HIGH VOLTAGE POWER SUPPLY FOR A SMALL PLASMA SPECTROMETER

By

Quetzal Luebke-Laroque, B.S. Mechanical Engineering

A Thesis Submitted in Partial Fulfillment of Requirements

For the Degree of

Master of Science

in

Electrical Engineering

University of Alaska Fairbanks

August 2023

APPROVED:

Denise Thorsen, Committee Chair
Donald Hampton, Committee Member
Richard Wies, Committee Member
Denise Thorsen, Department Chair
Department of Electrical and Computer Engineering
William Schnabel, Dean
College of Engineering and Mines
Richard Collins, Director
Graduate School

Abstract

This thesis describes the design, manufacturing, and testing of a high voltage power supply for an electron spectrometer to be used on a small spacecraft. The challenge was reducing complexity, size, and power consumption enough to be useful for a CubeSat or a sounding rocket sub-payload.

The power supply has two purposes: i) it produces DC at roughly 2000 V for a microchannel plate detector and ii) it produces exponential voltage sweeps from about +4000 V to below 5 V for an electrostatic analyzer. This power supply uses solid state relays for a charge pump and for producing voltage sweeps with capacitive discharge. The detector power supply is software and hardware regulated to provide tunable output up to 2400 V, with voltage drift and ripple of around ±1%.

The power supply is managed by an MSP430 microcontroller. It includes several automatic, hardware-based, closed-loop controls to stabilize performance across a range of temperatures, input voltages, and component parameters. All digital switching, signal changes, measurements, and monitoring are done in a single control loop. The design includes thermal protection, current limiting, a hardware watchdog, transient voltage suppression, and hardware-based overvoltage protection.

The circuits survived thermal vacuum testing, operating within specifications for several weeks during many cycles of -40 °C to +70 °C. The supply operated a real spectrometer inside an electron beam chamber. The consistent voltage sweeps allowed the spectrometer to accurately distinguish electron beam energies.

This power supply is a successful proof of concept for powering a microchannel plate detector and an electrostatic analyzer at over 4000 V using a single 64 mm diameter circuit board, consuming 0.5 W or less from a small spacecraft.

Table of Contents Page

List of Appendices Page

List of Figures

List of Tables

Acknowledgements

I am thankful to Dr. Denise Thorsen for being a great teacher, and for managing the Space Systems Engineering Lab where I work. She has probably been the single largest positive influence on my life outside of my family. Working in the SSEP lab has been truly life-changing, even long before I started on this thesis.

I have learned from many great student mentors, including Katie Aikens and Matt Pacheco, who introduced me to the lab as a mechanical sophomore. Morgan Johnson and Jordan O'Dell both worked as lab managers before me and taught me many skills. Also, Brian Bieshelt in the SSEP lab has saved me a lot of time by machining the shielding for this power supply.

Craig Pollock and Don Hampton have enabled me to work on the ACAPS (Alaska CubeSat Auroral Plasma Spectrometer) project for many years with a high degree of autonomy and all the material support I have needed. This thesis is only one small part of the fun I have had working for Craig and Don. I am thankful to Craig for bringing ACAPS into the SSEP lab and mentoring me for 5 years now.

Dr. Richard Wies's Power Electronics Design course taught me many of the basic ideas which I used in this thesis, and it gave me useful hands-on experience with power supplies.

Dr. Hui Zhang's Space Physics course was a great survey of concepts used in my thesis. I honestly was not qualified to take it, and the math still frightens me, but I learned a lot of concepts relevant to this thesis and ACAPS overall.

I am thankful for scholarships and fellowships from the University of Alaska, the Alaska Space Grant Program, and UA Scholars/Alaska Performance Scholarship. All this funding has given me financial freedom during college, and now that freedom lets me live a flexible debt-free lifestyle. Even before I graduated from West Valley High School, the university was supporting me through the Alaska State High School Science Symposium.

My family has encouraged my interest in technology from a very young age. I am particularly thankful for my father and grandfather for encouraging me to build model railroads, do science experiments, build a 3D printer, work on cars, and eventually build my own house. All those experiences designing and building things in the real world surely developed skills I used for this thesis.

I am thankful for my friends, Keegan, Cole, Jared, Tane, and Greg for putting up with my crazy projects and supporting me through so many years. Some of us have known each other for over ¾ of our lives at this point.

Chapter 1. Introduction

This thesis describes the design and testing of a High Voltage Power Supply (HVPS) for the Alaska CubeSat Auroral Plasma Spectrometer (ACAPS). ACAPS is a project at the University of Alaska Fairbanks' Space Systems Engineering Lab (SSEP) and the Geophysical Institute. Early concept design started in 2018, and since 2020 it has been funded under NASA Heliophysics (H-TIDeS). I have done mechanical design and electron beam testing with ACAPS, particularly focusing on additive manufacturing (3D printing) with conductive plastics.

Explanation of the name "ACAPS":

Alaska: Developed at the University of Alaska Fairbanks.
CubeSat: A standardized type of very small orbiting spacecraft.
Auroral: To be used for studying the space weather related to Aurora ("Northern Lights").
Plasma: Measures charged particles, meaning negative electrons and positive ions.
Spectrometer: An instrument for measuring the energy of particles.

ACAPS is designed to measure electrons in outer space that cause the aurora and other space weather. As an imaging spectrometer, it measures electrons' direction of arrival as well as their energy.

However, this thesis is only about the HVPS for ACAPS. The design of the spectrometer itself is outside the scope of this thesis, but it will be described in this chapter as background on why the HVPS exists.

Why do we care about electrons in space?

Space weather can be disruptive to modern technology, and even destructive. Understanding it can help forecast problems and minimize the impact. Also, studying space weather around Earth can help us understand the physics of plasmas in general. 99% of the universe is plasma, and there are still many scientific uncertainties about how plasma flows over time and space.

How can weather in space be disruptive?

Our modern society directly depends on orbiting satellites to function. However, those satellites are vulnerable to many space weather hazards. The world depends on satellites for:

- Navigation such as GPS and emergency beacons
- Internet and TV service
- Weather observations for forecasting and disaster warnings
- Military satellites for early warning, reconnaissance, and communications
- Commercial photography and radar imagery for agriculture, and industry
- Scientific data which informs political decisions and helps develop new technology.

Thus, understanding and forecasting space weather is important, in the same way that forecasting storms on the ocean is important for the safe passage of ships. Just like weather within the atmosphere, space weather is hard to predict and can be violent. NASA and the National Oceanic and Atmospheric Administration (NOAA) monitor threats caused by space plasma [1].

- High-energy charged particles can damage materials and disrupt electric circuits. [2]
- Unbalanced flows of low energy plasma can cause charge to build up on exposed surfaces, causing destructive electric arcs. This can damage solar panels on spacecraft. [2]
- High densities of plasma in Earth's ionosphere can disrupt communications, occasionally making parts of the radio spectrum useless here on Earth. This is especially a problem at high latitudes.
- Scientific measurements on earth and in space can be disrupted by electrical noise from auroras and other movements of electric charge in space.

In addition to all these dangers for satellites and technology, civilization on Earth is vulnerable to more direct space weather impacts. Any long conductive structure can be damaged by induced current from space weather, especially power lines. As the Aurora clearly shows, Earth is not completely insulated from the storms of outer space. Earth does have a magnetic field that redirects the Solar wind, but it does not provide complete protection. In fact, the Earth's magnetic field creates some new space weather features of its own, which are especially apparent near the poles. Figure 1.1 [3] below shows the magnetic environment around Earth, with green boxes where ACAPS will study:

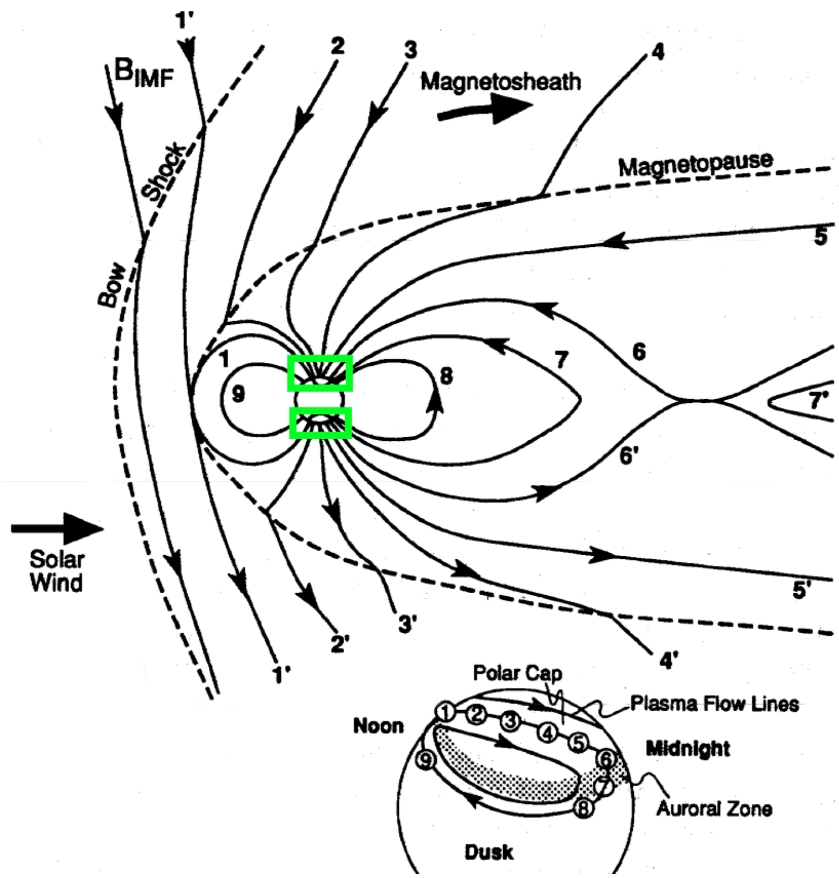

Figure 1.1: Earth's Magnetosphere [3]

As seen in Figure 1.1, the Interplanetary Magnetic Field (IMF) interacts with Earth's magnetic field. The solid black lines are magnetic field lines. The area enclosed by the Magnetopause (dotted line) is called Earth's magnetosphere. The magnetosphere extends tens of thousands of kilometers away from earth, while low-orbiting satellites and most sounding rockets reach only a few hundred kilometers above the earth's surface.

Magnetic fields are a core part of space weather because charged particles follow magnetic fields. Nearly all particles around Earth are closely tied to magnetic field lines, and they spiral around the field lines like a corkscrew. When aurora happens, it is caused by electrons crashing into the upper atmosphere as they spiral down the field lines toward Earth's polar regions.

What do scientists want to learn?

ACAPS is proposed to answer three main questions about the aurora: [4]

1) What is the typical scale size of discrete auroras?
2) What is the nature of the energy distribution of electrons that produce discrete auroras?
3) How do the scale size and energy distribution evolve in time?

These three questions require an instrument to analyze aurora-causing electrons across time and space. They demand very good resolution in time and space because aurora can have surprisingly small and fast-moving features, sometimes too small and fast to be visible to the naked eye. For example, Figure 1.2 below shows auroral arcs only about 1-km wide on the left, but also spirals that are up to 200-km wide on the right!

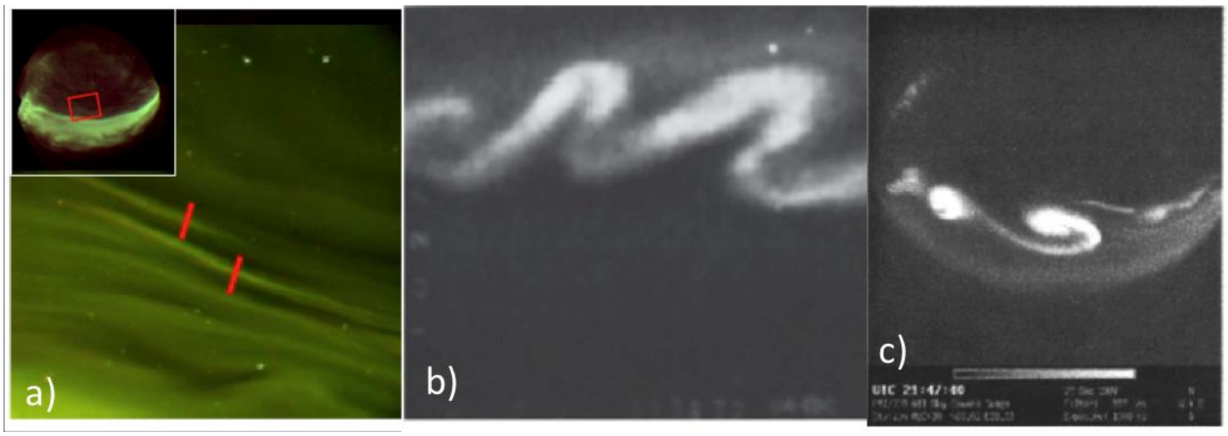

Figure 1.2: Various scales of aurora features, <1km to >100 km [4]

Auroras can "dance" across the sky and change brightness in very short times, less than 1 second. This means the aurora-causing electrons are changing dramatically. Because auroras occur over both large and small scales of time and distance, there is uncertainty about many aspects of aurora formation. Scientists need high-resolution data about electrons to understand the smallest aurora features. This demands time resolution of well below 1 second, and distance resolution well below 1 km.

Why is a spacecraft necessary?

Even though there are many ground-based cameras, spectrometers, radars, lidars, and magnetometers that study aurora, the three questions listed above remain unresolved. One of the major problems of studying space weather is location. A full understanding of plasma cannot be gained by observing at a distance. Direct measurements of particle density, direction, and energy must come from an instrument inside the plasma. Unfortunately, launching satellites into orbit is expensive, and even small science missions generally cost several tens of millions of dollars. A satellite may gather data very precisely and reliably for years, but it can only be in one place at a time. Sounding rockets are somewhat cheaper than orbiting satellites, but they only can gather data for a few minutes before falling back to Earth.

Why use a *small* spacecraft?

First, the cost of launching just a few kilograms into space is much lower. The smallest satellites often cost well under $1 million to design, build, and launch, which is generally less than even a single sub-orbital rocket flight. For example, NASA's CubeSat Launch Initiative and the Air Force Research Laboratory's University Nanosatellite Program both fund universities a few hundred thousand dollars for educational small satellite missions. This reduced cost per satellite allows more risky or exotic technologies to be tested. For example, ACAPS is focused on testing 3D printed conductive plastics on space plasma instruments, which is too risky for larger projects to try. Finally, small size allows many satellites to be launched to operate as a group.

Why use multiple small spacecraft?

While a single expensive satellite might gather high-resolution plasma data along many points in its trajectory, its measurements will always have **ambiguity**. That is because the spacecraft is moving as it takes measurements, so it is difficult to separate changes over time from changes across distance.

For example, if you are driving a car through a rainstorm and suddenly the raindrops stop hitting your windshield, the cause is uncertain. Because you are moving fast, you do not know if it is because the rainstorm has truly ended or you drove out of the rain. The only way to resolve the ambiguity is to compare with another car somewhere else on the road to see when the raindrops stopped for them. Space weather is like this rainstorm example, except satellites travel about 300 times faster than cars. A typical low-earth orbit is 7 to 8 km/sec (about 17,000 mph).

The only way to resolve the time-distance ambiguity with plasma measurements is by taking measurements in multiple places at once. Several large missions have used a pyramid-shaped formation flight of four satellites, such as the European Space Agency's Cluster and NASA's Magnetospheric Multiscale. Those missions are operating today, providing wonderful data about Earth's magnetosphere. However, each of those satellites weighs hundreds of kilograms and are therefore inevitably expensive. While they have answered many questions about space weather, many new questions have emerged. It is not affordable to build a huge network of full-size satellites to monitor all forms of space weather around Earth. There is an incentive to build small, low-cost spacecraft to study space weather, so that dozens or even hundreds can be launched at a reasonable cost.

Multi-point measurement is also relevant for sounding rockets because rockets can deploy sub-payloads. These sub-payloads are each like a small satellite, containing their own science instruments and communications. A single sounding rocket with sub-payloads can briefly create a dense network of many measurements spread across a region of tens of kilometers. This can be useful for providing high-resolution snapshots of space weather in one location. Rocket sub-payloads are small, generally around 100 mm in diameter, which makes them similar size to the smallest CubeSat orbiting satellites. Therefore, these sub-payloads are easily deployable in space in large numbers.

ACAPS is primarily designed to work on an orbiting satellite but can also be adapted for a rocket.

What instruments are used to study electrons in space?

Electrons can be divided into three ranges based on energy with different types of instruments used to measure each. There is no single type of instrument that is useful across the whole energy range. Figure 1.3 shows a spectrum of electron energies with examples of the electron source and the type of instrument commonly used in that range. Note that the ACAPS instrument and auroral electrons are highlighted by bold text boxes because they are most relevant to this thesis.

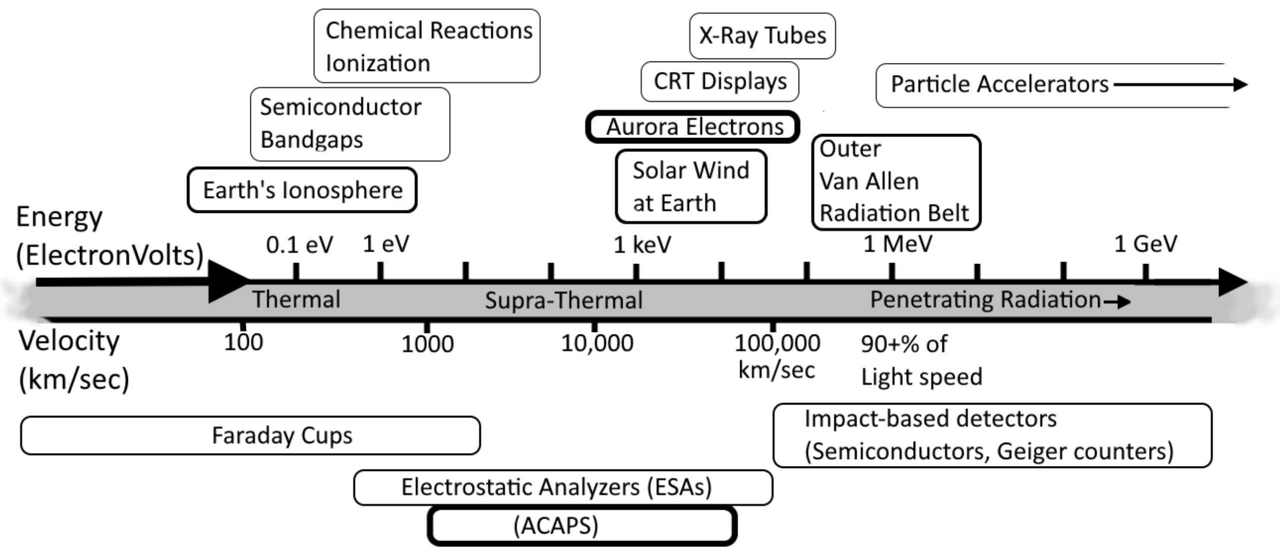

Figure 1.3: Electron Energy Spectrum

The common unit of energy per electron used in Figure 1.3 is the electronvolt (eV). If an electron is attracted toward something that is charged up to +1 V, the electron will gain kinetic energy of 1 eV. The energy of 1 eV is extremely small, about 1.6 x 10^{-19} J. The commonly measured energies of electrons range from around 0.01 eV to over 1,000,000,000 eV. All but the most energetic electrons carry imperceptible amounts of energy and can only be detected with very sensitive electronics after amplification. Three categories of electron energy are listed below:

1) *Low range*: Electrons that are often considered as a bulk flow of plasma instead of being treated like individual particles. They flow in tight spirals around magnetic fields, travelling at velocities governed by the temperature of their plasma. These are often called “thermal electrons”. They are

slow by plasma standards, but still far faster than any speed relevant for humans. Instruments used to detect low energy range electrons include Faraday Cups and Langmuir Probes. These instruments generally allow the electrons to approach from a wide range of angles and are mostly concerned with measuring small energies accurately, but are less concerned with imaging across angles.

