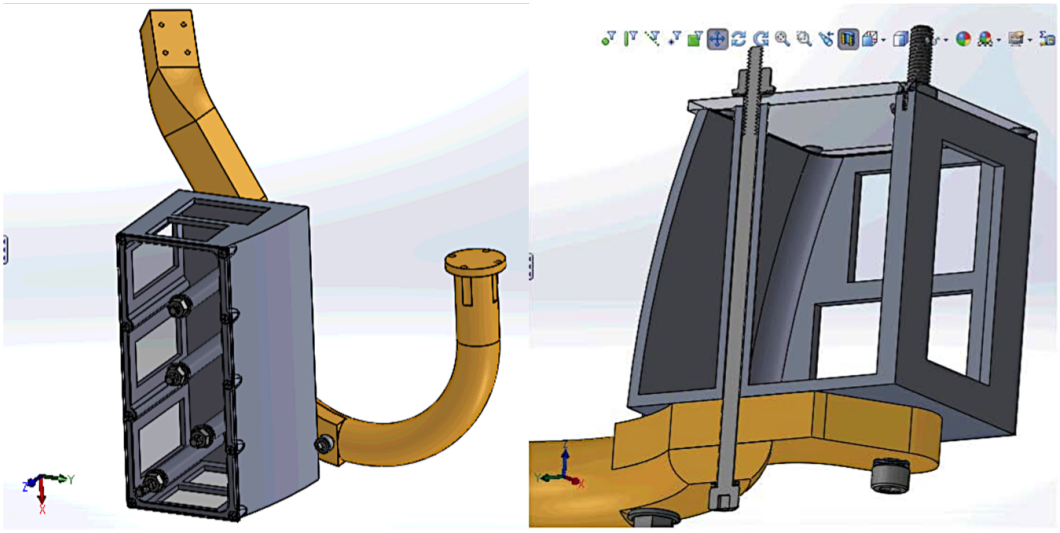

Figure 4.25: View of the Flight coldlink and charcoal getter. The getter is pressed against the coldlink with four long screws inside hollow cylinders that will also keep the charcoal uniformly cold.

Flight Read-Out electronics

For the first FIREBall-2 flight, we will use Nüvü version 2 controller because it has been extensively tested and its performance has been consistent and reliable. For the first flight, I designed and built a pressure box enclosure for this controller, because

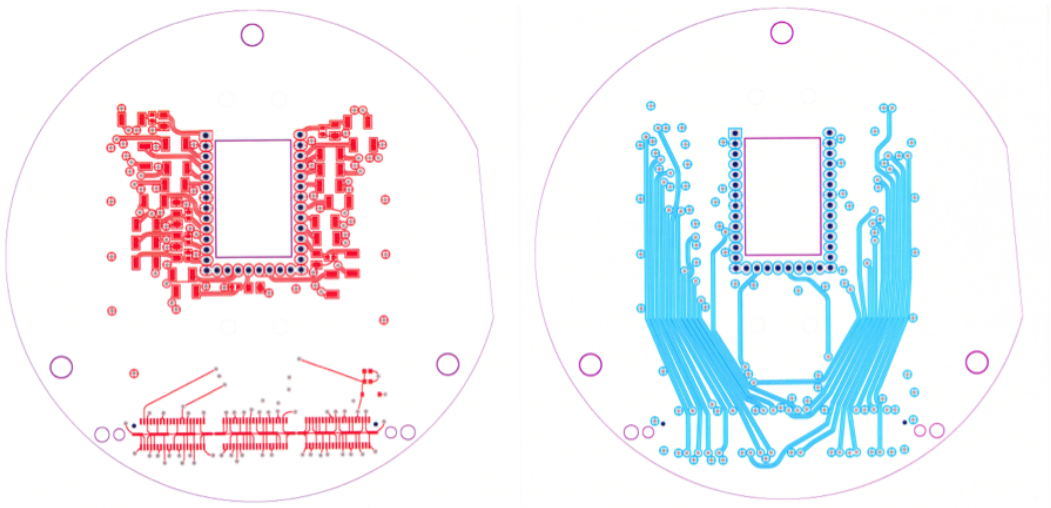

Figure 4.26: Flight PCB drawings with purple board cut-outs for the EMCCD at the center of the board and several screw clearances. Left: the top layer copper traces are shown in red. The electrical components can be seen as well, resistors and capacitors close to EMCCD pin sockets and the EQCD connector at the bottom of the board. Right: the middle layer copper traces are shown in blue. All signals are fed through the center of the board for noise isolation. This board has four layers, two of which are copper planes above and below the center signal traces for shielding.

some of its electrical components (capacitors) are not rated for low pressure and will explode at high altitude. The detector board inside the controller can operate up to 70°C before it shuts itself off, and the controller radiates about 60-W of heat. Therefore, the pressure box has been designed to attach heat fins to remove heat from the inside of the box during flight. During lab testing, a large radiator in combination with a fan has been attached to the side of the pressure box to remove heat efficiently. The overall geometry of the pressure box has to take into account the limited space inside the gondola due to the rotation of the vacuum tank. The pressure box is attached to the upper part of the vacuum tank, as this will result in the shortest distance from the controller to the EMCCD. This meets the maximal ribbon length requirement, which will preserve the signal integrity of the EMCCD clocks (see Section 4.3). The vacuum interface follows a standard DN 100 design, but the flange is tilted by 9° to keep enough distance between the power cable connector on the back of the flange and the vacuum tank. At the same time, this minimum tilt will leave enough room on the outside of the box for rotation of the vacuum tank close to one of the gondola trusses. Although the pressure box does not experience the large pressure gradient that the vacuum vessels (lab and flight) withstand, I still