2) *Middle range*: Electrons that have more energy than the main thermal group, but still not enough energy to penetrate through significant material. These are still well below the speed of light, but well above thermal velocities. Electrons in this range are sometimes called "supra-thermal", although that is a more specific definition. While these are a small fraction of the total population of electrons in the solar wind, they carry a large fraction of the total energy, so they are very important. Instruments used to detect middle energy range electrons include electrostatic analyzers, magnetic analyzers, and time-of-flight based instruments.

3) *High range*: Electrons that travel at a significant fraction of the speed of light. These are more frequently called "beta rays" in the context of radioactive materials. They carry enough energy to be detected more easily, but also are significantly penetrating and dangerous to life, materials, and electronics. Electrons in this range are difficult to manipulate. Instruments for high energy range electrons rely on directly detecting the results of the electron's impact, such as gas discharge tubes, semiconductor ionization, and scintillators.

ACAPS is an electrostatic analyzer (ESA) for auroral electrons in the middle range of electron energies.

How does an ESA work?

ESAs use attraction of charges to separate electrons of a certain energy level. Electrons are all negatively charged, so they will always be attracted by a positively charged object. A faster electron will require more positive charge to attract it by the same amount as a slower electron. More positive charge on an object means it must have a higher voltage, as shown in Figure 1.4. In this case, if the electrons started in an area of space at 0 V, they will be attracted toward the positively charged boxes. There will be an electric field which the electron flies through, where it feels an electrostatic force.

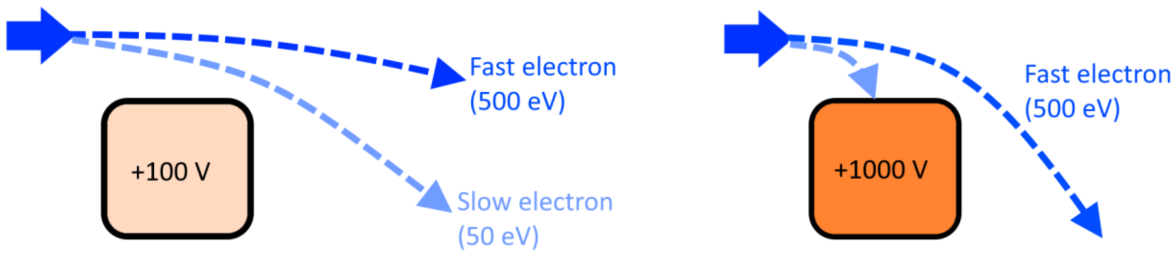

Figure 1.4: Electron Attraction to a positive charge

While being attracted or repelled, an electron's trajectory is very predictable, assuming the voltages attracting it are known. Electrostatic analyzers (ESAs) separate particles based on their energy per charge by taking advantage of the trajectory deflection shown above. An electron of the correct energy can orbit the inside of an ESA. Fortunately, all electrons have the same charge, so an ESA can directly select electrons of a certain energy range.

An ESA like ACAPS extends the idea of trajectory deflection shown in Figure 1.4 to a curvature of 90°. Three trajectories for electrons at various energies are shown in Figure 1.5.

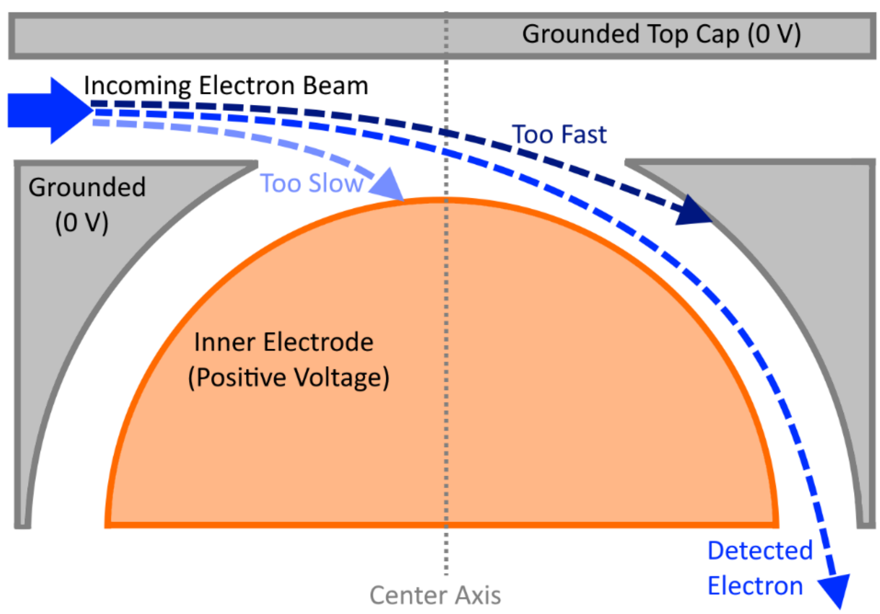

Figure 1.5: Basic Top Hat ESA concept

An electron that reaches the bottom of Figure 5 will be confined to a narrow range of energies. A faster electron would have hit the outer grounded electrode, and anything slower would have hit the inner electrode. That narrow range of accepted energies is defined by the voltage applied to the inner electrode. If that inner electrode voltage is increased, the ESA will select electrons with higher energy, and vice-versa.

The curved channel which the detected electrons travel through can be revolved around the ESA's center axis to create a half-sphere shape. This idea is called a Symmetric Quadrasphere or "Top Hat" ESA and was developed in the 1980s [5] [6].The Top Hat style of ESA is useful because it can "see" in a field of view that extends all around the instrument, like a panorama picture. This ability to collect and distinguish particles from many directions at the same time while also measuring their energy is very powerful. Figure 1.6 below shows the latest ACAPS design as a cross-section view. The overall diameter is 63 mm. This is a screenshot from a SolidWorks CAD model of the ESA, stretched approximately to scale if printed on standard size paper.

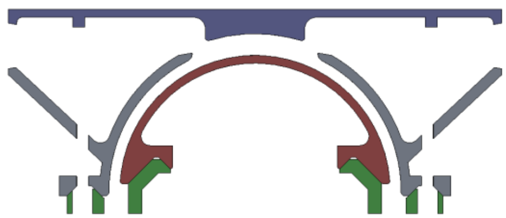

Figure 1.6: ACAPS Section View, to-scale

Figure 1.7 below shows a cutaway view of the same ACAPS ESA 3D model.

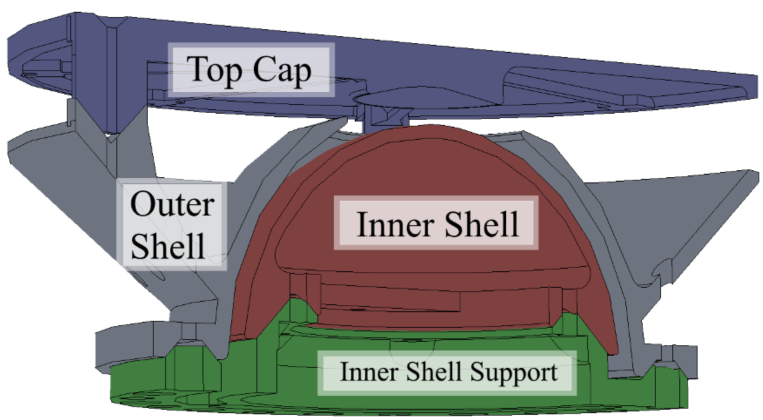

Figure 1.7: ACAPS cross-section, labeled

The red part labeled "Inner Shell" is where positive high voltage is applied. It is this positive voltage which determines which energy of electrons the ESA will detect at any given moment. The Top Cap and Outer Shell are grounded on ACAPS, while the Inner Shell Support is a plastic insulator.

Why can't an ESA be used across the whole energy spectrum?

ESAs don't work well at **higher energies** (above about 50 keV), because it becomes difficult to collect the electrons with any reasonable voltage. The required voltage for an ESA is directly related to the desired electron's energy, which can be both helpful and harmful. For example, a typical ESA design would require 30 - 60 kV to capture 300 keV electrons. Such absurd voltages would require very large spacing between electrodes and wires to prevent arcing and would require a lot of electrical power to operate. Also, those high energy electrons penetrate through materials significantly, instead of just being absorbed at the surface, meaning they can be detected by other means. At high energies, semiconductor-based detectors are dominant. They rely on the energetic particle's ability to create a trail of ionization as it penetrates a semiconductor. Semiconductor detectors (and the related older style of Geiger tube detectors) don't have to influence the electron's trajectory at all, unlike an ESA, so they can detect arbitrarily high energies.

ESAs also don't work well at very **low energies** (such as 0.1 eV), because the electrons' trajectories are so difficult to control. Even a weak magnetic field (such as that carried by the solar wind) can be enough to noticeably curve the electron trajectories. Also, it becomes difficult to work with the very low voltages involved. Measurements will be disrupted by tiny amounts of electric field leakage or by changes in the spacecraft's ground voltage [2]. Equalizing a spacecraft's ground voltage with respect to the surrounding plasma is possible, but demands significant resources (power and space) because it requires actively emitting a beam of particles. Small spacecraft are unlikely to have such equipment.

To minimize the magnetic curvature problem, some exotic ESAs have been built, such as Laminated ESAs [7] or simply ESAs with very small overall size [8]. These instruments sacrifice performance at higher energies, though. The low energy range of electrons are more generally measured by Faraday Cups, Retarding Potential Analyzers, or Langmuir Probes. Those types of devices all rely on attraction and repulsion of plasma based on voltage, but they don't have the confined geometry of an ESA.

What is the high voltage power supply for?

The purpose of this high voltage power supply is:

1) Produce voltage sweeps for the ESA, covering a wide range of electron energies.
2) Supply a DC high voltage to the detector Micro Channel Plates (MCPs) so the electrons can be amplified and detected.

ESA Voltage Sweeping

Covering a wide range of electron energies requires a variable voltage on the inner electrode. The standard way to do this is by rapidly sweeping through many different voltages to cover the entire desired electron energy range. Voltage sweeping is important because the ESA is quite selective. At any single voltage, the ESA can only detect electrons within about a ±10% energy range. Electrons outside that range will have been filtered out by hitting the walls of the ESA, as seen in Figure 1.5. Voltage sweeping allows the ESA to sample a wide range of energies in a short period of time, allowing scientists to form an electron energy spectrum from every sweep.

For ACAPS, the ratio between electron energy in eV and ESA voltage in volts is approximately 5-7. This ratio is based on the size of the gap between the grounded electrode and the positive electrode. An instrument with a narrower gap would have a higher ratio. The target range of energies for ACAPS is about 10 eV to 20,000 eV. Assuming an energy-voltage ratio of 5 eV/V, the ESA power supply needs to generate a sweep of voltages from roughly 2 to 4000 V. Figure 1.8 shows the conceptual decreasing voltage sweep that was eventually designed for in this work.

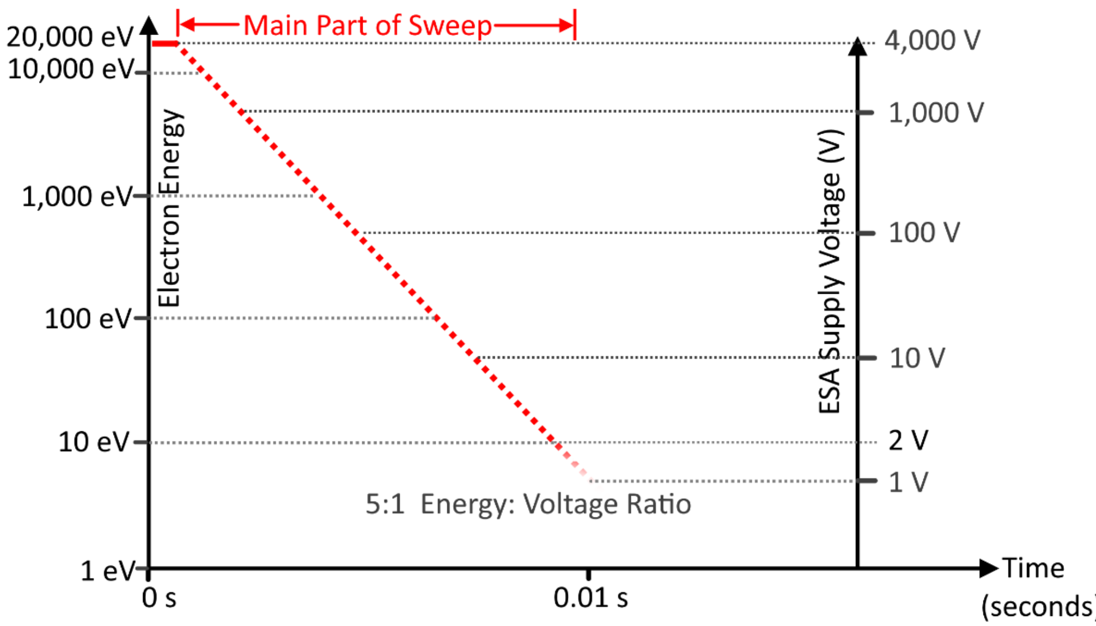

Figure 1.8: ACAPS voltage sweep concept

The sweep period of 0.01 seconds means there will be 100 sweeps repeated per second, providing very good time resolution. Note that the electron energy and ESA voltage are parallel axes in Figure 1.8, both using logarithmic scales (not linear).

Electrons of interest to scientists appear at both ends of this voltage sweep, so high proportional energy resolution is important at 100 eV as well as 10,000 eV. Also, the ESA's intrinsic energy resolution is proportional, not an absolute number of eV. Both of these warrant an exponential voltage sweep instead of linear, so that the proportional rate of change of voltage remains constant. For example, this sweep shown in Figure 1.8 has the voltage always decreasing by a factor of 2 every 0.84 ms.

To acquire an evenly sampled spectrum across these energies, the voltage sweep should be exponential, not linear. If sampling is done at energy steps, an evenly sampled spectrum means each step's voltage changes exponentially, not linearly. A linear sweep would spend too little time sampling the lowest energies, while an exponential sweep samples each order of magnitude equally. For example, a linear sweep would spend equal time measuring between 10 eV and 100 eV as it would between 10,010 eV and 10,100 eV. That means it would spend only 0.5% of its time sampling the lowest energy order of magnitude, and 90% of its time sweeping through the highest energies. However, an exponential sweep spends equal time sampling between 10 eV and 100 eV as it does between 1000 eV and 10,000 eV. Every ESA found in the literature uses some kind of exponential voltage sweep, but there are some differences in how the sweep is generated.

See Chapter 3 for the circuit which generates the voltage sweep. In addition to producing sweeps for the ESA, there needs to be a high voltage supply for the electron detector. Before describing how the detector power supply works, the detector itself must be explained.

How are the electrons detected?

This is an important question because the ESA does not detect electrons, it only filters them by energy. The detector for ACAPS uses **Micro-Channel Plates (MCPs)** to amplify the charge of individual electrons. MCPs are complicated but necessary, because transistor-based circuits are not able to directly measure the charge of a single electron. Noise and leakage currents are far too large. An impressively low transistor leakage current, such as 1 femtoamp (10^{-15} A), is still about 6000 electron charges per second, enough to wash out any individual electron's signal. Even with hypothetical leakage-free electronics, a single electron worth of charge would be lost in thermal noise of any amplifier circuit. MCPs are used in all kinds of detectors because of their amazing single-particle sensitivity. MCPs are such useful devices because they can amplify a single electron with minimal noise. However, they must be operated in a very clean vacuum, such as outer space.

How does an MCP work?

An MCP is a specially coated thin glass plate with many tiny holes going through it. The holes are almost microscopic, (about 10-25 μm in diameter), and each hole can be the site of a chain reaction if it is triggered. An MCP can be triggered by almost any energetic particle, such as X-rays or ultraviolet light, neutral atoms, ions, or, in the case of ACAPS, electrons. Figure 1.9 below shows 4 stages of the electron chain reaction in a single MCP channel.

After surviving a journey from space, into the satellite, and through the curved channel in the ESA, an electron will impact the wall of a hole in the MCP. This causes several new "secondary" electrons to be emitted. The high voltage applied on the bottom side of the MCP attracts these new electrons and accelerates them down toward the output. Those new electrons will hit the wall somewhere else, releasing even more electrons in a very fast chain reaction. The channels are 40-60 times longer than their diameters, so a sequence of impacts and emissions can multiply many times.

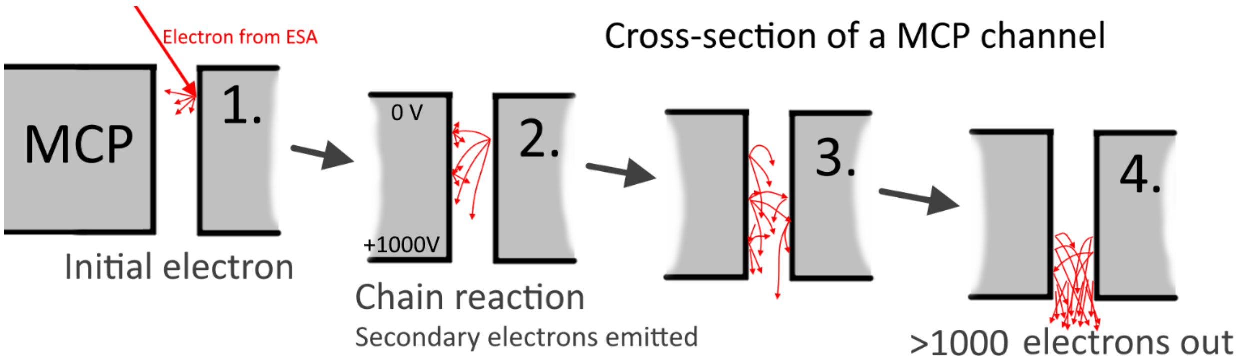

Figure 1.9: MCP chain reaction

A single microchannel plate generally has an electron gain between 1,000 and 10,000. To amplify even more, a second MCP is usually added, turning thousands of electrons into millions. At that level of amplification, each electron entering the detector will produce an output pulse of around 10,000,000 electrons, or 10^{-12} Coulombs, which can be loosely expressed as “1 milliamp for 1 nanosecond”. That is still a tiny charge but is certainly measurable by modern silicon electronics.

How is an MCP powered?

To fuel the chain reaction of electron emissions, about 1000 V is applied from the input surface to the output surface of the MCP plate. This positive voltage on the output attracts electrons, accelerating them through the plate. Without a high voltage supplying a continuous flow of new electrons, the MCP would not be able to output more electrons than it takes in.

The high voltage supply needs to provide a consistent voltage across each MCP in series, meaning about 2000 V total for the two plate stack planned on ACAPS. It is also helpful to bias the input of the first MCP slightly positive, to help attract electrons toward the channels, adding another 100 V to the total. Finally, the detector downstream of the second MCP may be biased slightly, again to help attract electrons toward it.

Because an operating MCP emits far more electrons than it receives from the ESA, the power supply must provide a small current to replenish the electrons which are emitted out of the channels. The current load will vary based on how many electrons ACAPS is detecting at any moment. Even if an MCP is not detecting any particles, it conducts a steady amount of “dark current”, usually a few µA per square centimeter of MCP area. This dark current also varies with temperature and varies dramatically between

MCPs, even from the same manufacturer. MCPs sometimes require gradual voltage adjustment over time as they age. Usually, voltage must be increased as the MCP's electron gain gradually declines.

The power supply's output voltage and current must be adaptable to meet the needs of various MCPs, because the specific MCP stack detector for ACAPS has not been designed yet. Chapters 2 and 4 will discuss the MCP supply in greater detail.

What is special about this design situation?

The main constraints of the HVPS for ACAPS are as follows:

1) It must fit into the size and mass limits of a small satellite.
2) It must use very little power to work on a small satellite.
3) It must operate in the harsh environment of space.
4) It must produce stable, consistent voltages to allow precise measurements with the ESA and detector.
5) It must have minimal impact on the spacecraft's overall design, so that it can be realistically integrated into an overall space science mission.

This list has been turned into Design Requirements in Chapter 2. Chapter 2 discusses systems engineering for the power supply as a whole. Chapter 3 explains the sweeping supply, and Chapter 4 explains the detector supply. Chapter 5 explains the microcontroller and other low voltage electronics. Chapter 6 explains the protection circuits. Chapter 7 is dedicated to circuit board design, including mechanical/thermal design, and Chapter 8 is testing results. Chapter 9 is unfinished ideas, and Chapter 10 is the conclusion.

Chapter 2. Overall Design Basics

At the start of this design, it was obvious that no “off-the-shelf” option existed. As far as could be determined from the available literature and product searches, HVPS for space are normally custom designed for each mission. I found no mass-produced and publicly sold space-quality HVPS. In the future, this may change as the commercialization of small satellites continues, but as of 2021, custom design was necessary. However, this custom HVPS design does take advantage of many off-the-shelf components

Objective

The power supply will generate, manage, and measure voltages for both the ESA and the detector of ACAPS. Low voltage power management, data handling, and communications will be handled by other circuits.

Figure 2.1 below shows an overall block diagram of the final design for this thesis.

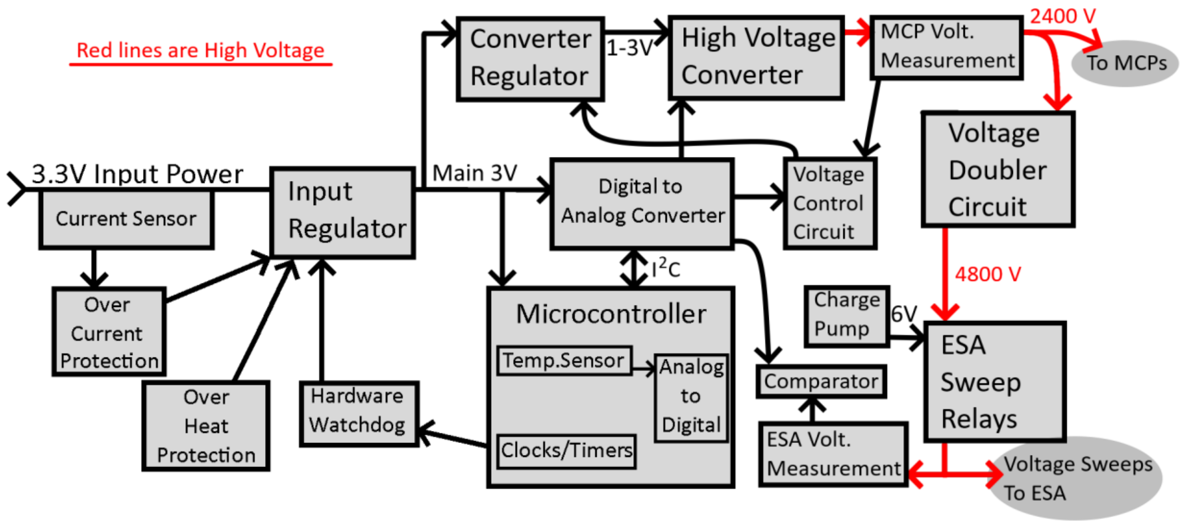

Figure 2.1: Overview Block Diagram

In later chapters, there will be detailed schematics for each rectangular box in Figure 2.1. Chapters 3-6 will each discuss portions of this diagram, which will be highlighted at the beginning of the chapters.

Requirements

Table 2.1 on the following page lists the requirements and their rationales for this HVPS. Table 10.1 in Chapter 10 summarizes whether each requirement was met based on testing results.

Before discussing all the engineering decisions made to meet the design requirements, some additional background on the HVPS design features will be discussed in this chapter.

Table 2.1: Design Requirements List

#	Requirement	Rationale
1	Supply an exponential voltage sweep across the range of < 10 to > 5000 V, at least 20 times per second, to a load of 10 to 50 pF. Sweep must be repeatable within ±5 % of nominal voltage at all points throughout the sweep.	This is for the ESA. 20 sweeps per second (20 Hz) would already be better than many instruments, giving a resolution of 400 meters on an 8 km/sec spacecraft. 5000 V allows detecting about 30 keV electrons, which is necessary for auroral science. 10 to 50 pF is an estimated range of reasonable ESA capacitances, including the cables. Large ESAs with closely spaced electrodes have higher capacitances. The ESA has about ±10% energy bandwidth, so ±5% is chosen so that the HVPS is not a large contributor to overall energy uncertainty.
2	Supply a DC voltage which is tunable between 1500 V and 2500 V, at up to 40 µA. Must be stable within ±1% voltage across all operating conditions and all current loads.	This is for the microchannel plate detector. MCPs nominally take 1000 V per plate (2000 V total), but some adjustment room is needed. 40 µA implies a minimum allowed resistance of about 30 MΩ per MCP. Measuring the current is also required. 1% voltage stability is important because MCP gain varies dramatically with small changes in voltage. 1% was suggested in a design review.
3	Use only standard voltages of 3.3 V and 5 V from the host CubeSat. Consume less than 1.5 W average power during operation.	3.3 V and 5 V are standard voltages. 0.5 W is the target power consumption. 1.5 W would be demanding for a small satellite, but still potentially usable if the ESA only operates for short periods of time.
4	Do not apply any mechanical demands on the host spacecraft, beyond those already imposed by the ACAPS spectrometer (ESA).	This allows the power supply to be easily included on any spacecraft which can support the rest of ACAPS.
5	Survive a range of -60 °C to +100 °C. Operate within specifications from -40 °C to +60 °C.	These are reasonable ranges of temperatures experienced in a CubeSat's "Tuna Can" volume.
6	Survive the vibration environment of a rocket launch, based on the NASA General Environmental Verification Standard (GEVS).	All spacecraft must survive launch.
7	Recover safely from high voltage breakdowns at the outputs.	This means it must not overheat, be damaged, and produce unexpected outputs when arcs occur.

Basic Parts of a Power Supply

Most HVPSs contain these components:

1. Transformers
2. Charge Pumps
3. Voltage Dividers
4. Capacitors

This HVPS includes all four components in some form. Each component is briefly explained in the following sections before going into more detail in later chapters.

Transformers

A common example of a high voltage transformer is the ignition coils for internal combustion engines. Those “coils” are transformers, which take in a pulse of 12 V or 24 V at relatively high current, and output 20,000 volts or more at low current, which is sent to a spark plug. Electric fly swatters use a similar simple transformer circuit. The technology of producing sparks with a transformer has been commercialized with ignition coils for well over 100 years. Clearly, creating a few thousand volts or more is not inherently a design challenge, considering that the parts are cheap enough to mass-produce $13 bug zappers. The challenge is producing that high voltage safely, efficiently, consistently, and with small size and mass.

Transformers require alternating current (AC) to work, ideally a smooth sine wave, but at least a square wave. This is an inconvenience on a spacecraft where the power source is a DC battery. Before a transformer can be used, DC must be “inverted” to AC. Thankfully, modern transistors can achieve this by switching current through the transformer on and off very quickly. That idea is the foundation of nearly all power supplies today, which is why they are called “switching” power supplies. Switching power supplies are useful at both low and high voltages.

Unfortunately, to achieve a high voltage output from a low input, a transformer needs many loops of wire wrapped around a magnetic core. Transformers are often the largest and heaviest part of a power supply because they need a magnetic core, many loops of wire, and insulation on all of those windings. High voltage can arc through a thin layer of insulation, causing another design challenge.

If high voltage DC is desired, the transformer’s output must be “rectified”, to convert it back to DC. This usually involves rectifying diodes and is not very difficult. However, a small amount of “ripple” will remain in the output after rectification. In other words, a rectifier’s output is not quite a perfect DC signal. Improving the smoothness of the DC output is a major challenge for miniaturizing a power supply. Minimizing ripple requires physically large capacitors, even more so at higher voltages.

Also, the AC output of a transformer can be fed into a charge pump before it is rectified.

Charge Pumps

Charge pumps convert AC to DC just like any rectifier, but they also can increase the voltage in the process. There are many varieties, but all charge pumps use multiple capacitors to momentarily store charge, and some sort of switch or diode to asymmetrically conduct current. The simplest charge pump is the diode-based voltage doubler, shown in Figure 2.2 on the following page.

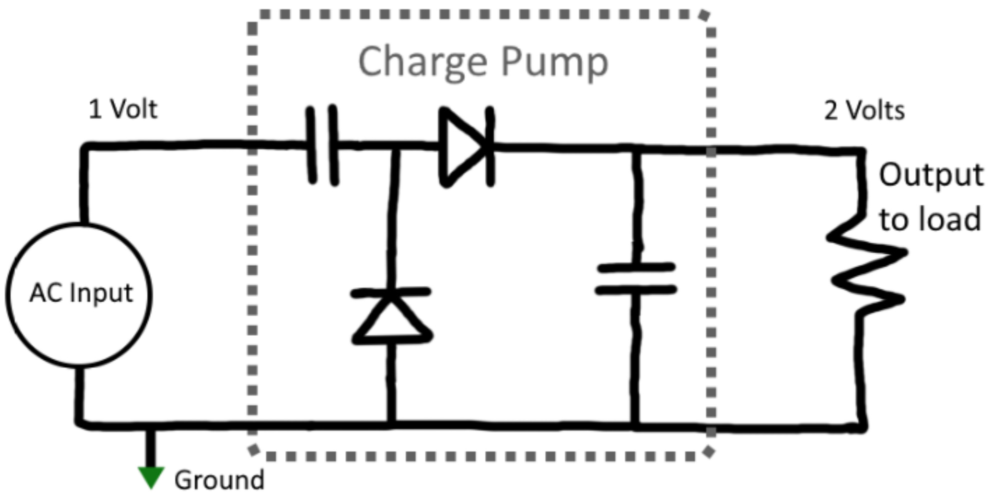

Figure 2.2: Basic Charge Pump Voltage Doubler Schematic

Charge pumps work because a capacitor stores a relative voltage between its two pins, NOT necessarily an absolute voltage relative to ground. This concept is sometimes known as a “flying capacitor”. Charge pumps effectively rearrange capacitances in a circuit. They charge capacitors in parallel but discharge them in series. Most charge pumps are designed using the circuit shown in Figure 2.2. The pattern of diodes and capacitors can repeat indefinitely. Even in the early 1900s, charge pumps with many stages reached millions of volts for particle accelerators [9].

For applications that demand very little current, a charge pump is exceptionally effective, because the capacitors can be small and efficient, even at very low power levels. For example, the TC1240A from Microchip which I used in this thesis is roughly the size of a green pea but can supply up to 0.2 W and has efficiency above 86% for any load [10]. Charge pumps thrive in the low power range, where other types of power converters struggle. Transformer-based power supplies can work well at higher power levels, but often have poor efficiency below around 1 W, and they are likely to be larger. Charge pumps are especially well suited to high voltage, low current supplies. At high voltage, the losses from diodes are minimized. For example, a voltage multiplier starting with 5 V would lose 20% of its voltage just from two 0.5 V voltage drops of diodes, but a multiplier starting at 1000 V only loses 0.1%.

As an alternative to diodes, charge pumps can also be actively controlled with switches (relays or transistors). This allows the charge pump to take in DC instead of requiring AC input like Figure 2.2 above. This type of design is discussed in more detail in Chapter 4.

Voltage Dividers

After a high voltage has been generated, it needs to be measured. A voltage divider is usually necessary to bring it back down to a safe voltage for measurement and control. Whether the measurement is to be analog or digital, generally a signal of just a few volts is desirable. The required voltage divider ratio in a HVPS is often 1000 or more. Figure 2.3 below shows a simple resistive voltage divider with a 3000:1 voltage ratio and several useful features.

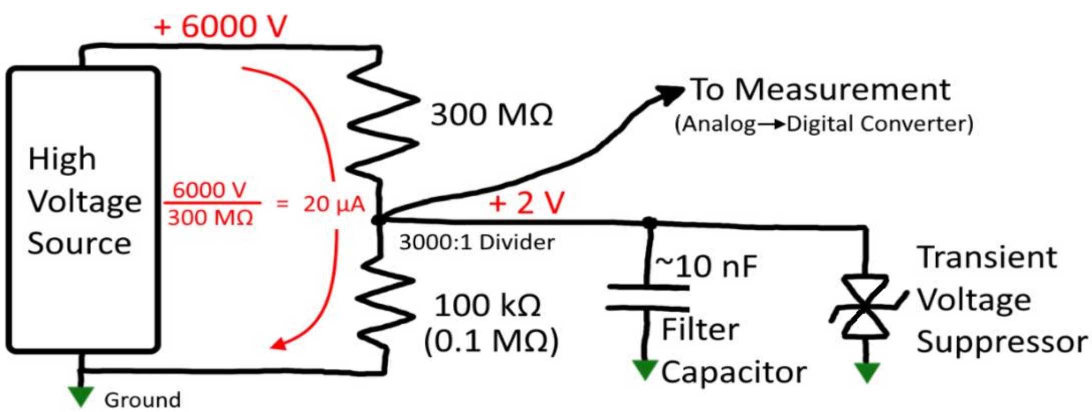

Figure 2.3: Generic Voltage Divider Schematic

While they are somewhat wasteful of power, simple voltage dividers are common on HVPSs because of their simplicity and consistency. However, designing a good resistor-based voltage divider is more complicated than just selecting two resistors.

It is necessary to protect the measurement electronics on the low voltage side of the divider from accidental voltage spikes. Voltage spikes can couple through the divider's high side resistance to the output, because all resistors have a small capacitance, around 10^{-12} or 10^{-13} F. Suppressing spikes requires matching capacitor impedances of the divider, not just resistances. A simple way to achieve this is by adding a much larger capacitance such as 10^{-8} F on the low voltage side of the divider, which also results in some filtering of the output voltage.

For additional protection beyond just a capacitor, a transient-voltage-suppressor (TVS) should be included. TVS devices are usually diodes specially designed to carry short, sudden currents. This is like protecting against static electricity in any ordinary circuit. I used large numbers of TVS diodes in this thesis to protect nearly every vulnerable signal.

The electronics used for measuring the divider's output must have very high resistance compared to the divider itself. Otherwise, the measurement resistance in parallel may noticeably alter the divider ratio. For example, Figure 2.3 has 100 kΩ on the low side, so the measurement hardware must be many MΩ to be considered insignificant. Connecting a 10-MΩ measurement device (such as an oscilloscope) would alter the divider ratio by 1%. Providing an amplifier at the voltage divider's output can "buffer" the output, preventing a measurement device from altering the overall resistance. Buffered outputs were used for all measurements in this thesis.

Often, many individual resistors need to be stacked in series to make sure they withstand the applied voltage. For example, the resistors used in this thesis could only handle about 1000 V each. Some resistors behave nonlinearly at voltages well below where their failure point. This nonlinear behavior must be avoided for the divider's resistance ratio to stay consistent.

All resistors change resistance with temperature. If the two parts of a voltage divider do not have perfectly matched temperature coefficients (measured in parts per million, ppm, per degree C change), the divider ratio will be inconsistent. Thin film resistors are good, but still have temperature coefficients around 25 ppm/°C. This means a voltage divider's ratio with two resistors might vary by 0.5% across a 100 °C range of temperatures experienced in space. Such variation is still within the ±1% needed to meet Requirement 2, but it is enough variation to be measurable.

Capacitors and Breakdown

All power supplies, namely HVPSs use capacitors for storing charge and filtering high frequency noise. Any capacitor relies on conductive surfaces being close together, with an insulator between them. In ordinary low-voltage capacitors, the required insulating gap can be microscopic. At low voltages, many layers of conductive surface can be stacked like sheets of paper, allowing grain-of-rice sized ceramic capacitors to easily achieve 10^{-5} F of capacitance at around 10 V. Unfortunately, achieving anything close to that at high voltage is very difficult. The trouble is dielectric breakdown (also known as arcing). Any insulating material can be vulnerable to this, air, plastic, glass, and even a perfect vacuum, surprisingly enough. To avoid breakdown in a high voltage capacitor, the insulator gaps must be much larger, and special insulating materials must be used. The rule of thumb used in this work is that insulators become questionable beyond 4000 V per mm of thickness.

One of the largest design challenges was minimizing the amount of capacitance required, since 3000-V or 6000-V capacitors are so large. For example, a 3000-V ceramic capacitor of $4\text{x}10^{-9}$ F is about the size of a US dime coin, despite having over 1,000 times less capacitance than the "grain of rice" sized 10-V 10^{-5} F capacitor.

Not only is breakdown a problem with capacitors, but also with other electrical components. A circuit board itself can be a breakdown concern, either through the board material or across its surface. Along surfaces, a rule of thumb is 1000 V per mm of distance for a spark to travel. Strictly applied, this rule would mean a 6000-V wire on a circuit board should have a 6-mm radius circle of empty space around it. This need for spacing between components quickly consumes huge areas of a board. This will be discussed more in Chapter 7 about the board design. See Figure 7.10 in Chapter 7 for a real example of spacings between voltages on a circuit board.

Prior Space High Voltage Supplies

The community of HVPS designers for space applications is small and sometimes insular. Even basic information about HVPS for publicly funded (NASA) missions is difficult to find. Despite science data being openly available for most spacecraft, the design of their power supplies is usually not available. Often, there are many papers or even an entire book about the design of a spacecraft and its science instruments, but essentially no information about the power supplies supporting those instruments. I believe this is because HVPS design for space is done by a few competing groups of engineers at major institutions or contractors, who have little incentive to share their work. Still, a general impression of the state of the art could be developed for the work in this thesis.

Where to put the power supply?

The normal place for any power supply would be inside the main part of a spacecraft, on a circuit board within a stack of boards. Most likely this would be below the MCP detector. The available board area is about 90 x 90 mm on a “CubeSat” small satellite, or about a 70-mm circle for a sounding rocket. That is plenty of board area, but unfortunately space within the main spacecraft body is valuable for many other needs, such as batteries, radios, computers, and other science instruments. The available area is different on a satellite versus a sounding rocket sub-payload. Positioning the power supply below the detector was the default option, but a more efficient design was desired.

Power supplies are also sometimes placed in the hollow space within the ESA’s inner electrode (see Figure 1.7). This may be viable with larger ESAs, but only about a 30-mm diameter circle is available for ACAPS. Therefore, such a small board area would be difficult for ACAPS. In hindsight, it may have been possible, but the design challenge was beyond my abilities at the start of this thesis.

Instead, a third unconventional option of placing the power supply on top of the ESA was used, meaning it is extending farther out of the spacecraft. This was an appealing idea to save volume within the main body of the spacecraft. It is also a flexible location, because it does not matter what type of spacecraft is hosting ACAPS. On the "top cap" of the ESA, there is a flat circular surface. Normally this surface doesn't have any practical purpose other than acting as a grounded wall to block energetic particles.

However, it is possible to attach a circuit board there. The board area is limited to a 64-mm circle, because of the size constraints of a CubeSat small orbiting satellite. For a sounding rocket sub-payload, the diameter would be similarly limited to around 70 mm.