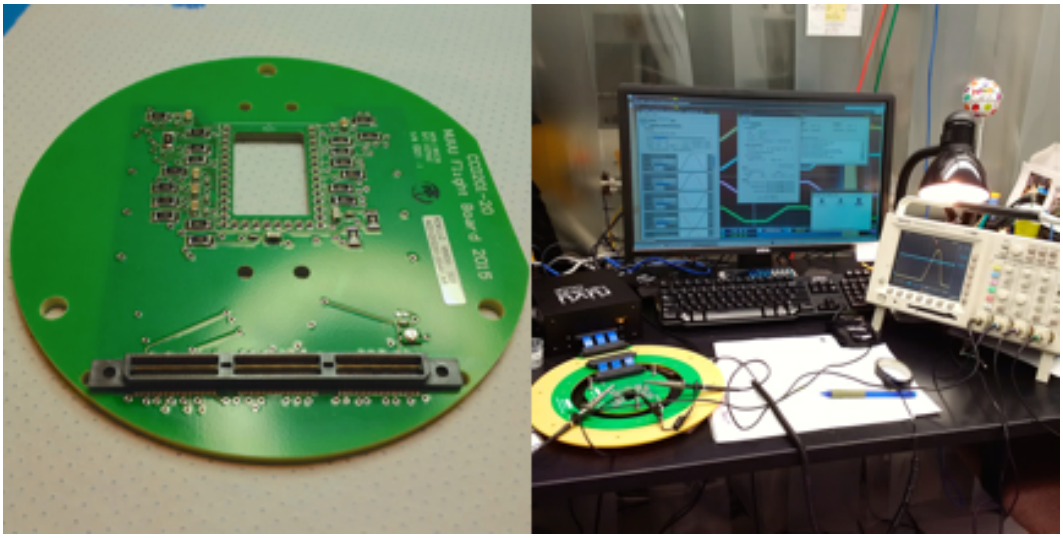

Figure 4.27: Left: Flight PCB. The electrical components (resistors and capacitors) can be seen close to EMCCD pin sockets at the top and the EQCD connector at the bottom of the board. The adjustable capacitor for the HV clock is just above the EQCD connector to the right. Right: Board verification and continuity testing of the flight PCB. The smaller diameter flight PCB is at the center of the interface ring, which is also the wire feed-through into the dewar. It will hold the PCB in place during lab testing and double as the vacuum seal. A second interface ring was made to test the flight PCB with the version 3 Nüvü controller.

paid close attention to the design of the O-ring grooves and wall thickness and I anticipate no complications at altitude.

The controller is attached to the pressure box via four small stand-offs. These keep air circulating at the bottom of the box and hold the controller rigidly in place during flight. To improve the heat removal from the detector board inside the black controller housing, we have removed the top part of the Nüvü enclosure. The pressure box has been closed for testing and will remain closed until after the flight to keep the electronics shielded and prevent any damage.

Flight Hardware Tests and Adjustments

I verified the geometry of all of the hardware components with the team of mechanical engineers at LAM in France, who are responsible for the design and manufacture of the flight vacuum tank. All of the components have been manufactured and shipped to France, where they have been integrated into the flight tank and are currently being tested. The CryoTel MT cryocooler was successfully installed and operated, but because initial tests showed that we would operate the cooler at

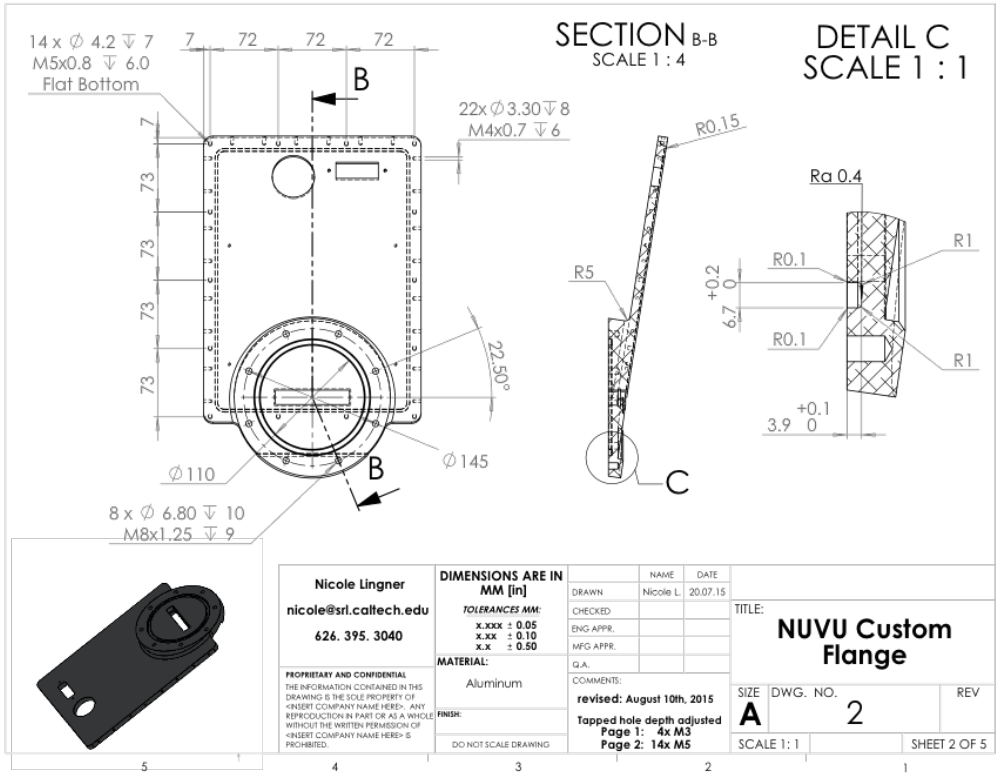

Figure 4.28: The pressure box design was difficult to machine, but Accurate Manufacturing Company in Glendale delivered a flawless final product.

the maximum input cooling power, we decided to purchase the larger CT cooler to ensure proper operation.

A more difficult task was to find a vacuum seal solution for the SAMTEC ribbon cable. It is important to reduce as many connectors in line with the detector signal as possible to minimize the noise that can be introduced at these interfaces. After a few sample runs with Douglas Electrical Components, I decided to seal the cable directly at the cable pigtail-PC board interface. This also allowed for the shortest cable distance inside the pressure box. The controller has been tested inside the cooled box and performed soundly during the SIMEON test described below.

The assembled flight vacuum tank has been undergoing tests for several weeks as of this writing. So far, a few vacuum leaks have been identified and fixed and the vacuum vessel has been kept under vacuum and cooled on and off over many weeks. The initial pressure inside the tank could only be maintained by a combination of the turbo pump and molecular getters, due to the high outgassing rate of multiple surfaces inside the tank. After several weeks of repeatedly heating and cooling the

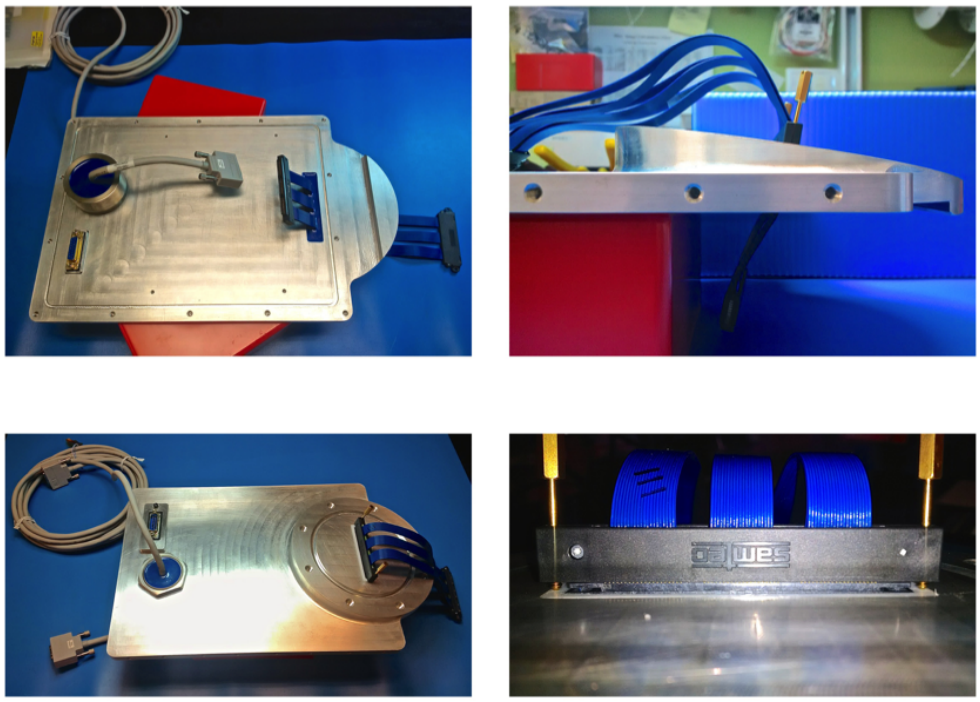

Figure 4.29: Final Nüvü pressure box flange, with sealed-in cables for all required electrical wiring. Top Left: Inside face of the pressure box with the grey video cable and blue SAMTEC ribbon connector for the controller signals. Bottom Left: Outside face of the pressure box and "DN 100" vacuum interface with SAMTEC ribbon cable. Top Right: Side view of the 9° flange tilt. The screw holes for air fin attachments are visible along the side of the flange. The SAMTEC ribbon attachments on top goes to the vacuum tank and below is the pressure box enclosure. Bottom right: SAMTEC ribbon cable connector attached to the potted (sealed) connector inside the flange.

tank to remove most of the water from all surfaces, in combination with charcoal pumping for even more efficient water removal, the tank has now been sufficiently cleaned. The new CT cryocooler has been installed and operates at a constant input power of 100 W, with a coldtip temperature of -170°C, which brings the charcoal getter to -165°C at a reject collar temperature of 33°C. The EMCCD temperature will be kept at about -110°C and is regulated by two omega heaters. We are currently using an off-the-shelf cold ribbon to conduct the heat from the back of the EMCCD clamp to the coldlink. With this setup we are able to cool the charcoal sufficiently to reduce the vacuum pressure to 10^{-7} Torr and hold it below 10^{-6} Torr for much longer (several days) than 24 hours, which is our requirement.

Two flight hardware tests have been performed in Toulouse inside the SIMEON climactic chamber. During this test the vacuum tank and various other flight hardware

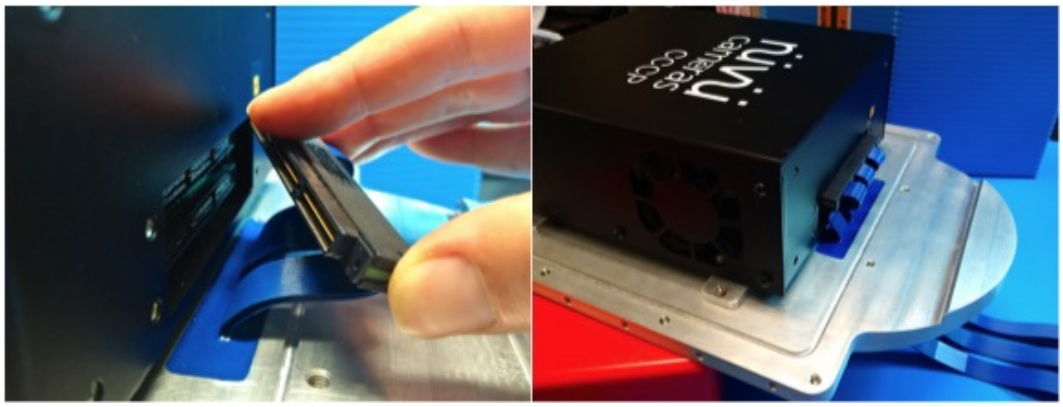

Figure 4.30: Left: Nüvü controller and blue signal cable. One SAMTEC ribbon cable has been sealed directly into the pressure box flange by removing the black plastic shell and potting it directly at the PCB-wire interface. The light green PCB inside the connector is visible on the side the black shell in the left image. Right: The Nüvü box is attached to the pressure box by small stand offs, one is visible here.