As discussed later in Chapter 7, placing the HVPS out on top of the ESA has a few disadvantages with wiring, but fitting all components on that ~60-mm circular area was a success.

Interfaces: What will the HVPS connect to?

This HVPS was designed without any knowledge of the detector or data gathering electronics. They simply had not been designed yet, so the HVPS was designed to work flexibly and independently.

Figure 2.4 below shows the basic interfaces of the HVPS, which are listed below:

- The only input voltage that this power supply truly needs is 3.3 V, which is a standard voltage used by most modern electronics, especially computers. This means it can be tested alone, which is necessary since the other hardware does not exist yet.
- UART (Universal Asynchronous Receive/Transmit) is for communicating with other electronics on the spacecraft. This UART design was chosen because it requires only two wires, is very simple to work with, and is like USB. The HVPS should be able to receive commands, send status information, and voltage/temperature readings through the UART. UART can conveniently be used with a USB adapter to communicate with any computer.
- SBW/JTAG are the two wires required for reprogramming and debugging the circuit. These are not expected to be used in space at all but are very helpful with testing on the ground.
- The HVPS will also output analog divided versions of its high voltages. This gives flexibility for an external detector circuit to read those voltages itself.

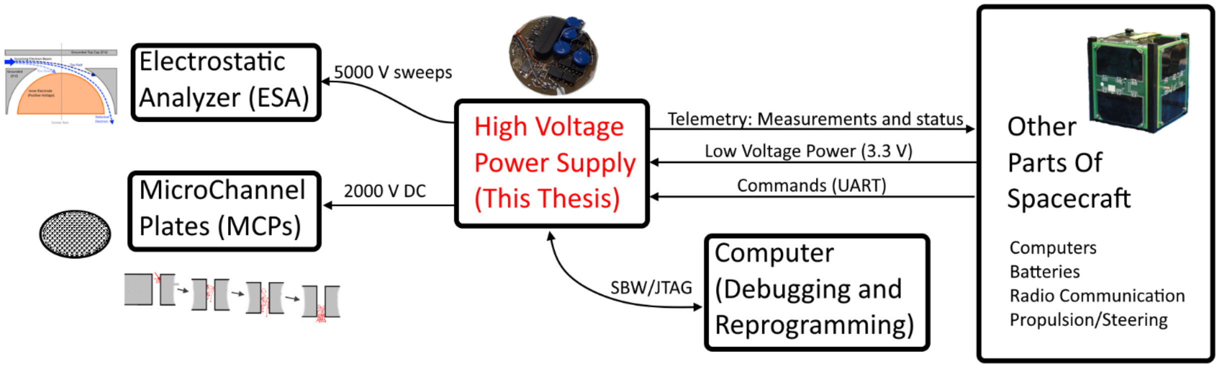

Figure 2.4: Basic Interfaces of the High Voltage Power Supply

Chapter 3. ESA Sweeping Supply

The ESA and MCP supplies are closely related, and they are combined in the final design. However, back in 2020 at the start of this thesis, they were considered as two separate circuits. The design began with the ESA sweeping supply (see components in HVPS block diagram of Figure 17) as I assumed that producing rapidly changing sweeps would be far more difficult than producing tunable DC voltage. Chapter 4 will focus on the MCP supply.

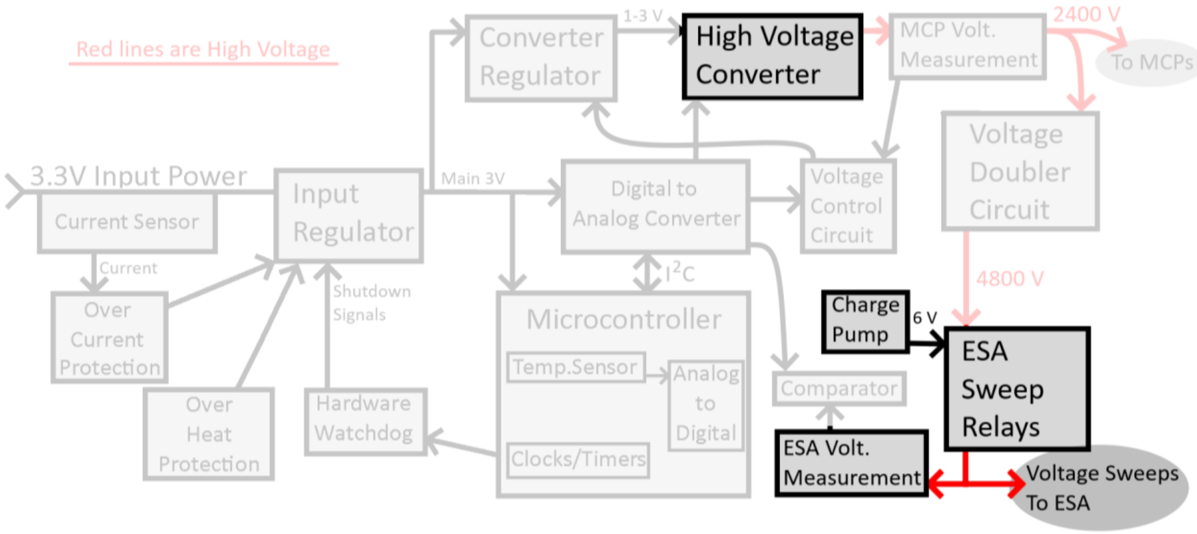

Figure 3.1: ESA Sweeping Supply Portions of HVPS Block Diagram

How to create the high voltage?

I quickly decided that custom-designing a transformer-based switching power converter with a voltage multiplier was not reasonable. It is certainly possible, and most power supplies for space applications seem to be fully custom designs for the oscillators, transformers, and voltage multipliers. However, I chose not to do a fully custom design because there are commercially available options. I preferred to spend more time designing and testing more specialized elements of the power supply, rather than reinventing well-established technology for generating the voltage.

The commercially available self-contained high voltage converters I looked at are not designed for space applications, but they are still impressive enough to attract attention. Specifically, a company called EMCO has been making miniaturized converters for many years, since at least 1991. Craig Pollock introduced them to me back in 2018, long before I was even an electrical engineering student. Craig had considered using them for a spacecraft in the past but was never able to make such a move.

EMCO had recently been purchased by XP Power, but they still produce small DC-DC high voltage converters. At the time I started this thesis, there was no good competition for EMCO/XP Power in the realm of converters at > 3000 V, < 1 W of less than 12.7 mm (½ in) height.

I chose XP Power's A-series of converters, shown below in Figure 3.2. This choice was almost entirely based on their height of only 6.35 mm (1/4 in). I also considered products from Advanced Energy, Pico Electronics, and HVM. Some other product lines have 12.7 mm (1/2 in) cube shaped converters with a similar overall mass and volume, but I dismissed them because of their increased height.

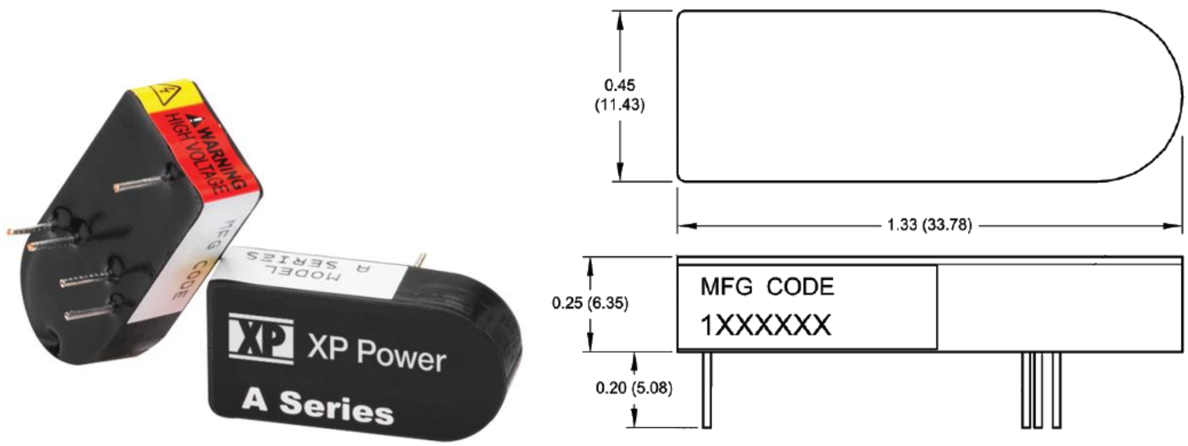

Figure 3.2: XP Power A-series Converter [11]

I modified the converter's wires to lay flat, further reducing thickness. In hindsight, I was overly focused on the height of the converter. Perhaps 12.7 mm (½ in) would have been acceptable after all. A limit on board height was never clearly defined, so I was simply trying to make every component as thin as possible.

The A-series converters use an oscillator, transformer, rectifier, and output capacitor. The converter takes in DC at a few volts and outputs DC at high voltage. This means all of the AC is contained within, so the converter acts as a "black box" of DC in, DC out. The remainder of the power supply circuit does not need to work with high frequency AC at high voltage. There are still many complications though, because the converter's output is difficult to efficiently control, but at least it is guaranteed to be a source of some kind of high voltage DC. More details about operating the converter will be found in Chapter 4.

How high should the sweep voltage go?

For science, higher voltages are desirable to increase the maximum energy the ESA can measure. Auroral electrons can be above 20 keV. Ideally, ACAPS would even reach 40 keV, which corresponds to about 7000 V on the ESA. At times I considered only using 3000 V for simplicity, but that wouldn't have fully met Requirement 1. The ESA can easily handle over 6000 V before experiencing high voltage breakdown, so it was not a limiting factor in the voltage decision.

However, practical design concerns pushed me toward lower voltages. Designing a circuit board becomes more difficult at higher voltages, because additional components like relays and resistors in series are required, and they must be spaced farther apart to avoid arcs.

Also, there is a significant drop off in availability of commercial parts beyond about 5000 V. It is common for optocouplers and relays to have 3750 or 5000 V isolation between input and output, but rarely anything more. Unfortunately, this isolation limit can not be overcome by simply stacking more components in series. In the smallest of sizes (under 15 mm thick), commercial board-mount DC-DC converters are only available up to 6000 V. Also, 6000 V is the highest common rating for the thinnest ceramic disk capacitors. I decided to design for 6000 V at first, and later downrated to about 4500-5000 V, based on concerns about reliability of high voltage components.

How to generate a sweep?

As seen in Figure 1.8 and defined by Requirement 1, the voltage sweep for the ESA needs to cover more than three orders of magnitude, from 5000 V down to a few volts. Covering that range at all is a challenge, and the need to sweep quickly, repeatably, and power-efficiently only magnifies the challenge.

Over the course of this thesis, I investigated three different voltage sweeping methodologies: direct sweeping, optocoupler sweeping, and RC sweeping.

Direct Sweeping

The direct sweeping method as shown in Figure 3.3 below was the first and crudest method I investigated.

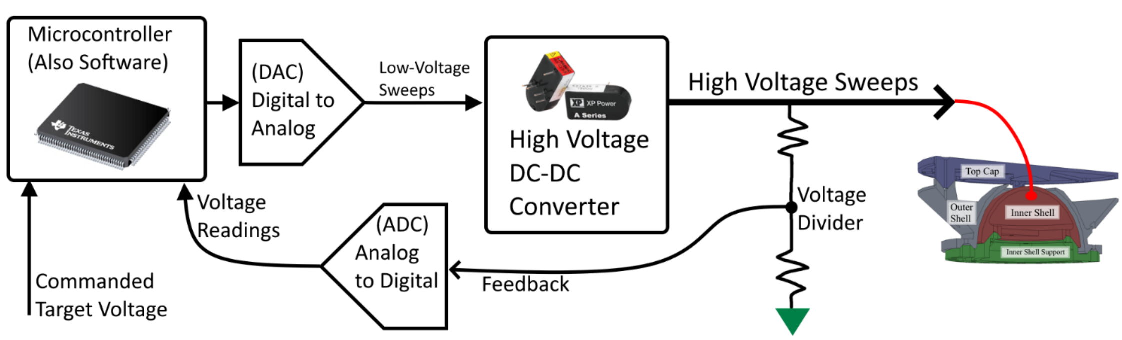

Figure 3.3: Direct Sweeping Concept

The core idea of direct sweeping is to first digitally create a low-voltage sweep with a digital-to-analog converter (DAC), and then amplify the voltage of that sweep with a DC-DC high voltage converter. In general, the precision and speed of modern digital-analog and analog-digital electronics incentivizes working digitally whenever possible. Direct sweeping could allow the consistency and unrivaled flexibility of digital-to-analog conversion to be scaled up to high voltages. However, there are several problems.

First, the digital-to-analog converter (DAC) needs to have a wide output range to cover the full range of sweep voltages. Because the high voltage sweep covers about 3 orders of magnitude, the DAC must also. That means the resolution for a single bit must be around 1/10,000th of the total range. This means 14, or probably even 16 error-free bits would be required to achieve sufficient resolution at the lowest voltages (2^{16} is 65,536 bits). Analog-to-digital converters (ADCs) of 16 bits or more certainly do exist, but working with such small voltage increments would be a challenge. If the maximum DAC output signal were 3 V, the low end of the sweep would be several thousand times smaller, about 1 mV. Such a small control signal is very vulnerable to noise and interference.

An even bigger problem is that DC-DC converters do not always maintain a reliable input to output voltage relationship at very low input voltages. In other words, they stop being linear or predictable at low voltages. The A-series converters I chose only have about 1 order of magnitude of useful adjustment range. For example, I found that the 6000 V output model (A60-P5) has a practical minimum of 500-1000 V, depending on load resistance. Even at the low voltages that it can achieve, the converter is practically impossible to control, and is very sensitive to variations of the load. Any converter that is designed for thousands of volts will not be engineered to output something like 2 V or 10 V.

As an alternative to the “proportional” converters, there are “regulated” converters that include a feedback control loop. While this does simplify the task of controlling a converter, these regulated options introduce other problems. When comparing products from XP Power, the regulated converters are also physically larger and have slower response times. The HVM Technology SK series might work to achieve low output voltages, except it is unreasonably large for this application: Not only is the SK series twice as tall as the A series, but it also covers over 3 times more circuit board area than the A-series shown in Figure 3.2. The XP Power C series is similarly too large.

As an extreme example, even large benchtop laboratory power supplies struggle at low voltages. For example, our lab’s PS350 digitally controlled 5000-V power supply can not reliably output less than 20 V. Even if its size, power consumption, and minimum output were acceptable, the PS350 also has a long output voltage settling time, which is far too slow for the requirement of achieving even a 20-Hz sweep rate. This is because it wasn't designed for rapid sweeping, it was designed to output constant DC voltage. Power converters are generally designed to filter out variations in the output of more than a few Hz, because the designers consider those frequencies to be noise. For this HVPS however, those frequencies are important since the desired ESA voltage-sweeping rate is approximately 20-200 Hz.

Why are converters slow to respond?

The main reason for a slow response time on a converter is output capacitance: There are usually large capacitors on the output which help for filtering DC, but those capacitors are suddenly very harmful when trying to sweep.

Repeatedly charging a capacitor can take a lot of power, and that power is usually simply wasted when the capacitor discharges. Figure 3.4 below shows how the energy stored in a capacitor depends on voltage squared.

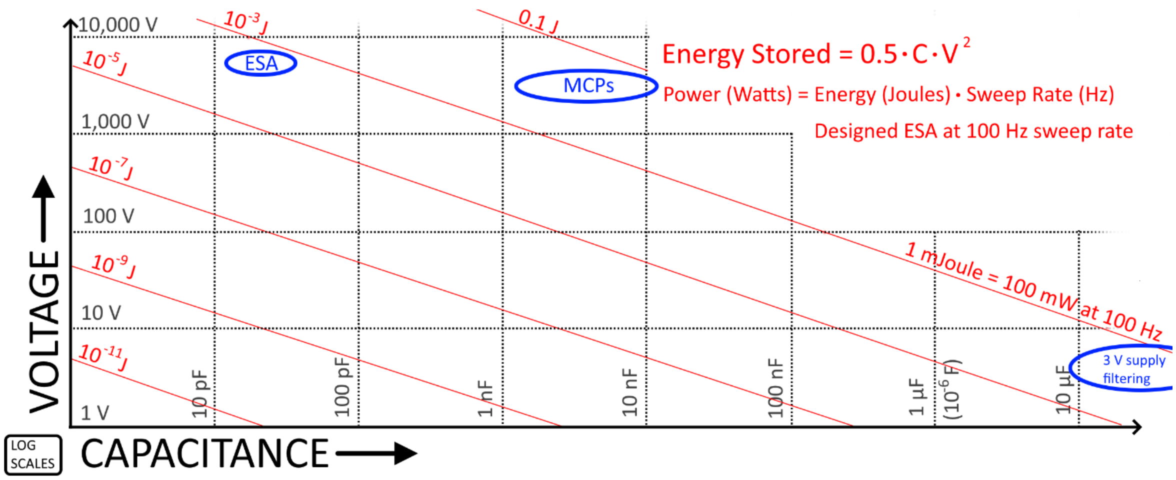

Figure 3.4: Capacitor Energy and Power

For example, charging to 5000 V at a rate of 100 times per second on a relatively small converter output capacitance of 1 nF would consume 1.25 W. That is already more power than I want the entire HVPS to consume, even assuming the converter was 100% efficient.

The capacitance of the ESA itself is estimated as 10-50 pF, much smaller than the output capacitors in high voltage converters. This means nearly all of the converter's power would be used simply charging the converter's own capacitance, instead of the actual ESA. As shown in Figure 3.4, an ESA of 20 pF charging 100 times per second up to 5000 V should only consume 0.025 W, which is very reasonable.

As a first exploration of commercially available high voltage sources, I purchased a +1200 V proportional DC-DC converter, the A12-P from XP Power. It operates on 0.7 V to 12 V DC input. When operating linearly, the output is conveniently about 100 times the input voltage.

Figure 3.5 below shows output voltage over time as the converter is turned on and off. The green trace is the input voltage and yellow is the output voltage, both measured through a 10:1 voltage divider. This voltage divider was measured by an oscilloscope probe with about 10 pF capacitance, which is insignificant compared to the converter.

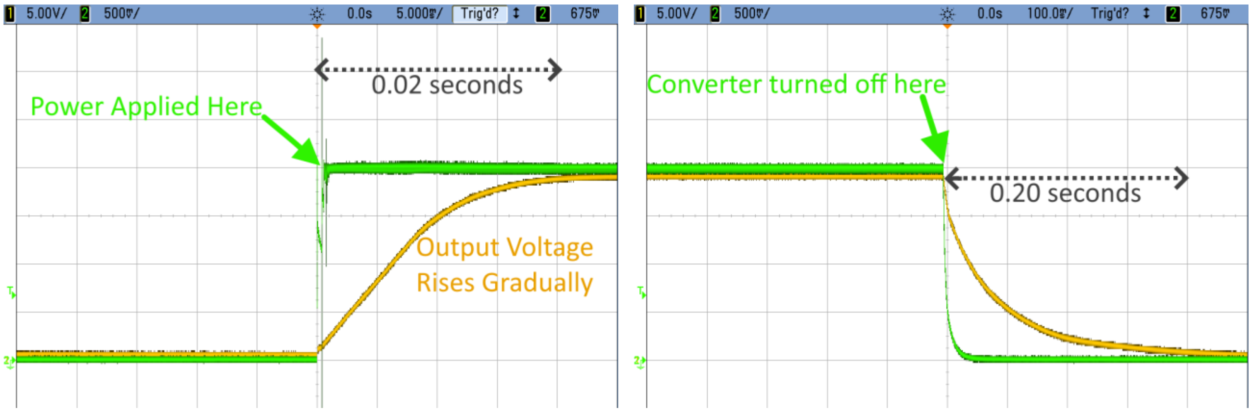

Figure 3.5: XP Power A12-P Converter transients.

Interestingly, the rise time shown above is about 10x faster than the fall time, but both are unfortunately too slow. Rise time is about 20 ms to reach 95% voltage, although it rises faster with higher input voltages. Fall time is several hundred ms, depending on the load resistance. I also found that this A12-P converter shows nonlinear behavior below about 20 V output (2% of maximum output), so it is definitely not capable of operating across three orders of magnitude.

Optocoupler Sweeping

Optocoupler sweeping was the next option I considered, after dismissing RC sweeping for being unreasonable because of its insufficient speed, poor low voltage performance, and excessive power consumption.

An optocoupler is simply a light-sensitive semiconductor combined with an LED. It works similar to a relay, but with no moving parts or magnetic fields. Increased illumination of the input LED will decrease the effective resistance of the output, allowing more current to flow. Another way to think of it is that more current through the input LED allows more current through the output. Optocoupler sweeping uses optocouplers in a "push-pull" configuration, controlling the ESA voltage using an electronically adjustable voltage divider. A “push-pull” circuit means there are two optocouplers acting as variable resistors. The high side optocoupler can “push” the voltage toward the converter’s output, and the low side optocoupler can “pull” the voltage down toward ground, or 0 V (see Figure 3.6 below).

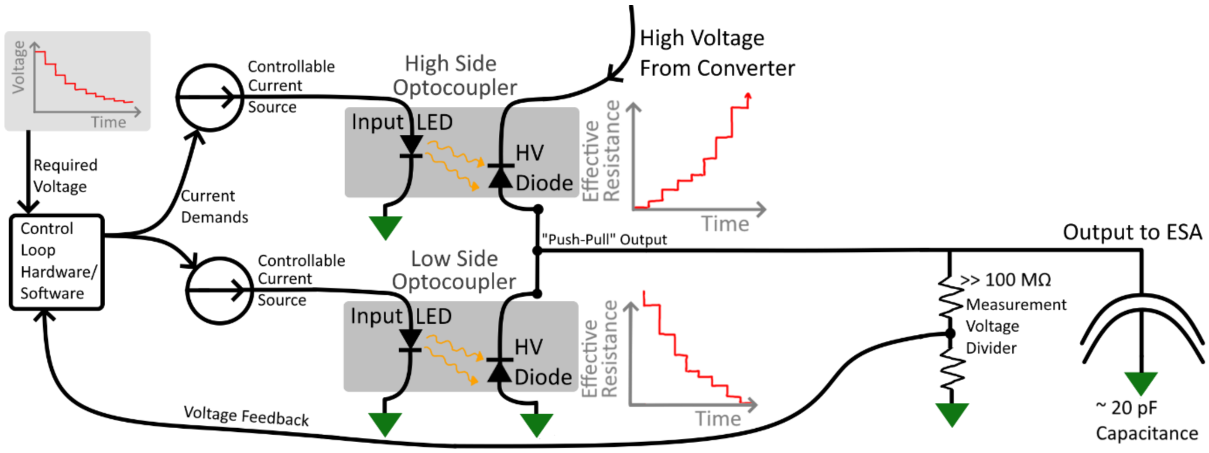

Figure 3.6: Optocoupler Sweeping method concept schematic

Optocoupler sweeping can be much faster than direct sweeping because the high voltage converter's output voltage does not have to change to produce a sweep. Instead, only the optocouplers need to change their resistance rapidly. I considered building custom optocouplers with individual LEDs and photodiodes, but then learned about commercially available options.

Commercial Optocouplers

Southwest Research Institute (SwRI) offers the SW1502A optocoupler designed for high voltage space applications. It fits within a 13 mm x 13 mm x 9 mm volume, a large but maybe acceptable size for this application and its voltage and radiation ratings are far higher than this project requires. It has low enough leakage current (< 100 nA) to achieve extremely low sweep voltages and reasonably low output capacitance (5 pF) [12]. Amptek also offers a widely used optocoupler called the HV801, which is even smaller, just over 10 mm x 10 mm x 5 mm. It also has an acceptable leakage current, output capacitance, and voltage rating [13].

If optocouplers are so good, why consider anything else?

Optocouplers appear to have been the industry standard method of making voltage sweeps on large and expensive spacecraft for at least 20 years. However, there are two disadvantages to optocouplers: power consumption and complexity.

Output power consumption: The main point of using optocouplers is to command the voltage accurately and quickly. However, the response time of this circuit is limited by output resistance times capacitance, often simply called RC. Inductance is not a significant factor. For the output voltage to adjust quickly, the RC time constant must be much shorter than the time between voltage steps. In each single time constant, the voltage error will decrease by 1/e, or about 63%. For this power supply design, the C (capacitance) is constant, because the circuit board, the optocouplers, and the ESA itself will have unavoidable "parasitic" capacitance, about 20-50 pF total. To achieve faster response, R must therefore be reduced. Unfortunately, having a lower resistance means more current will flow through the optocouplers, therefore more power will be consumed. Figure 3.7 shows a conceptual range of voltage responses to a stepped voltage sweep at different effective optocoupler RC time constants.

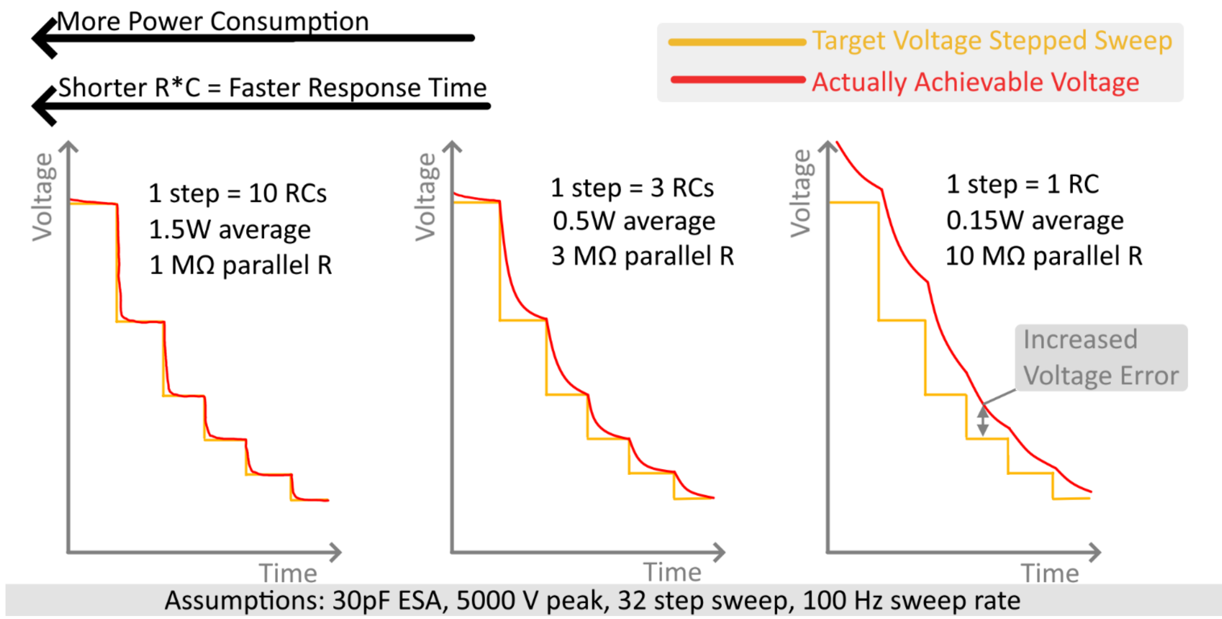

Figure 3.7: Optocoupler Sweeping waveform comparison with various RCs

At one extreme (left side graph), the RC time is held constant at 1/10th the duration of a voltage step, so the optocouplers can practically change the ESA voltage (red line) instantly. This means the optocouplers can accurately create any arbitrary sweeping pattern. In a perfect world, all sweeping power supplies would be like this. Unfortunately, it requires too much power. Such a small RC might be more feasible at lower voltages, or if the sweeping rate was something slower, like 10 Hz.

The middle case is somewhat more achievable, but the power required is still excessive.

At the other extreme (right side graph), the RC time constant defined by the optocouplers is equal to the step length. With this dramatically increased RC, the ESA voltage is always lagging, following an overall exponential decay curve. The ESA voltage never has time to reach its target level before the target changes again. This means the optocouplers are not serving their purpose of accurately controlling the output voltage.

I analyzed the power consumption required for optocoupler sweeping. This analysis was done by enforcing a constant RC for each of 32 voltage steps in a sweep. Power consumption for each step was calculated, based on the series combination of the two optocoupler resistances. To calculate average power, I averaged the power from all steps of the sweep. Figure 3.8 shows that power consumption varies throughout the sweep, because power depends on the series resistance of the two optocouplers, while RC response time depends on their parallel combination. When the output voltage is halfway between maximum and zero volts, the combined resistance of the two optocouplers is lowest, so power consumption is highest.

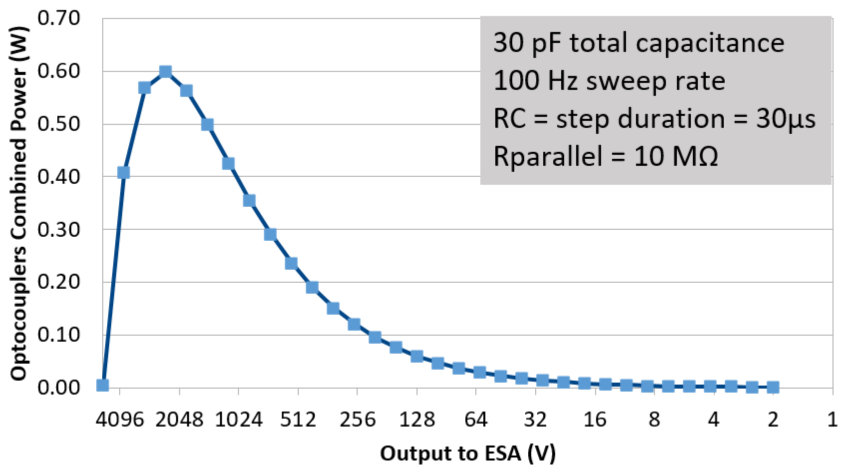

Figure 3.8: Optocoupler power consumption during a voltage sweep at a constant effective RC

Based on the data shown in Figure 3.8, only 37.5 mW average is theoretically required to charge the 30 pF ESA capacitance 100 times per second up to 5000 V. Meanwhile, even a very low-power example of optocoupler sweeping with RC equal to the step length still consumes 147 mW. Therefore, optocoupler sweeping consumes far more power than is necessary, even when it isn't controlling the voltage very accurately. Also, the exponential voltage decay curves shown in Figure 3.7 are not necessarily bad, since an exponentially decreasing voltage is exactly the kind of voltage sweep we want for the ESA. This realization led me to question the purpose of optocoupler sweeping, if the basic voltage curve can be achieved without any optocouplers at all.

Complexity: The SwRI and Amptek high voltage optocouplers both require significant and variable input current to operate. The ratio of applied input current to allowed output current is called current transfer ratio (CTR), and it is generally around 1%. Unfortunately, CTR is not very consistent, varying with applied voltage, age, temperature, and radiation exposure. The current through an optocoupler is not predictable enough to be assumed accurate within a few percent, which is what ACAPS needs. Instead, the creation of a voltage sweep requires closed-loop control, where the output voltage is constantly being measured and fed back into the circuits which control the optocouplers (shown in the Figure 3.6 schematic). This closed-loop control circuit must operate at high speed compared to the sweep without causing oscillations, adding significant complexity to the design.

RC Sweeping

After deciding to avoid optocoupler sweeping because of its high power consumption and complexity, I chose to use RC sweeping for this thesis. RC sweeping is the logical extension of the optocoupler method with an extremely long time constant, even longer than is shown on the right side of Figure 3.7. Instead of trying to force the output toward a certain target value, the RC time constant is allowed enough time to produce the desired exponential voltage sweep automatically. RC sweeping trades controllability for efficiency and simplicity. Figure 3.9 below shows a simplified schematic of RC sweeping with realistic component values.

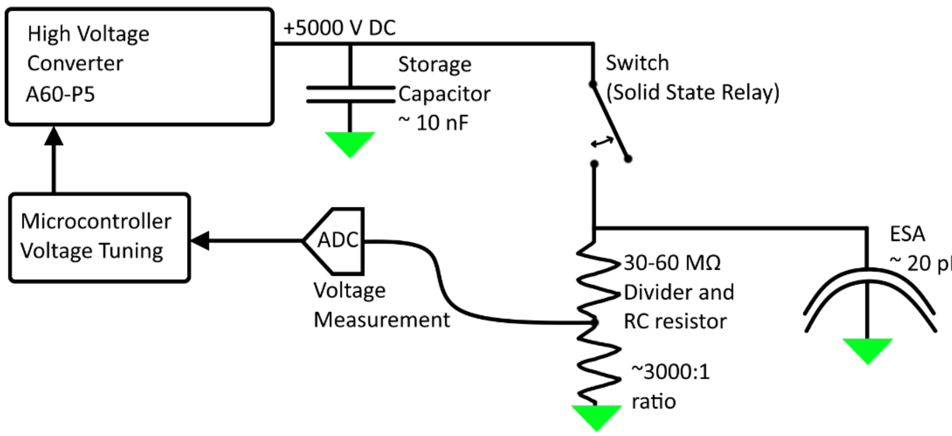

Figure 3.9: RC Sweeping Basic Circuit Schematic

RC sweeping produces a consistent exponential decay voltage curve by design. No active control or regulation is required. There is no concern about semiconductors or switches being nonlinear, because the only nonlinear device is the switch, which is turned off during the sweep.

RC sweeping is also known as “Capacitive Discharge”. It was more common in the early days of spacecraft, probably because computers weren’t advanced enough to actively manage a power supply's output. I know of at least three examples of satellites with RC sweeping supplies between 1978 and 2007, all for ESAs associated with Los Alamos National Laboratory (LANL) [14] [15] [16]. I have also been told that sounding rockets associated with University of New Hampshire in the 1960s used RC sweeping [17].

Today, I believe simple RC sweeping is not obsolete, at least on small spacecraft. It is still simple, cheap, and easy to design. RC sweeping operates at the physical limit of minimal power consumption, given a required sweep rate and a minimum ESA capacitance. It has no particular issues with high voltages, and the only limit to its low-voltage performance is leakage current and coupling through the relays. Today, modern computers make it easier for any nonlinearity in the RC sweep to be calibrated out.

For some applications, a major disadvantage of RC sweeping is the lack of control. The shape of the sweep is not governed by software, so it cannot be commanded to produce any arbitrary sweep shape. It can only ever produce exponential decay voltage curves because it is governed by the discharge of a constant capacitance through a constant resistance. For ACAPS, this limitation is not a problem, since we would likely only use one voltage sweep in any case.

Since the capacitance is fixed at about 20-50 pF based on the ESA, the required resistance can be calculated to achieve a certain RC sweep rate. To fit a sweep range of 5000 V to 2 V at 100 Hz requires about 10 exponential time constants within 10 ms, so the RC time constant should be around 1 ms. For an estimate of 30 pF ESA capacitance, about 30 MΩ of resistance would be required. It definitely makes more sense to select a resistance as needed and assume constant capacitance, because adding extra capacitance instead would increase power consumption (as seen in Figure 3.4).

High Voltage Switching Options

Once I had chosen to use RC sweeping, I needed to choose a high voltage switch. To start a voltage sweep, a switch needs to disconnect the RC circuit from the rest of the power supply, so that the ESA capacitance can begin discharging through a resistance. Then at the end of the sweep, the switch must connect again to recharge the ESA capacitance and prepare for the next sweep.

The requirements for this switch are quite demanding:

- At least 5 kV operating voltage between output terminals (Note this is not the same thing as isolation.).
- The switch must consume minimal power, ideally under 50 mW since the target total power for the whole supply is around 0.5 W.
- It must switch quickly enough to allow 100 Hz sweeping. This means the turn-on and turn-off times should be around 1 ms or faster.

I evaluated mechanical relays, diode-like devices, transistors with opto-isolation, and solid-state relays. See Table 3.1 on the next page for a comparison of switch types.

Table 3.1: High Voltage Switching Methods

Parameters	Mechanical Reed Relays	Thyristors or SCRs or TRIACs	Transistors: Opto-Isolated (BJT + LED)	Solid State Relays
Size	50 x 10 x 10 mm Only need one.	Need 4-6.	About 5 x 5 x 2 mm Need 10.	10 x 8 x 3 mm Need 4.
Voltage per Item	6000 V+ available	1200-1600 V	600 V	1200-1500 V
Speed	2 ms at best	Few µs at best?	Very fast (µs)	Under 1 ms
Average Power Throughout Sweep	>100 mW? (25 mA at 24 V, for 20% of the time = 125 mW)	Uncertain, probably much worse than transistors.	Very low, limited only by gate drive electronics.	~ 10 mW (2 mA average at 5 V with 4 relays in series)
Leakage Current	Near Zero	Assumed to be bad, often completely ignored on datasheets.	Uncertain, probably similar to SSR or better, except when very hot.	< 10 µA guaranteed. < 1µA normally.
Other Concerns	Limited mechanical lifetime. (still $>10^9$ cycles). May produce electrical noise, needs shielding. May not switch cleanly (bounce).	Require current pulses to turn on. Needs isolated drive electronics. Generally designed for much higher currents	Each need separately isolated drive electronics. Complicated circuit, lots of parts	Turn-on time depends on input current. Either fully on or off, no linear region.

While **mechanical relays** are becoming obsolete in most realms of engineering, they are still somewhat competitive at high voltages. This is because high voltage breakdown (arcing) limits how small semiconductors can be. Normally, a low-current transistor can be literally microscopic, but such miniaturization is not possible when several thousand volts must be applied across the silicon. High voltage semiconductors generally also have slower switching speeds than their low voltage counterparts have.

Mechanical relays have significant power consumption because they need to run current through a coil to create a magnetic field which physically moves the contacts. Using higher input currents can achieve faster switching speeds, but at the price of increased power consumption. Even with quite high input currents, the fastest relays I researched had switching times of about 2 ms. At a sweep rate of 100 Hz, that would still mean 40% of the total time is spent switching on and off between sweeps. Switching time is wasted because the ESA will not detect electrons during that time. The power consumption required by mechanical relays is excessive for this small of a power supply. They are generally designed for 12 V or 24 V, which is inconvenient since most other electronics only need 3.3 V or 5.0 V.

Transistors (both MOSFETs and BJTs) can be used at high voltage. Unfortunately, they are not available in small enough sizes beyond 1000-1500 V. This is not a disaster, as MOSFETs can be safely added in series to increase their voltage limit. However, transistors do not have input-to-output isolation, unlike relays. This is a problem because it means each transistor needs its own isolated control circuit. To turn on a normal n-channel MOSFET, the gate must have a voltage higher than the source. That voltage difference only needs to be a few volts, but the source may be already at several thousand volts above ground. Voltage isolation can be done with a large coupling capacitor and a rectifier, or with a LED and photodiode. The latter case is essentially a custom-built optocoupler.

Thyristors, SCRs, and TRIACs are other types of silicon devices related to transistors. They are able to work at high voltages and currents, but they are generally large, leaky, and difficult to control. Despite being semiconductor devices, they require large currents to activate, much like mechanical relays. They are better suited for higher power applications, like power supplies operating at many watts or even kilowatts. They are also difficult to stack in series like transistors because they are not isolated.