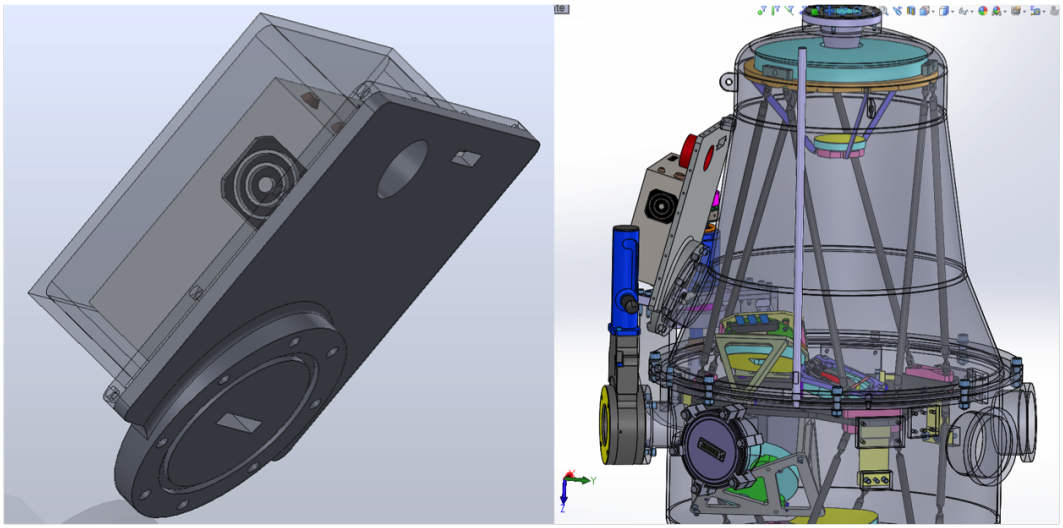

Figure 4.31: Left: SolidWorks model of the Nüvü pressure box with CCCP inside and ribbon feed-through requirements. Right: SolidWorks model of the Nüvü pressure box attached to the flight vacuum tank, to verify all required clearances.

(flight computers, etc.) are enclosed inside a vessel that can mimic flight conditions by adjusting the ambient pressure and temperature. Multiple temperature sensors are attached to various parts of the hardware to check temperature gradients and operating conditions. For this test the equipment was first exposed to atmospheric pressure at 40°C to mimic the gondola's environment while it waits on the launch-pad during midday at Fort Sumner. The pressure and temperature are then reduced, as would be the case during the balloon ascent. The first stage of this test takes

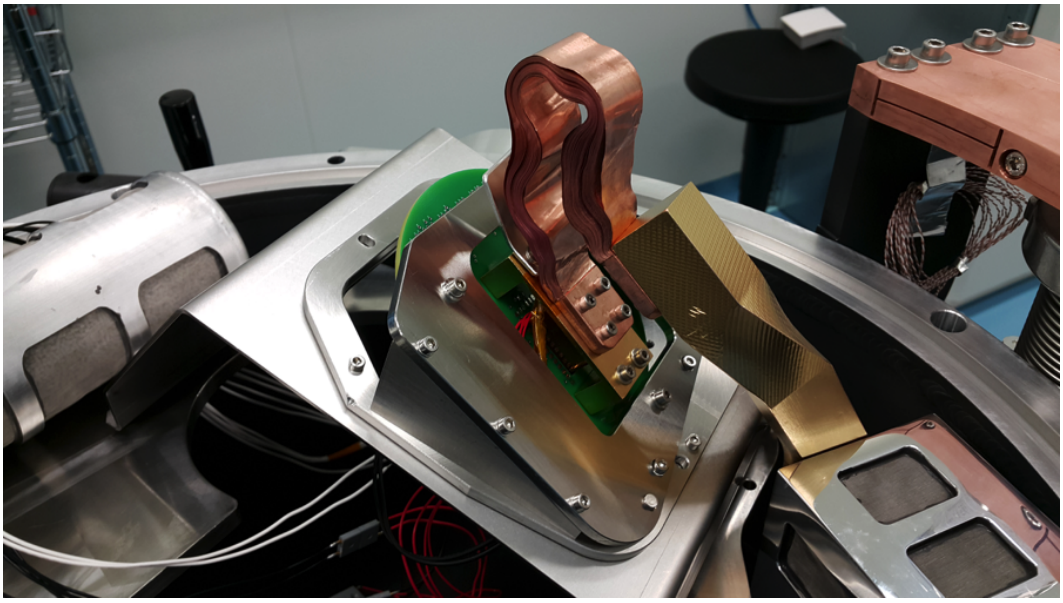

Figure 4.32: A first look at the finished hardware showing the bottom half of the vacuum tank. The cylindrical zeolite getter lies on top of a metal shield to the left. The back side of the mounted green PCB can be seen at the center with a temporary copper cold braid that is attached to the gold plated copper cold clamp at the back of the EMCCD and to the gold plated coldlink on the opposite side. To the right is the polished charcoal getter that is attached to the coldlink, and the cryocooler is mounted to the outside of the tank, partially visible on the right. On top of the cryocooler are copper blades that hold the cooler in place and remove heat from the heat rejector.

the temperature down to 20°C with a pressure of about 0.2 Torr, while during the second stage the temperature drops to 15°C at a pressure of about 0.02 Torr, and the third stage reduces the temperature to -45°C as the experiment reaches the mission ceiling. The test then brings the chamber back to ground conditions as would be the case during descent. The test was repeated twice with similar parameters.

All of my equipment performed with no complications during this test. The cover of the Nüvü controller pressure box stayed below 30°C, which keeps the air inside the pressure box at an acceptable operating temperature for the detector board, and it ran continuously during the test. To ensure safe flight conditions, we will reduce the temperature by adding heat pipes to the enclosure and conduct some of the heat to the vacuum tank. Some adjustment has to be made to the heat removal of the Sunpower CryoTel CT cryocooler. The temperature of the heat rejector was higher than we would like it to be for efficient operation (see Figure 3.12 for nominal performance of the MT cooler). It should stay below 35°C to achieve a high enough lift for our

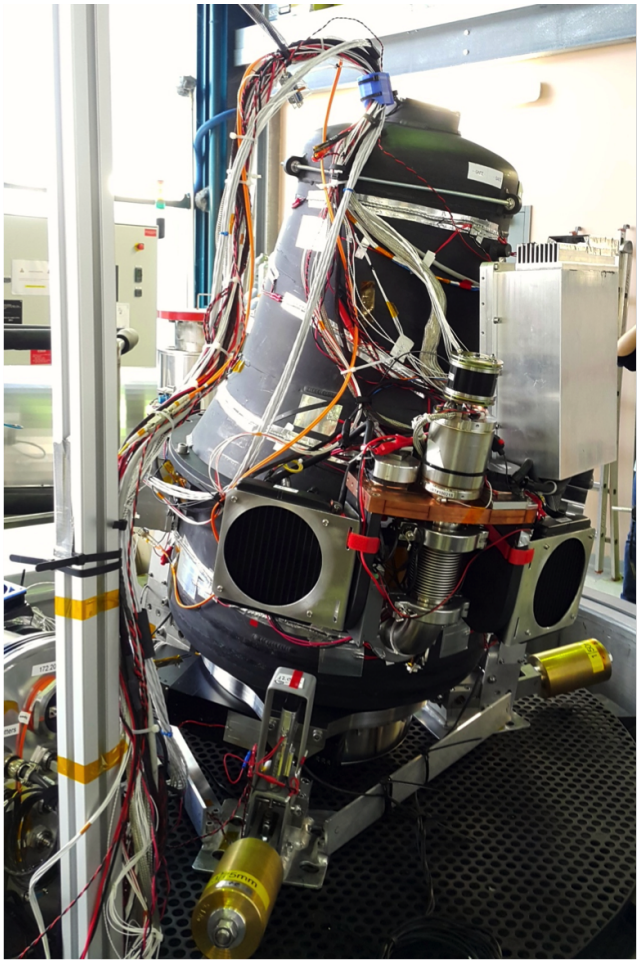

Figure 4.33: SIMEON test images. Vacuum tank on the SIMEON platform. The circular objects are heat radiators. The cryocooler with copper blade attachment to the heat rejector are visible at the center of the image and the aluminum Nüvü controller pressure box is attached to the top of the vacuum tank on the right.

operation requirements. The bottom part of the vacuum tank remained very cold and will require heating during flight, and we will conduct some of the heat from the Nüvü pressure box as well as the cryocooler heat rejector to it, providing more effective heat removal from the reject collar and eliminating the need for additional heaters.

We purchased the version 3 Nüvü controller because it offers several improvements that are very valuable to our experiment. The main aspect of the controller is enhanced (finer) control over the waveforms with a user-friendly software interface and an extensive user manual. Two other key improvements were the driving force behind our upgrade: the version 3 controller hardware is pressure resistant and

radiation hard. Electronic components inside the original hardware were unable to withstand low pressure, but these have been replaced and the controller is now able to operate at high altitude without a pressure container. FIREBall-2 is a pathfinding mission and tests hardware for a future UV satellite. To that end we have developed technologies that will enable NASA's future missions. For a satellite, flight radiation hardness is required, and being able to fly the controller without a pressure vessel is desirable for weight and volume considerations. The new Nüvü controller requires a new PCB interface board, due to changed connectors and signal outputs from the controller. I designed an interface ring that will allow testing of the current flight PCB with the new controller. Testing of the new controller is currently underway, and once it has been fully investigated, verified, and optimized for our application, it will be integrated into the second FIREBall-2 Flight.

Chapter 5

CONCLUSION

I designed and built the Scientific Camera and Cooling System for our balloon experiment FIREBall-2. This instrument investigates the processes that built cosmic structure and the galaxies that trace it. The FIREBall-2 Scientific Camera includes a PCB, which is used as a rigid harness holding the sensor in place. It incorporates RC filters close to the sensor pins to reduce any noise ripple left in the signal trace. I incorporate grounding plates as well as grounding paths around the signal traces to provide a low inductance path for high frequency current, to further quiet any noise in the transmission lines.

The imaging chip is a Delta-Doped CCD201-20 from e2v that has been modified to be used as a 1K by 2K (1024 by 2048 pixels) sensor. The EMCCD is currently under development at JPL's Micro Devices Lab. I operate this next-generation detector at 10 MHz with read-out electronics from Nüvü Camēras. Nüvü's controller for counting photons (CCCP) was chosen to achieve extremely low detector noise by reducing the clock induced charge and is attached to the PCB with a SAMTEC EQCD high-speed coax cable.

A gold plated copper clamp supports the EMCCD that is cooled with a CryoTel MT mechanical cryocooler (Stirling engine) from Sunpower. The initial EMCCD operating temperature was -120°C, with a heat lift of about 7 watts at 30°C reject temperature, which I have verified in the lab. To perfect the camera, the temperature had to be adjusted continuously and two 10-Watt Omega heaters regulate the EMCCD temperature. These heaters are attached to the copper clamp just below the EMCCD for quick and easy tuning of the sensor's operating temperature. During the first phase of testing, the cooler was operated at a constant input cooling power of 45 watts, which corresponds to 73-Watt battery power and consequently keeps the 60 Hz vibrational noise to a very low level. This configuration maintains the coldfinger at -160°C, which places the attached 120mL charcoal getter at an optimal pumping temperature of -155°C . The sorption pump, together with a zeolite desiccant, maintains good vacuum and I use indium at all interfaces to minimize any cooling loss from insufficient surface contact. During subsequent tests the input cooling power was raised to 55 watts, with no obvious effect on the vibrational noise.