Solid state relays have input-output isolation like a mechanical relay and behave very similarly, but without the moving parts. Many solid-state relays use optical coupling similarly to the optocouplers previously discussed. The main difference is that these solid-state relays are an “on or off” type of device, while optocouplers are designed to operate somewhere between fully off and fully on.

The solid-state relay I chose is the Panasonic AQV258H5 (see Figure 3.10). I chose it based on its small size, low leakage current, and low input power consumption. It is rated to 1500 V between the output pins, with 5000 V of isolation between input and output. When turned on, the resistance is only a few hundred ohms, which is comparatively near to zero, since the other resistors in this power supply are millions of ohms. These solid-state relays can be stacked in series very easily because of their input-output isolation. The input side behaves exactly like an LED, requiring only about 5 mA to turn on. That is a low enough current that it is no special challenge to control.

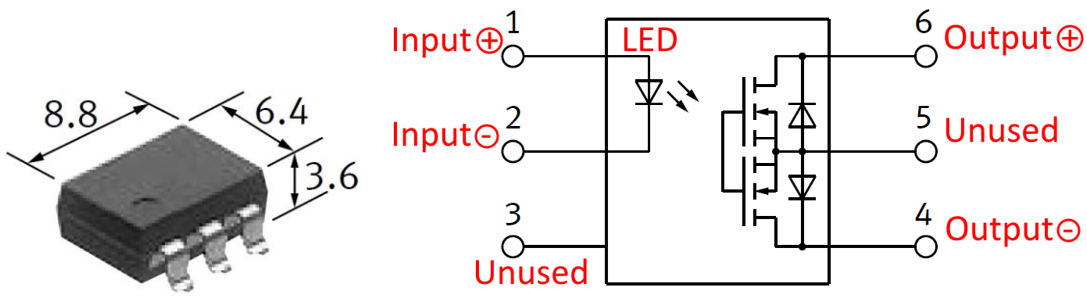

Figure 3.10: Panasonic AQV258H5 Solid State Relay (units of mm) [18]

How much relay leakage current is too much?

Relay leakage current is a concern because it flows into the ESA RC sweep circuit, pushing its voltage higher than what RC decay alone would cause. Figure 3.11 shows the flow of the relay leakage current which might disturb the idealized ESA voltage sweep.

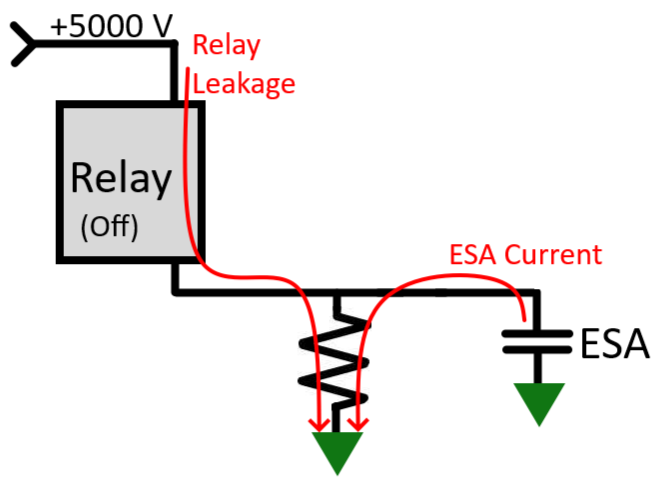

Figure 3.11: Panasonic AVQ258H5 Relay Schematic with Leakage Current

There is no exact limit on the acceptable leakage current because the lower end of the ESA voltage sweep is somewhat flexible. Relay leakage current has the most impact at the lowest voltages, near the end of the sweep. The design level of 30 MΩ at 5000 V would have 167 µA leaving the ESA capacitor at the start of the sweep. 167 µA is enough to make any reasonable relay leakage current appear completely insignificant. However, once the voltage has swept down to about 1 V, there will only be 0.03µA of current leaving the ESA capacitance, yet the relay leakage current will have remained constant. That is 5000x less current from the ESA capacitance, therefore a relay's leakage current can be much more influential. In this example, 1 µA of leakage current would be insignificant at 5000 V, but would ruin the sweep completely at 1 V.

The leakage current of the AVQ258H5 relays becomes significant for output voltages below a few volts. Note that as the sweep approaches zero, the voltage difference (load voltage) across the relays reaches a maximum. As a rule of thumb, I wanted to keep leakage below 10% of the intended ESA current, so that the RC time constant would not vary by more than 10% due to leakage. At the lowest required voltage of 2 V with an assumed 60 MΩ resistance, this means the leakage should be kept below 0.004 µA (4 nA). Figure 3.12 shows leakage current from the relay datasheet.

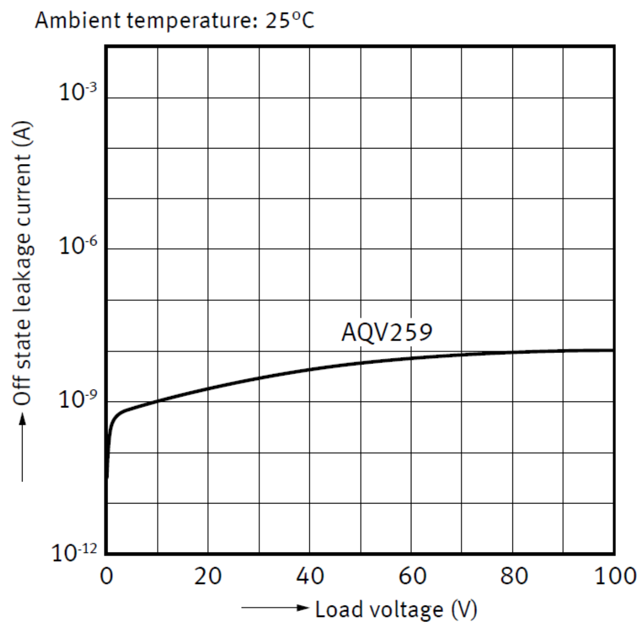

Figure 3.12: Panasonic AVQ258H5 Relay Leakage Current [18]

Unfortunately, Figure 3.12 does not provide a full picture. It implies that leakage current will be constant beyond 100 V across the relay, but testing shows this is not true. Also, Figure 3.12 shows leakage current at 25 °C, but leakage current in electronics almost always increases with temperature.

In addition to this graph, the datasheet also specifies 10 µA as a maximum leakage current, but that is probably a very pessimistic assumption. With very simple tests on the relays, I determined that leakage is normally far below 1 µA. At the time I selected these relays, I was aware that the leakage might not be satisfactory at high temperatures, but I couldn't find a better alternative. All semiconductor devices experience increased leakage current when operated at higher temperatures, so this is not a concern unique to solid-state relays.

How are the relays controlled?

The input side of the relay truly does behave like a diode. Unfortunately, diodes have a few non-ideal properties. First, the voltage required across a diode to achieve a certain current changes significantly with temperature. Figure 3.13 on shows the LED voltage drop and turn-on time versus temperature relationships for these relays.

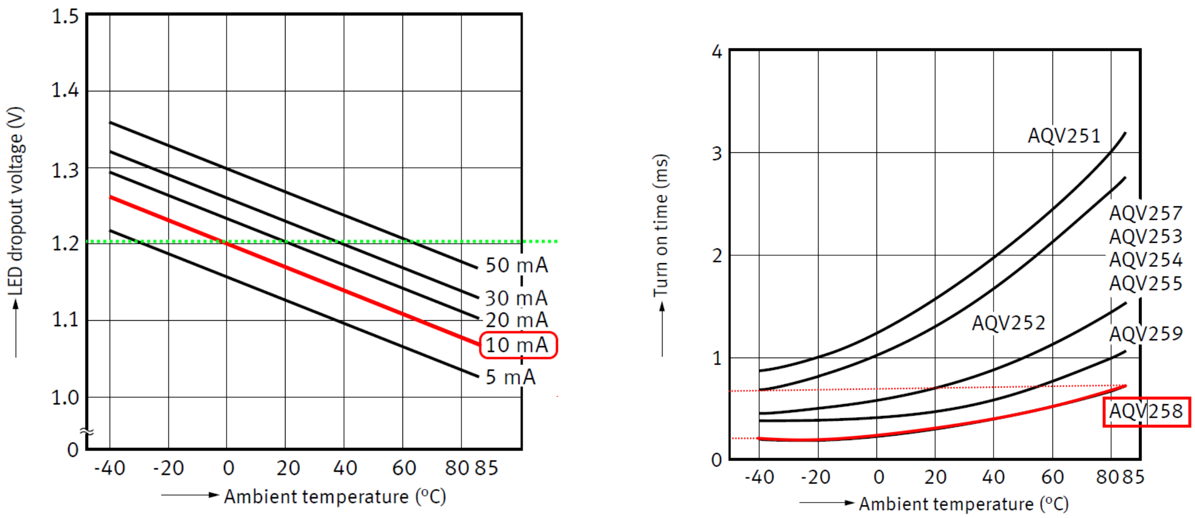

Figure 3.13: AVQ258H5 Relay LED Voltage (left) and Turn on Time (right) vs. Ambient Temp.

Unfortunately, the change in LED voltage drop across temperatures shown in Figure 3.13 is too dramatic to be ignored. For example, 1.2 V per relay (green line) causes too little current to even turn on the relay at -40 °C. However, at +80 °C, that same 1.2 V per relay would probably overheat and destroy the relay because of excessive current beyond 50 mA. This means that the voltage across each relay cannot simply be held constant. Also shown in the left side of Figure 3.13 is the very exponential relationship between voltage and current. A tiny increase in voltage applied to the LED causes a large change in current. Doubling the current takes only about an additional +0.03 V. This exponential change in current is the same characteristic that all diodes have, which makes series resistance or current control necessary.

Finally, the time required for the LED to turn on or off the relay varies dramatically with current, and with temperature. I chose 10 mA as a design target because the turn-on times are sufficiently short. Any higher current is just a waste of power.

I decided to actively control the current through the ESA relays. Figure 3.14 shows the schematic of the final RC sweeping circuit design. An n-channel MOSFET controls how much current is allowed to flow through the relays' LEDs. The MOSFET can switch off completely to turn off the relays. This design relies on closed-loop control. An OP747 amplifier measures the actual relay current through a 10 Ω shunt resistor and adjusts the MOSFET's gate voltage to maintain the target relay current.

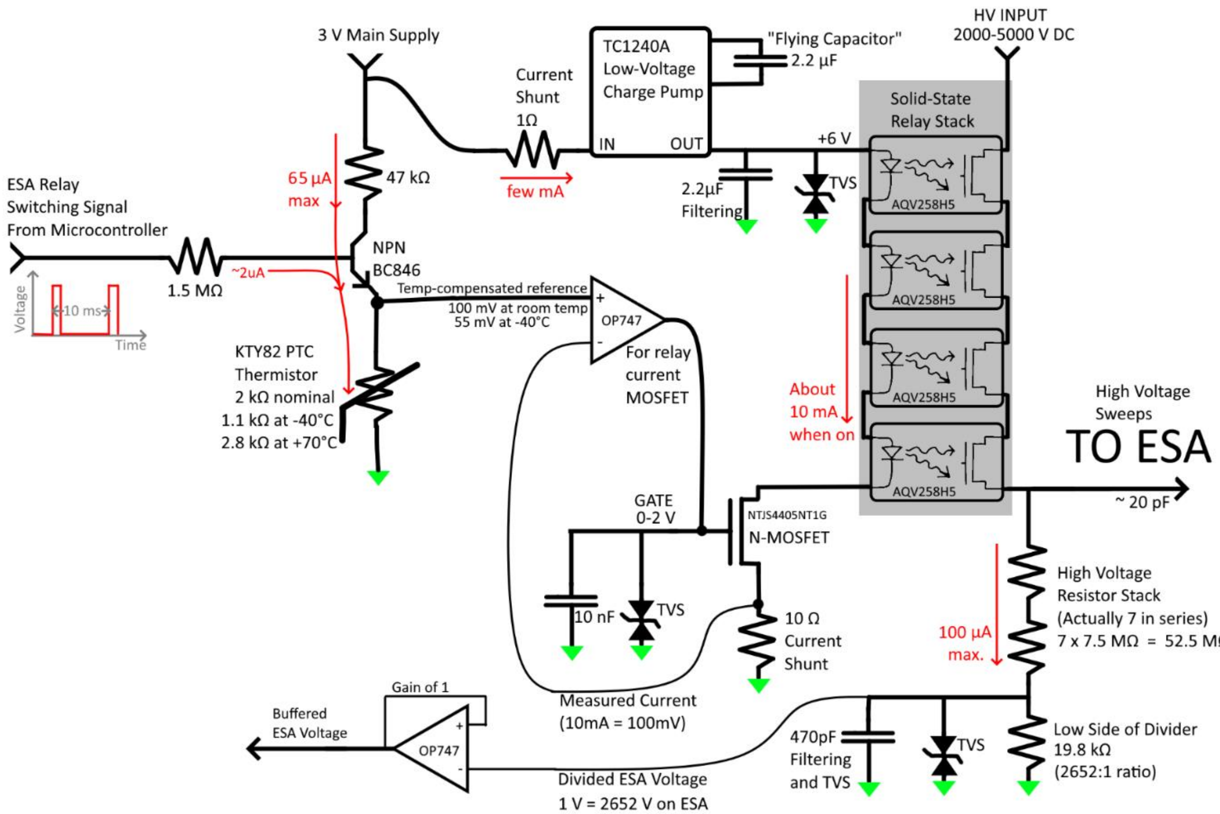

Figure 3.14: RC Sweeping Method Detailed Schematic

How does this circuit switch the relays off?

The circuit shown in Figure 3.14 can also turn the relays off completely in addition to adjusting their current when on. A square wave signal feeds the base of a NPN transistor (BC846) in the middle of the reference voltage divider which sets the target relay current. Whenever the incoming square wave is 0 V, the NPN transistor turns off, making the target relay current become zero. The amplifier "sees" this and brings the relay MOSFET's gate voltage to zero, turning it off.

How does this circuit deal with temperature?

Temperature compensation was definitely necessary. Figure 3.13 shows that a relay could take about 3 times longer to turn on when hot versus cold, despite the current being the same. The way to compensate for this is using higher relay current at higher temperatures. A thermistor is used to adjust the target relay current. In Figure 3.14, the thermistor's resistance increases as the current increases. A voltage divider made from the thermistor and the 47 kΩ resistor produces a reference voltage that is temperature dependent. At room temperature, this divider produces about 100 mV and feeds it to an OP747 amplifier controlling the MOSFET. However, when the thermistor heats up, its increased resistance causes that divider reference voltage to rise, as labeled on the left side of Figure 3.14. The amplifier compares that reference with the actual voltage across the 10 Ω relay shunt resistor. If the shunt has less voltage on it than the reference, the amplifier increases the MOSFET's gate voltage until the current increases sufficiently. This works because the shunt is converting relay current into a measurable voltage.

This relay control circuit required some trial and error to achieve successful implementation and operation. Figure 3.15 shows the first iteration of the design on a prototype board. This picture shows a circuit very similar to the schematic of Figure 3.14, but using a different amplifier. I tested it at up to around 100 °C using a heat gun to verify that the relay current would increase with temperature. The current remained in the reasonable range of 5-25 mA at all temperatures I tested.

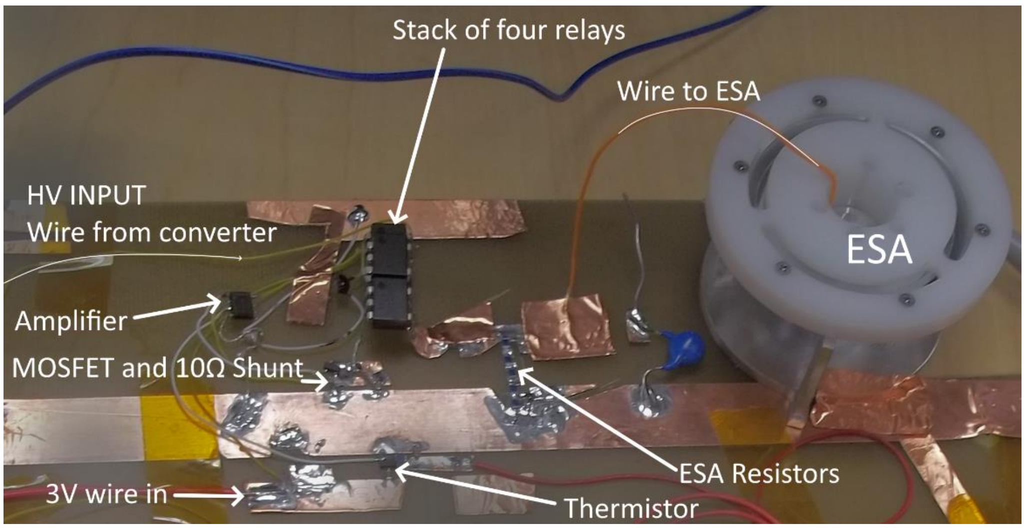

Figure 3.15: Early relay testing setup with ESA

Thermistor Selection

Figure 3.16 shows three different temperature compensation designs that I tested. I tried different types of temperature-sensitive circuits to match the "LED Operate Current" versus temperature curve for the relays. At first, I had looked at using an NTC (Negative Temperature Coefficient) thermistor on the high side of the voltage divider. Unfortunately, the NTC thermistor changed voltage far too dramatically, making it a poor fit for the "ideal" resistance shown in Figure 3.16. I also looked at using a Resistance Temperature Detector (RTD), but it had the opposite problem; not enough resistance change with temperature. Finally, I settled on a PTC (Positive Temperature Coefficient) thermistor as being the best fit, remaining within ±40% of the ideal resistance at all temperatures. This is shown on the right side of Figure 3.16.