A thin 36-layer (76 micron each) flexible copper strap conducts heat from the bottom of the CCD via the copper clamp. A Lakeshore controller (Omega controller for the flight) is used to monitor the temperature across several important surfaces with multiple DT-670 temperature sensors. The detector operates in a vacuum of 10^{-6} Torr, which is achieved in the lab using a Varian V-250 turbo pump as the roughing pump in combination with molecular getters. Once the dewar is valved off, the pressure is maintained by the activated charcoal adsorption and the molecular sieve zeolite. I use a 354 series Ionization Vacuum Gauge to monitor the pressure inside the dewar, and scripts automatically record the ion gauge and temperature diodes with a dedicated Scientific Camera computer (OS: Scientific Linux) through serial (RS-232) connectors and the HyperTerminal Minicom.

The lab camera has been verified and tested extensively and achieves a low noise level of 0.002 events pixel^{-1} frame^{-1}. The flight camera system is currently at LAM in France where it was integrated into the FIREBall-2 gondola and has undergone hardware as well as environmental testing. All hardware functions properly and the flight vacuum tank maintains the required pressure of 10^{-6} Torr over 24 hours. The CryoTel MT cryocooler has been upgraded to the larger CryoTel CT model and we operate the unit at 100-Watt input cooling power. This keeps the coldtip temperature at -170°C and the charcoal getter at -165°C. Additional heat pipes will be used to distribute heat more evenly around the vacuum tank, and heat will be removed from the heat rejector of the cryocooler and used to keep the lower part of the vacuum tank warm enough for operation. The experiment is scheduled to fly out of Fort Sumner in September of 2017.

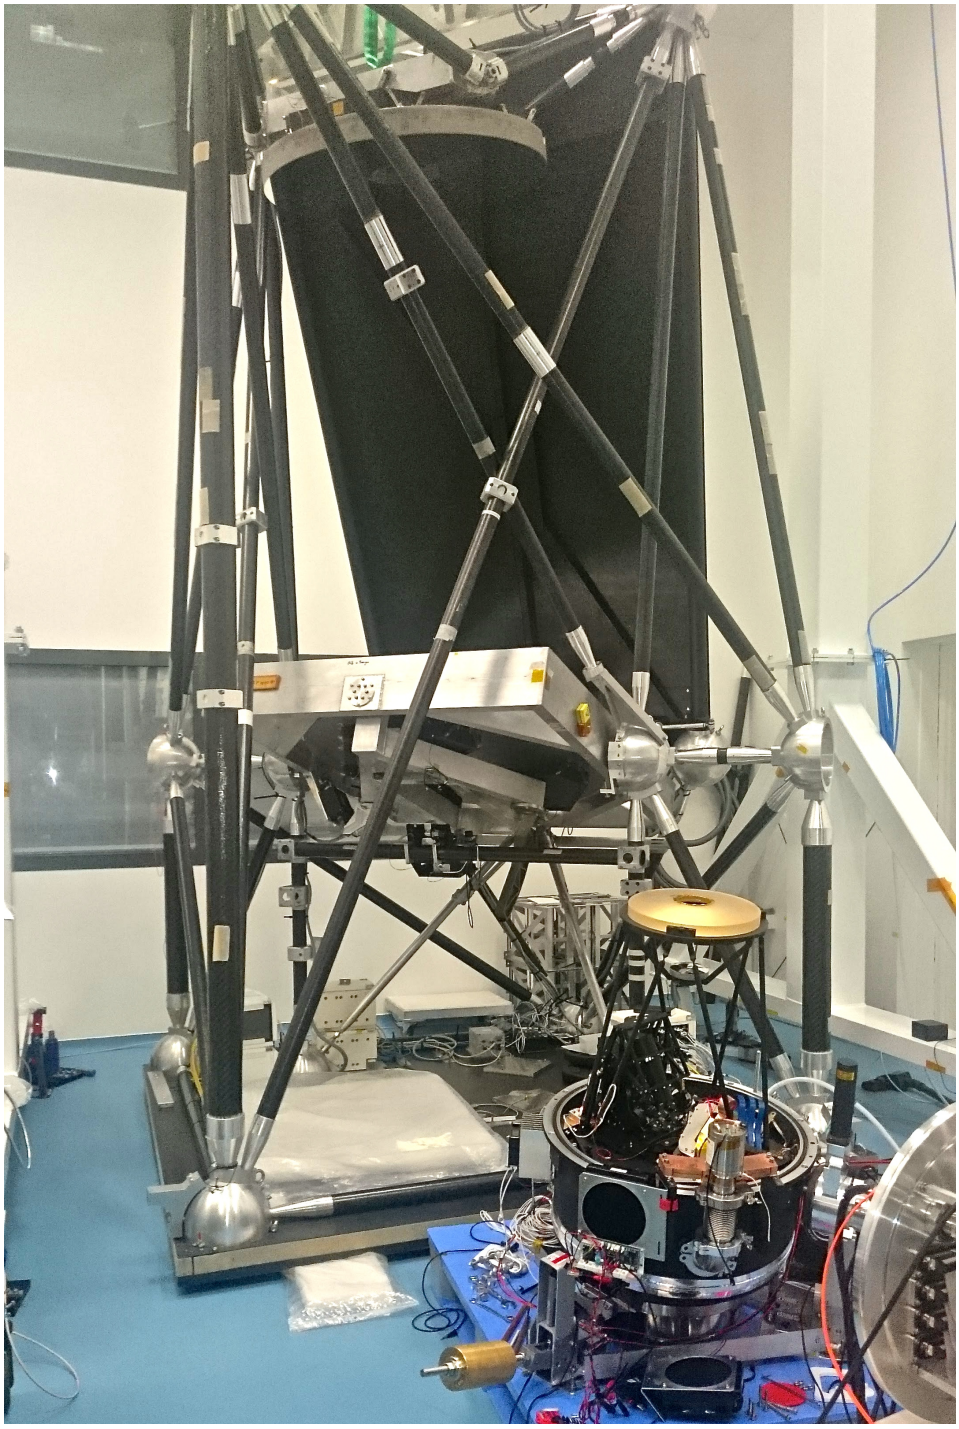

Figure 5.1: This is the status of the FIREBall-2 experiment as I finish writing my thesis in April, 2017. The Balloon Gondola has been assembled in the high bay cleanroom at LAM, France, and the wonderful vacuum tank can be seen in front of it to the right.

BIBLIOGRAPHY

Abe, K et al. (2003). "Measurements of proton, helium and muon spectra at small atmospheric depths with the BESS spectrometer". In: *Physics Letters B* 564.1, pp. 8–20.

Alcazar, J et al. (2000). "Protons in near earth orbit, AMS Collaboration". In: *Phys. Lett. B* 472, p. 215.

Basden, AG, CA Haniff, and CD Mackay (2003). "Photon counting strategies with low-light-level CCDs". In: *Monthly notices of the royal astronomical society* 345.3, pp. 985–991.

Charlton, J. and C. Churchill (2000). "Quasistellar Objects: Intervening Absorption Lines". In: *Encyclopedia of Astronomy and Astrophysics*. Ed. by P. Murdin. DOI: 10.1888/0333750888/2366.

Cherenkov, Pavel A (1958). "Radiation of particles moving at a velocity exceeding that of light, and some of the possibilities for their use in experimental physics". In: *Nobel Lecture, December 11*.

Chief Engineer, NASA Office of the (2010). *Vacuum Seals Design Criteria, NASA Practice No. PD-ED-1223*. https://oce.jpl.nasa.gov/practices/1223.pdf. [Online; accessed 11-Apr-2017].

Commons, Wikipedia (2017). *Cherenkov*. http://sunpowerinc.com/cryocoolers/. [Online; accessed 11-Apr-2017].

Daigle, Olivier, Claude Carignan, et al. (2009). "Extreme faint flux imaging with an EMCCD". In: *Publications of the Astronomical Society of the Pacific* 121.882, p. 866.

Daigle, Olivier, Oleg Djazovski, et al. (2012). "Characterization results of EMCCDs for extreme low-light imaging". In: *SPIE Astronomical Telescopes+ Instrumentation*. International Society for Optics and Photonics, pp. 845303–845303.

Daigle, Olivier, Jean-Luc Gach, et al. (2008). "CCCP: a CCD controller for counting photons". In: *SPIE Astronomical Telescopes+ Instrumentation*. International Society for Optics and Photonics, pp. 70146L–70146L.

Dekel, A et al. (2009). "Cold streams in early massive hot haloes as the main mode of galaxy formation". In: *Nature* 457.7228, pp. 451–454.

Frank, S et al. (2012). "Observable signatures of the low-z circumgalactic and intergalactic media: ultraviolet line emission in simulations". In: *Monthly Notices of the Royal Astronomical Society* 420.2, pp. 1731–1753.

Hamden, Erika T, Frank Greer, et al. (2011). "Ultraviolet antireflection coatings for use in silicon detector design". In: *Applied optics* 50.21, pp. 4180–4188.

Hamden, Erika T, April D Jewell, et al. (2014). "High efficiency CCD detectors at UV wavelengths". In: *SPIE Astronomical Telescopes+ Instrumentation*. International Society for Optics and Photonics, pp. 91442X–91442X.

Hamden, Erika T, Nicole Lingner, et al. (2015). "Noise and dark performance for FIREBall-2 EMCCD delta-doped CCD detector". In: *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 96010O–96010O.

Illustris Project (2017). *Simulation*. `http://www.illustris-project.org/static/illustris/media/illustris_box_z0124_4panel.png`. [Online; accessed 11-Apr-2017].

Janesick, J. R. (2007). *Photon Transfer*. SPIE.

Janesick, James R (2001). *Scientific charge-coupled devices*. Vol. 83. SPIE press.

Jewell, April D et al. (2015). "Detector performance for the FIREBall-2 UV experiment". In: *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 96010N–96010N.

Kindall, SM (1963). *A study of vacuum cryosorption by activated charcoal*. Tech. rep. DTIC Document.

Kleiner, Dane et al. (2016). "Evidence for Hi replenishment in massive galaxies through gas accretion from the cosmic web". In: *Monthly Notices of the Royal Astronomical Society*, stw3328.

Kleinknecht, Konrad (1998). *Detectors for particle radiation*. Cambridge University Press.

Laboratories, IR Labs Infrared (2017). *Scientific and Technical Products Catalog*. `http://www.infraredlaboratories.com/Products.html`. [Online; accessed 11-Apr-2017].

Lane, Neill W (2005). "Commercialization status of free-piston Stirling machines". In: *12th International Stirling Engine Conference, Durham, UK September*.

Lang, Dustin, DW Hogg, K Mierle, et al. (2009). "Blind astrometric calibration of arbitrary astronomical images". In: *Astrometry. net*, pp. 1–55.

Marinacci, Federico et al. (2010). "The mode of gas accretion on to star-forming galaxies". In: *Monthly Notices of the Royal Astronomical Society* 404.3, pp. 1464–1474.

Martin, D Christopher, Daphne Chang, Matt Matuszewski, Patrick Morrissey, Shahin Rahman, Anna Moore, and Charles C Steidel (2014). "Intergalactic medium emission observations with the Cosmic Web Imager. I. The circum-QSO medium of QSO 1549+ 19, and evidence for a filamentary gas inflow". In: *The Astrophysical Journal* 786.2, p. 106.

Martin, D Christopher, Daphne Chang, Matt Matuszewski, Patrick Morrissey, Shahin Rahman, Anna Moore, Charles C Steidel, and Yuichi Matsuda (2014). "Intergalactic Medium Emission Observations with the Cosmic Web Imager. II. Discovery of Extended, Kinematically Linked Emission around SSA22 Lyα Blob 2". In: *The Astrophysical Journal* 786.2, p. 107.

Martin, D Christopher, James Fanson, et al. (2005). "The Galaxy Evolution Explorer: A space ultraviolet survey mission". In: *The Astrophysical Journal Letters* 619.1, p. L1.