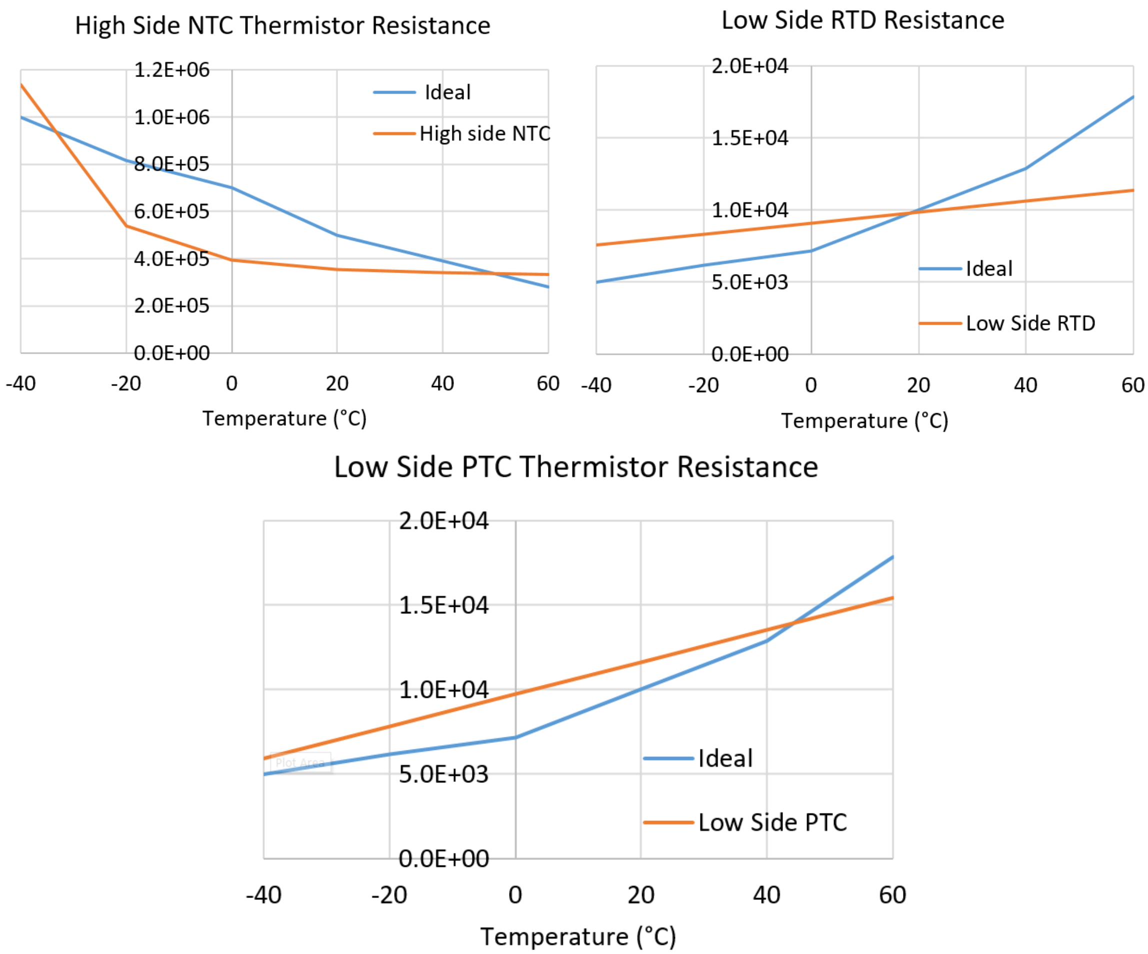

Figure 3.16: Temperature Compensation Comparison using NTC, RTD, and PTC Thermistors

Relay Stacking:

To turn on, each relay requires between 1.2 V and 1.4 V between its input pins, so a string of four relays in series requires 5-6 V to operate. I eventually decided to use only 3.3 V input for this circuit. This meant I had two options: either boost the relay input voltage or operate the relays in two parallel strings. Both options result in additional components on the board. I chose to boost the voltage because commercial low-voltage charge pumps are easy to use. I chose the TC1240A capacitor-only voltage-doubling charge pump, which I mentioned earlier in Chapter 2. The TC1240A is so small and has such low power requirements that it has little impact on the design, only taking 4 mm x 6 mm of board area. Without any load, I measured it consuming less than 100 µA (100 µV across the 1 Ω shunt resistor).

Sweep Components:

RC sweeping must rely on good quality resistors and capacitors to make sure the voltage sweep curve is as stable as possible, as well as good software to measure and calibrate out variations. Any change of the RC time constant is undesirable, especially those changes which are difficult to predict.

Fortunately, the capacitance seen at the sweeping supply's output is not likely to change significantly, because an ESA is basically a vacuum gap between two conductors, so it is very close to being a perfect capacitor. Vacuum is the most ideal environment for stable capacitance, although thermal expansion of the ESA may still alter its capacitance slightly. I expected the larger variations to be from capacitance within the circuit board itself, but I did not attempt to quantify that.

Thankfully, power consumption is really based on the capacitance alone, so the choice of resistors does not impact power very much. However, resistors are more of a stability issue: they change slightly across temperatures and over time. Thin film surface mount resistors have very good temperature stability [19]. For the ESA sweep resistors, I chose to use the Vishay MCU series of surface-mounted 0805-size resistors. They have a good temperature coefficient of ± 25 ppm/°C (parts per million per degree). Another advantage is that they come from a product series that also has low-resistance options, which allowed me to use them for the low-side of the voltage divider as well. To achieve a time constant of about 1 ms with a built-in capacitance of about 20-50 pF means the resistance should be around 30 MΩ. To ensure that the stack of resistors could withstand up to 6 kV without breaking down, I always used at least four 7.5-MΩ resistors in series to give at least 30 MΩ. I have found no problems with each resistor having 1500 V across it, although I later reduced the maximum voltage on each resistor to 1000 V for extra reliability.

Across the entire operational temperature range of -40 °C to +60 °C, the worst-case thermal variation of the resistance for four 25-ppm resistors is around 1%. (1% is 10,000 ppm). 1% is a very pessimistic calculation because it assumes all four resistors have maximum errors which add up in the same direction. Actual testing shows that the overall error in the series resistors is not as bad as this analysis.

Sampling the Sweep:

The lower right part of Figure 3.14 shows the voltage divider used for sampling the ESA voltage. For the final revision's ESA resistance of 52.5 MΩ, I used a low-side resistance of 19.8 kΩ, giving a divider ratio of 2652:1. The temperature coefficient of that tiny low-side resistance might not affect the total resistance very much, but it does influence the accuracy of the voltage divider ratio.

Voltage Divider Capacitance:

As explained in Chapter 2, capacitance on the voltage divider is a concern. When subjected to rapid voltage change (like a relay turning on), the parasitic capacitance on the voltage divider causes impedance to be equally shared across the divider resistors, since capacitors dominate the impedance at high frequencies. This can be a destructive situation, when a high voltage resistor divider that is supposed to have a 3000:1 ratio momentarily has a 2:1 or 5:1 ratio instead.

Figure 3.17 below shows an example of a voltage spike coupling through a voltage divider.

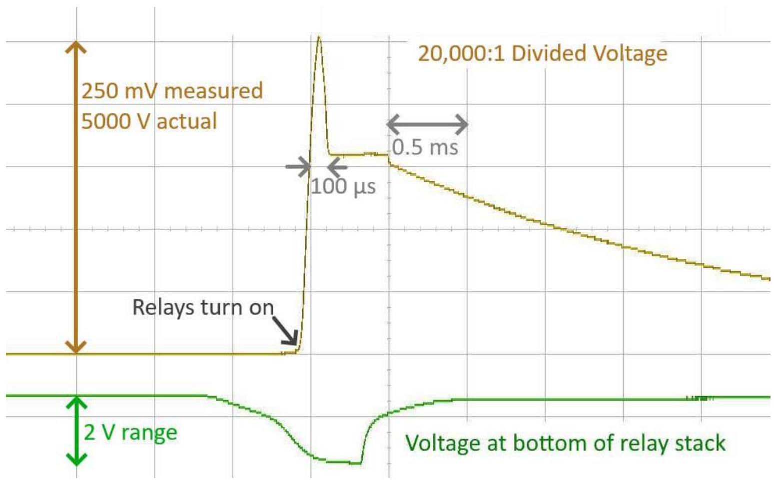

Figure 3.17: Voltage Spike from Capacitive Coupling through a Voltage Divider

This voltage spike was measured with an earlier version of the relay testing circuit shown in Figure 3.15. The voltage divider ratio was 20,000:1, with no protective devices installed. The yellow trace shows the measured voltage spike seen on the low side of the ESA voltage divider. That pulse of high voltage dissipates within about 100 µs, and in this case it was not large enough in amplitude to cause any damage. However, the only reason the voltage did not spike higher was the relatively slow turn-on time of the relays. If the relays turned on faster, the spike might destroy the OP747 amplifier used for measuring the ESA voltage. To avoid the capacitor impedance problem with the voltage divider, I added extra capacitance to the low side. At first, I used 10 nF. That stopped the voltage spikes, but it excessively smoothed the voltage waveform, so the divider was no longer a realistic sampling of the true ESA voltage. To get a better match, I used 470 pF instead. This is still enough capacitance that no voltage spikes appear. A more thorough design would tune the capacitance based on the theory of transmission lines to maintain constant impedance in the divider, but I simply used trial and error. It is difficult to accurately design, because I don't know the parasitic capacitances of the resistors and their circuit board.

As extra insurance against dangerous voltages appearing on the low voltage side, I added a transient voltage suppression (TVS) diode in parallel with the additional capacitance. Whenever the voltage exceeds approximately 4-5 V, the TVS diode will conduct current to ground until the voltage drops back down. A TVS diode is not perfect on its own because it takes time to activate but becomes more effective when combined with a capacitor to slow down any voltage changes. As seen in many of the HVPS schematics including Figure 3.14, I scattered TVS diodes on nearly all parts of the circuit, because they provide great insurance against wiring mistakes or accidental high voltage arcs and breakdowns.

Chapter 4. MCP Detector Supply and Converter Control

This chapter discusses the supply for the microchannel plate (MCP) detector and the circuit which manages the high-voltage converter on the HVPS as shown in Figure 4.1 below.

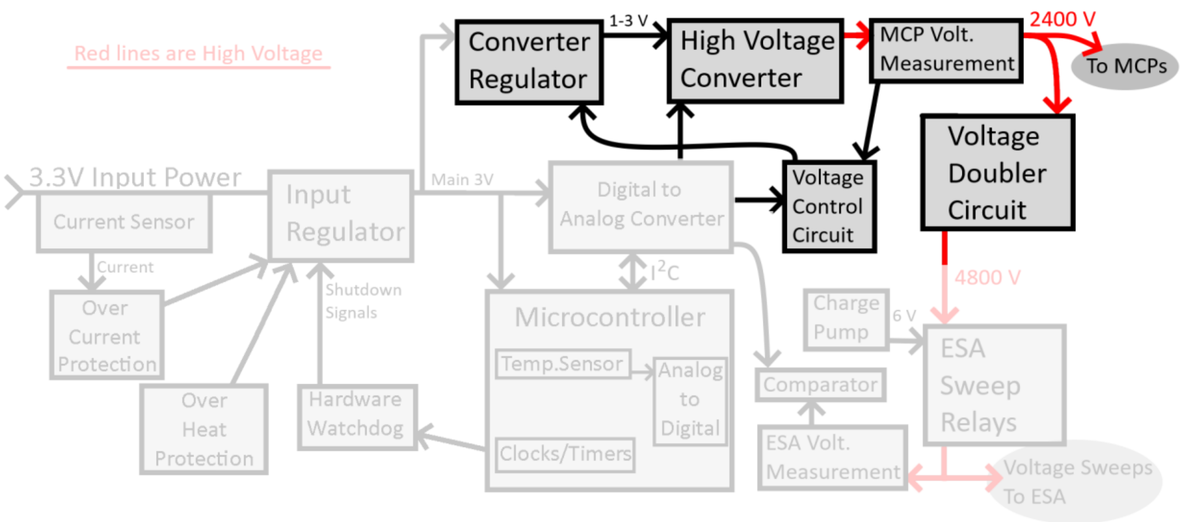

Figure 4.1: MCP Detector Supply and Converter Portions of HVPS Block Diagram

Detector Supply Basics:

As explained in Chapter 1, the detector supply is intended to power a stack of two microchannel plates (MCPs). The ACAPS-specific MCPs have not been designed yet, so I made a flexible design. This means it must produce a wide range of voltages and currents, because different MCPs have different effective resistances and may need different voltages.

Even for a single MCP, a tunable voltage is helpful, because the gain of the plate varies exponentially with the applied input voltage. As a rule of thumb, a 1% voltage increase will cause a 20% gain increase for a 2-stack of MCPs. This means the supply must have very low voltage variations to achieve stable MCP gain. For a spectrometer to produce high-quality scientific data, it is important that the detector gain is reasonably stable throughout the whole length of the mission. At the power supply design review in December 2021, I learned that MCP supplies on large spacecraft aim for 0.1% voltage stability. For this application of a small satellite with minimal resources, I chose a less ambitious target of ±1%, which means ±20 V at 2000 V.

MCPs have a wide range of sizes and resistances. Some are designed with extra-high resistance for lower current consumption, while others use lower resistance to be able to detect more particles at once. MCPs can be made to have resistances from just a few MΩ to over 1 GΩ, so I searched for a narrower range to design for. Examples of three MCP devices from different situation are listed below with their diameter and resistance range.

- The plates used on the Dual Electron Spectrometer of the FPI instrument for the Magnetospheric Multiscale spacecraft are 44-mm diameter and resistances from 18.2-36.4 MΩ [20].
- Hamamatsu manufactures MCPs at near 40-mm diameter with resistances from 2-30 MΩ of the high-dynamic range type, or from 20-200 MΩ for the standard types [21].
- J.L. Wiza's Microchannel Plate Detectors (Reprinted from Nuclear Instruments and Methods Vol 162) specifies 10^{-12} A/cm^2 current density. For a 40-mm diameter MCP at 1000 V, that corresponds to about 100 MΩ.

Based on these three relevant examples, I estimated a range of reasonable MCP resistances for ACAPS to use. ACAPS plans to use 40-mm diameter MCPs, with a central hole cut out. Because of this cutout and the low-power nature of the mission, the resistance for ACAPS MCPs is likely to be higher than the average spacecraft. I am guessing 100 MΩ each, but I also decided to do worst-case design for lower resistances. Considering that the MCP stack will certainly have two plates in series, I designed for a MCP stack resistance of 60 MΩ or greater (30 MΩ per plate). This is reflected in Requirement 2.

Relationship between Supplies:

One major choice was the relationship between the MCP supply and the ESA supply. In the early phases of design (until Summer 2022), I intended to use a 3 kV-rated converter for powering the MCPs, and a separate 6 kV-rated converter for the sweeping supply. After drawing detailed schematics, I realized it would be inefficient to operate two separate converters and that both parts of the circuit could operate off a single converter. This idea had the potential to reduce power consumption, complexity, and mass/volume demands. However, the shared converter added a new need for splitting the supply to two different outputs. I found three methods to split the supply.

1) The first option was to produce about 5000 V from a DC-DC converter for sweeping, and then use a simple resistor voltage divider to give 2500 V output for the MCPs. This simplicity would have come at a terrible cost of power consumption, since at least 50% of the high voltage power would be wasted through resistors instead of MCPs. I quickly started searching for more elegant solutions.
2) The next idea was operating the sweep output at the same voltage as the MCPs. This would be simple but sweeping at only 2000-2500 V would sacrifice part of the upper portion of the ACAPS electron energy range, violating Requirement 1.
3) The third option was to use a charge pump to change voltages. I chose a voltage doubler.

The Voltage Doubler

The voltage doubler is a type of charge pump which allows this design to supply 4000+ V to the ESA, instead of only around 2000 V which the MCPs require, while still using only one high voltage converter. Figure 4.2 below shows all components of the voltage doubler.

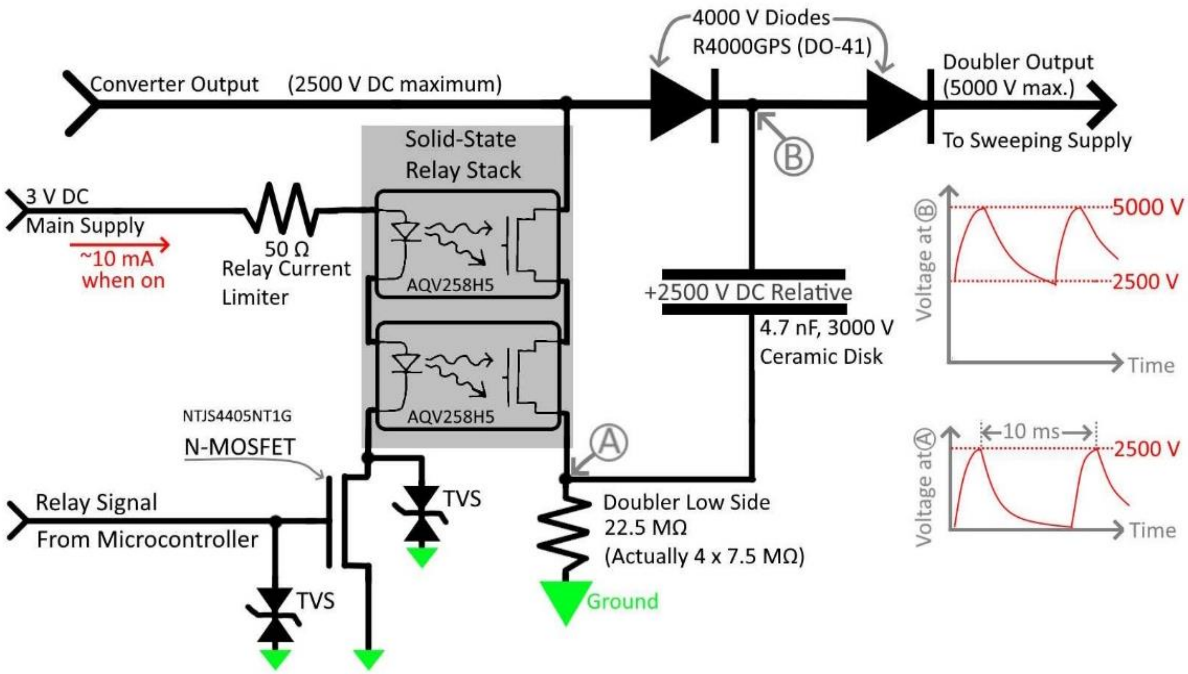

Figure 4.2: High Voltage Doubler Schematic

This design works by switching solid-state relays on and off to produce AC, and then coupling that AC through a capacitor to be added to the converter's DC. That combination of AC superimposed on DC has a peak voltage of twice the input voltage. Diodes are used to prevent that peak voltage from back feeding into the doubler circuit.

These doubler relays in Figure 4.2 do not use closed-loop current control, unlike the ESA relays. This is because the exact timing of the doubler relays switching on and off is not very important. I have verified that a 50-Ω resistor in series with the relays is enough to limit the current to around 10 mA, and it does not need to be any more accurate than that. These doubler relays are controlled by the same type of NTJS4405 n-channel MOSFET as the ESA relays use. However, in this case, the MOSFET is always either on or off, not somewhere in between.

Unlike most charge pumps, this voltage doubler only requires a positive input voltage. It also only requires one stack of relays. However, it is somewhat wasteful, because the low side resistance of 22.5 MΩ dissipates power (110 μA at 2500 V = 280 mW) during the time the relays are on. Fortunately, the relays only have to be on for a very short time to add another pulse of charge to the main doubler capacitor. I have synchronized the doubler relays with the ESA relays, so these doubler relays are on for a little under 1 ms out of every 10-ms sweep. In practice, the doubler increases power consumption of the entire power supply by 10-20%, generally between 30 mW and 70 mW. Avoiding the power wasted by the low-voltage side resistance is possible but requires two more relays. Those extra relays would take up a lot more board space than a string of four resistors.

By simply not turning on the relays, the doubler can be disabled. In that case, current flows straight from the converter through the two diodes to the sweeping supply, and the doubler consumes no power.

What are the voltage doubler's disadvantages?

The primary disadvantage of the voltage doubler is that it takes up a lot of board space. Including the doubler means two doubler relays, two more ESA relays to handle the extra voltage, two large diodes, one large capacitor, and a string of resistors. In total, this takes about 20% of the total circuit board area.

Also, it reduces reliability. Those relays, diodes, and capacitors can all fail in ways which might disable the entire power supply. For example, if the doubler relays fail shorted, they will apply a constant 22.5 MΩ of resistance to the converter's output, which will draw far too much power. Using the doubler to operate at higher voltages opens up more chances of high voltage breakdowns which might ruin the whole circuit. In testing, the doubler has been the source of many failed components.

Why not cut 5000 V in half instead?

It is relatively easy to make a voltage halver, just like a voltage doubler, but it's not a good idea. Charge pumps are more complicated at higher currents, and they introduce more instability to the voltage. The MCP supply needs higher current than the ESA supply, and stability is very important, so I wouldn't want to feed the MCPs power out of a charge pump. For the sweeping supply, the voltage doubler provides a less consistent voltage than the converter would give directly. That is acceptable because the starting voltage of a sweep doesn't particularly matter. The exponential sweep shape will still be the same, even if it starts at a slightly different voltage.