Martin, D Christopher, Mateusz Matuszewski, et al. (2015). "A giant protogalactic disk linked to the cosmic web". In: *Nature* 524.7564, pp. 192–195.

Matuszewski, Mateusz Konrad (2011). "The faint intergalactic redshifted emission balloon and the cosmic web imager: two integral field spectrographs designed to study emission from the intergalactic medium". PhD thesis. Citeseer.

Mihos, Chris (2005). *A Primer on Absorption and Optical Depth*. `http://burro.astr.cwru.edu/Academics/Astr221/StarPhys/opticaldepthprimer.html`. [Online; accessed 11-Apr-2017].

Milliard, B, J Donas, and M Laget (1991). "A 40-cm UV (2000 Å) balloon-borne imaging telescope: Results and current work". In: *Advances in Space Research* 11.11, pp. 135–138.

Milliard, Bruno et al. (2010). "FIREBALL: The Faint Intergalactic medium Redshifted Emission Balloon–Overview and 1st Science Flight Results". In: *Proc. of SPIE Vol*. Vol. 7732, pp. 773205–1.

Murray, Neil J et al. (2013). "Multi-level parallel clocking of CCDs for: improving charge transfer efficiency, clearing persistence, clocked anti-blooming, and generating low-noise backgrounds for pumping". In: *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 88600K–88600K.

Nikzad, Shouleh, ME Hoenk, et al. (2013). "High throughput, high yield fabrication of high quantum efficiency backilluminated photon counting, Far UV, UV, and visible detector arrays". In:

Nikzad, Shouleh, Michael E Hoenk, Frank Greer, et al. (2012). "Delta-doped electron-multiplied CCD with absolute quantum efficiency over 50% in the near to far ultraviolet range for single photon counting applications". In: *Applied optics* 51.3, pp. 365–369.

Nikzad, Shouleh, Michael E Hoenk, Paula J Grunthaner, et al. (1994). "Delta-doped CCDs: high QE with long-term stability at UV and visible wavelengths". In: *1994 Symposium on Astronomical Telescopes & Instrumentation for the 21st Century*. International Society for Optics and Photonics, pp. 907–915.

Österberg, Kenneth (2008). *Particle Identification*. `http://www.helsinki.fi/~www_sefo/accelerators/lectures/acc_exp08_8_101208.pdf`. [Online; accessed 11-Apr-2017].

Rauch, Michael and P Murdin (2006). *Lyman Alpha Forest*.

Robbins, Mark Stanford and Benjamin James Hadwen (2003). "The noise performance of electron multiplying charge-coupled devices". In: *IEEE Transactions on Electron Devices* 50.5, pp. 1227–1232.

Ross, Ronald G and Dean L Johnson (2004). "Effect of gravity orientation on the thermal performance of Stirling-type pulse tube cryocoolers". In: *Cryogenics* 44.6, pp. 403–408.

Rudie, Gwen Charlena (2013). "The intergalactic and circumgalactic medium surrounding star-forming galaxies at redshifts 2< z< 3". PhD thesis. California Institute of Technology.

Rudie, Gwen C et al. (2012). "The gaseous environment of high-z galaxies: Precision measurements of neutral hydrogen in the circumgalactic medium of z 2-3 in the Keck Baryonic Structure Survey". In: *The Astrophysical Journal* 750.1, p. 67.

Seals, Parker (2001). *Parker O-ring handbook*.

Shikaze, Yoshiaki et al. (2007). "Measurements of 0.2–20GeV/n cosmic-ray proton and helium spectra from 1997 through 2002 with the BESS spectrometer". In: *Astroparticle Physics* 28.1, pp. 154–167.

Space Sciences, Edge of (2017). *Zero Pressure Balloons*. http://www.eoss.org/. [Online; accessed 11-Apr-2017].

Sunpower (2017a). *Cryocoolers Overview*. https://commons.wikimedia.org/wiki/File:Cherenkov.svg. [Online; accessed 11-Apr-2017].

– (2017b). *CryoTel MT*. http://sunpowerinc.com/cryocoolers/cryotel-family/mt/. [Online; accessed 11-Apr-2017].

– (2017c). *High Performance Free-Piston Stirling Engines*. http://sunpowerinc.com/wp-content/uploads/2014/08/High_Performance_Free-Piston_Stirling_Engines.pdf. [Online; accessed 11-Apr-2017].

Tarzwell, Bob (2010). *PCB 101: Coefficient of Thermal Expansion*. [Wednesday, September 22, 2010, DMR Ltd.]

Technologies, e2v (2017a). *A1A-100013 Issue 1, June 2005*. [Online; accessed 11-Apr-2017].

– (2017b). *EM standard image sensors*. http://www.e2v.com/products/imaging/em-sensors/. [Online; accessed 11-Apr-2017].

– (2017c). *Low-Light Technical Note 4, Dark Signal and Clock-Induced Charge in L3Vision CCD Sensors*. http://www.e2v.com/content/uploads/2015/04/a1a-low-light_tn4_3_v1.pdf. [Online; accessed 11-Apr-2017].

Tumlinson, Jason et al. (2011). "The large, oxygen-rich halos of star-forming galaxies are a major reservoir of galactic metals". In: *Science* 334.6058, pp. 948–952.

Tumlinson, J. et al. (2013). "The COS-Halos Survey: Rationale, Design, and a Census of Circumgalactic Neutral Hydrogen". In: *The Astrophysical Journal* 777, 59, p. 59. DOI: 10.1088/0004-637X/777/1/59. arXiv: 1309.6317.

Tuttle, Sarah E et al. (2008). "The FIREBall fiber-fed UV spectrograph". In: *SPIE Astronomical Telescopes+ Instrumentation*. International Society for Optics and Photonics, 70141T–70141T.

Weitkamp, J (1989). "An introduction to zeolite molecular sieves". In: *Applied Catalysis* 52.3, N25–N26.