Controlling the High Voltage Converter

Proportional converters, like the A60-P5 which I chose, require some kind of closed-loop control. This is because their output voltage depends on temperature, load resistance, and other factors, and is therefore not fully predictable. Fortunately, there are two different ways to manipulate these converters. The most direct way is by changing the supply input voltage, but there is also a control pin. This control pin takes in a low-current voltage signal and alters the converter's output in response. Figure 4.3 below shows how the control pin alters the output voltage in a dramatic way.

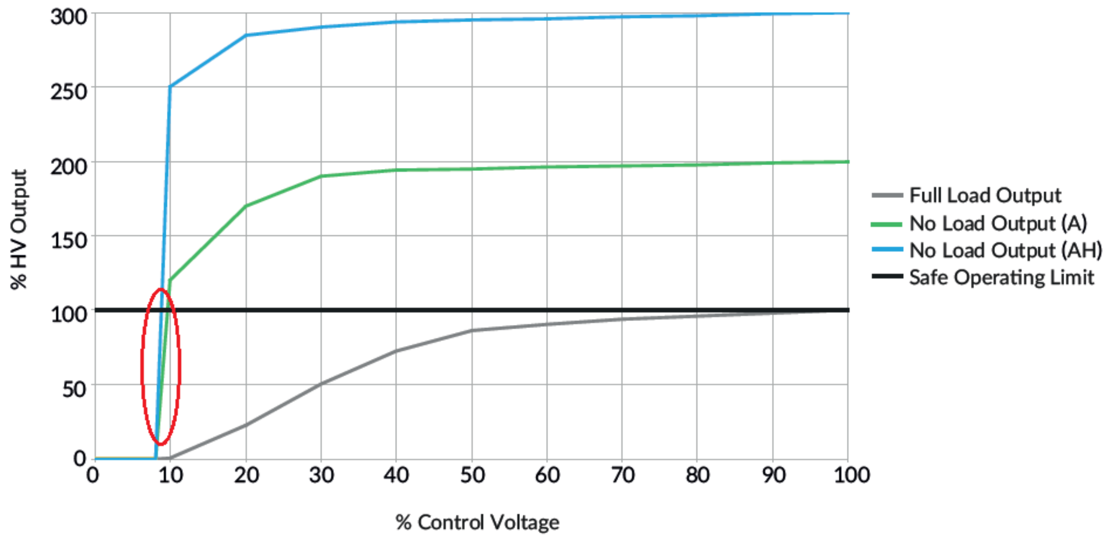

Figure 4.3: Converter Output Voltage versus Control Voltage [11]

The red area in Figure 4.3 highlights that at low control voltage and low load, the converter is extremely sensitive to tiny changes of the control voltage. This extreme sensitivity is frustrating to control, and also potentially dangerous. At low load, the converter's voltage output can easily shoot far above the safe operating limit, even at only 10-20% of the maximum control voltage.

Even worse than the sensitivity, I discovered that there is a region of instability at low load where it is literally impossible to maintain the converter output at a constant voltage. The output voltage will spike too high, but if the control voltage decreases even just slightly, the output drops all the way back to zero. At 5 V input with low load, there is a wide range of outputs up to almost 2000 V that are simply impossible to maintain. Surprisingly, this is still a problem even when the input voltage is reduced from 5 V down to 3 V, although the problem becomes much less severe.

Figure 4.4 on the next page shows the oscillation of an early software-based control loop failing to handle this instability. Figure 4.4 shows that as soon as the converter output voltage rises to the target, the control voltage logically began reducing to avoid an overshoot. However, instead of stabilizing, the converter voltage immediately stops rising and starts falling toward zero again. This oscillation continued forever.

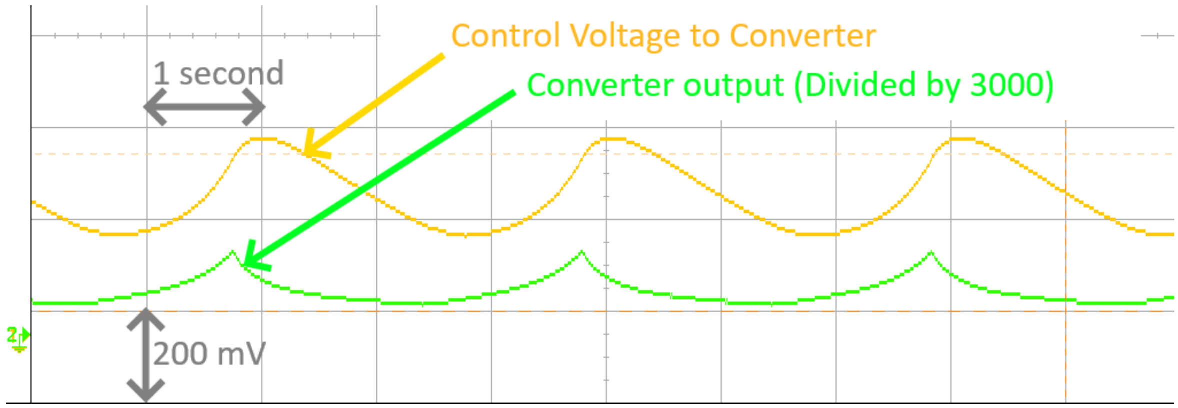

Figure 4.4: Converter Voltage Output Instability from a software control loop

I know this instability is simply a fundamental problem with the converter, not my control algorithm, because I also tried manually adjusting the control voltage at very slow speeds (effectively DC). No matter how slowly the adjustment was made, I found a range of unstable voltages. Even using a completely constant control voltage does not help with the instability, so the only solution to this problem was to decrease the converter's input voltage. While the converter was designed for +5 V input, my control circuit generally operates it between 1.5 and 2.8 V. At low load, this is conveniently more efficient than running at 5 V input. It achieves stable output by starting at only 3 V input, and then using a linear regulator to reduce the input voltage further.

To help with designing a controller that works efficiently, I measured converter performance at a wide range of loads, temperatures, control voltages, and input voltages. I made the following conclusions:

- Whenever possible, the control pin should be kept around 0.6-0.8 V. Anything under about 0.7 V might not turn the converter on at all. Meanwhile, any control voltage beyond 0.7 V is always less efficient, so it should only be increased when there are no other options. Unfortunately, this threshold changes with temperature, so I used 0.9 V to be safe.
- Higher control voltages are required at low temperatures, assuming all other conditions are equal.
- The converter is much more efficient with lower resistance loads (higher current). At best I measured efficiencies around 30-35% at 3 V supply voltage and 0.6 V control voltage with a 53 MΩ load. In other words from an efficiency perspective, it is advantageous to operate at the bottom right ends of the red lines (high supply voltage and low control voltage) in Figure 4.5.
- There will always be a minimum output voltage, below which the converter is unstable. Reducing the input voltage only reduces where that stability minimum lies but does not eliminate it.

Figure 4.5 shows a set of measurements I made for 300 MΩ and 53 MΩ at various control and input voltages. The purpose of these measurements was to see how the converter's characteristics change at different loads and different input voltages. The red lines are different output voltages.

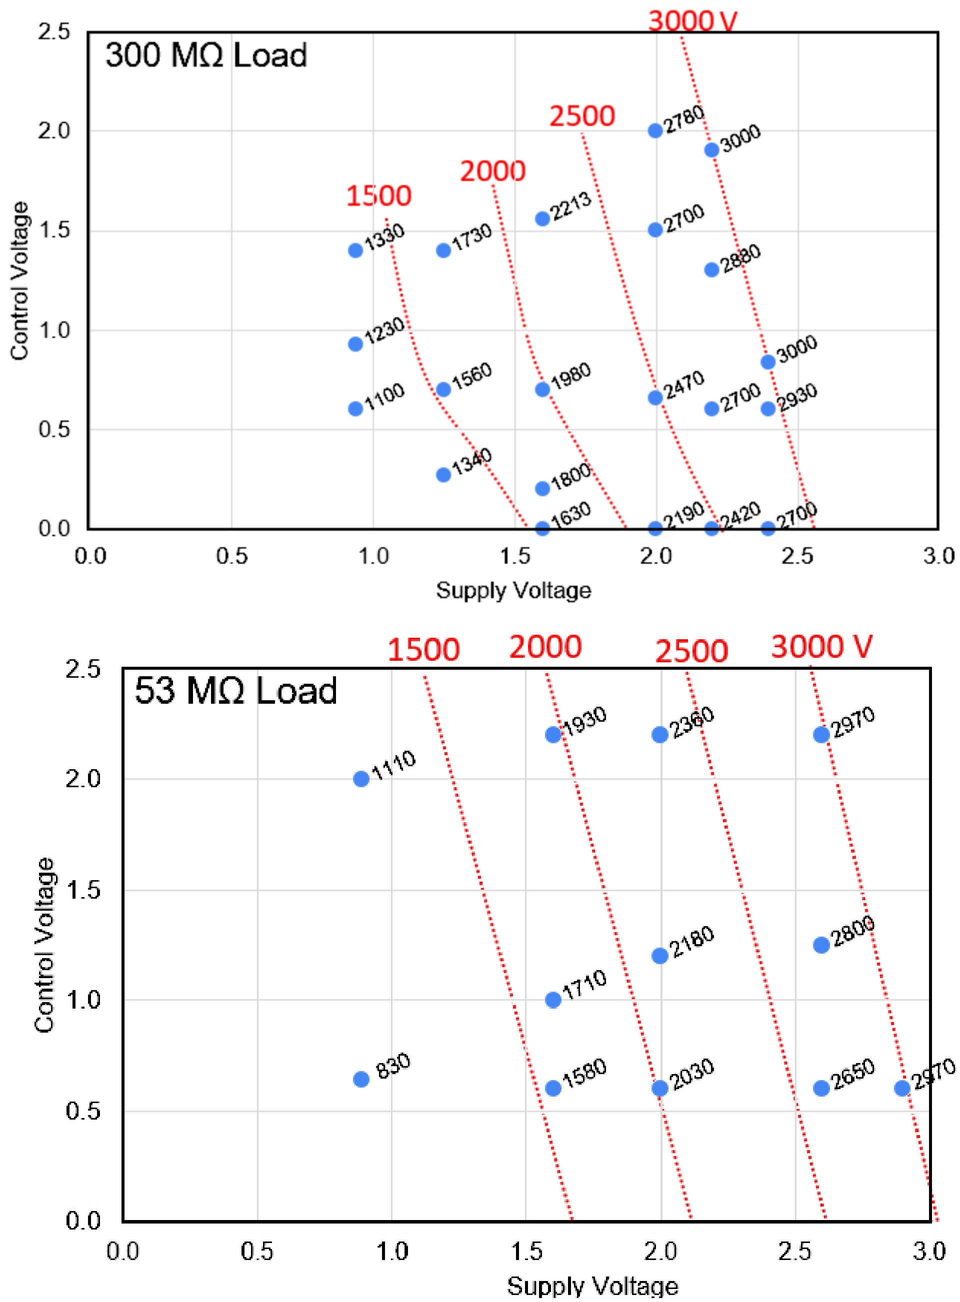

Figure 4.5: Converter output voltage (blue dots with red interpolation lines) for various control and supply voltage points at 300 MΩ load (left) and 53 MΩ load (right).

Figure 4.6 below shows the entire circuit which manages the high voltage converter.

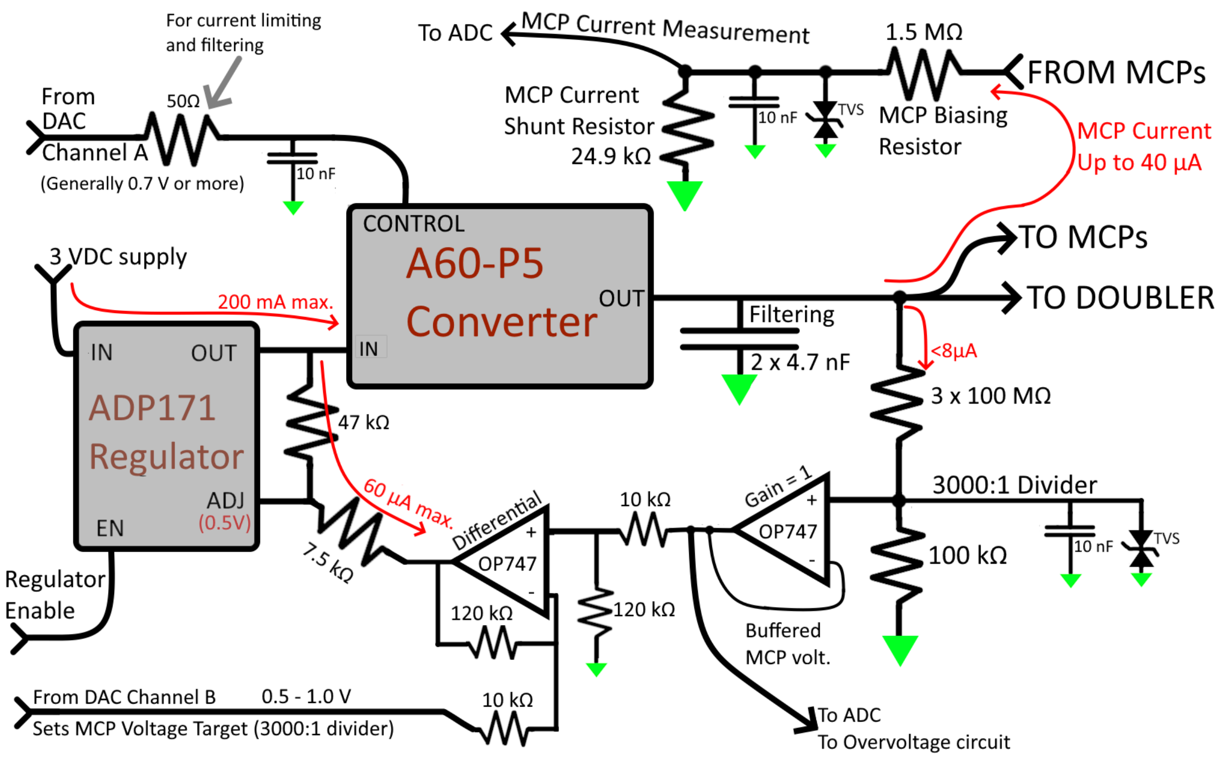

Figure 4.6: HVPS Converter Control Circuit Schematic

A digital-to-analog converter directly provides a voltage to the converter's control pin. The control pin is not part of any closed feedback loops. However, the converter's input pin is part of a feedback loop, managed by an ADP171 linear regulator (Figure 4.7 below). It takes in the 3 V main supply and dumps excess voltage until the desired voltage for the converter is reached. I chose the ADP171 because of its small size and two important features; an adjustment (ADJ) pin and an enable (EN) pin. The enable pin is used to cut off all power to the high voltage converter, which is useful for activating overvoltage protection, or for using software control to turn off the converter.

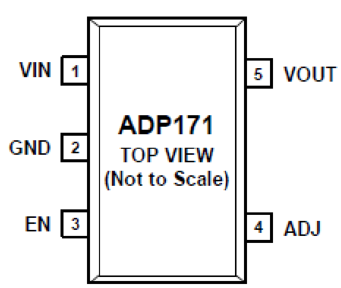

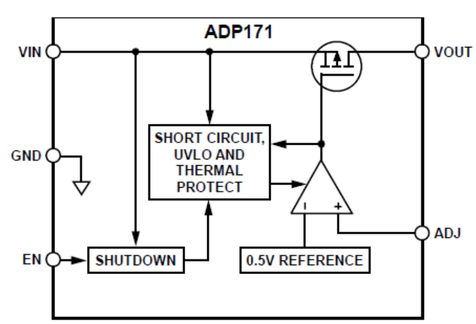

Figure 4.7: ADP171 Regulator Pinout (left) and Schematic (right) [22]

The ADJ pin is used to tune the converter with a feedback loop, as seen in Figure 4.6. The ADP171's ADJ pin always aims to be at 0.5 V. In other words, if the ADJ pin is less than 0.5 V, the ADP171 will allow more voltage through, until the ADJ pin voltage rises.

After the MCP voltage is produced by the converter, it is passed through a 3000:1 divider, and is buffered by an OP747 amplifier. The buffered MCP voltage is used for several things, including overvoltage protection and measurement by software with an ADC. Most importantly, this buffered MCP voltage feeds back into a control loop which manages the converter.

The other OP747 in Figure 4.6 acts as a differential amplifier with a gain of 12. This differential amplifier compares the measured MCP voltage to a target voltage, which is set by a digital-to-analog converter (DAC). The gain of 12 was found mostly through trial and error. It is not exact. A lower gain will cause the MCP voltage to be able to overshoot the target more. As it is, this design allows overshoots of around 20 V. If the gain was infinite, overshoots would be reduced, but it might oscillate.

Example Step-by-Step Converter Control Loop:

For reliability and consistency, I chose a control loop design that does not include any software in the loop. This decision came after months of struggling with direct control by software. I eventually decided that software is just not well suited for the task, mostly because I had constant trouble with overshooting, oscillations, and slow response.

The core parts of the feedback loop are a differential amplifier and the regulator's ADJ pin. A step-by-step description of the control loop function is as follows:

1. The target voltage is set by a DAC (Digital to Analog Converter) based on the 3000:1 MCP divider ratio. So, if 2200 V is desired, the target voltage would be digitally set to 0.733 V.
2. When the power supply is starting up, the divided MCP voltage will still be below that target of 0.733 V. In that case, the differential amplifier will output 0 V as it cannot go negative.
3. The resistors connected at the ADJ pin form a voltage divider from the regulator's output, with a ratio of 7.27:1. Since the ADJ pin is always trying to reach 0.5 V, it will try to reach an output voltage of 0.5 V x 7.27 = 3.63 V. However, the ADP171 is only supplied with 3 V, so that is all it can provide to the converter.
4. Later, once the MCP voltage starts to go above its target, the differential amplifier will start to output a positive voltage. For example, if the MCP voltage overshoots by 10 V, the divided voltage error is +3.3 mV. The amplifier multiplies this error by 12 to give an output of +40 mV instead of 0.

5. Unlike most voltage dividers, the low voltage side of this ADJ pin divider does not connect to ground. Instead, it connects to the output of that differential amplifier. Eventually the increasing amplifier output will push up the resulting output voltage of the voltage divider, and it will go above the ADJ pin threshold of 0.5 V.
6. Once the ADJ voltage goes above 0.5 V, the ADP171 regulator begins to cut off some voltage to the converter. The regulator will continue to reduce its own output voltage until the ADJ pin drops back to 0.5 V.
7. With less input voltage coming from the regulator, the converter cannot output as much, so the MCP voltage gradually falls in response.
8. Eventually, equilibrium is reached, and the MCP voltage will be equal to the target. The differential amplifier's output will settle at whatever is necessary to keep the ADJ pin at 0.5 V.

I have not had any issues with oscillation on this design, while I did have issues when software was in the loop. The converter's output responds so slowly that it is practically constant in comparison to the fast response time of the regulator and amplifier. The OP747 amplifier is relatively slow, but still achieves a gain of 12 up to around 30 kHz [23]. The regulator's internal amplifier is much faster, and I intentionally did not add any capacitors to any part of that voltage divider.

Filtering

The largest disadvantage of deciding to share a single converter between the two power supplies is voltage ripple. The sweeping supply is a very inconsistent load, compared to the MCP supply. The sweeping supply demands pulses of current each time the relays turn on. Every time the relays turn on, this sudden removal of charge causes a small voltage ripple on the MCP supply. Meanwhile the MCP stack is basically a constant resistance, and it demands a constant voltage to avoid variations of MCP gain. The only way to reconcile these conflicting supplies is to heavily filter the MCP supply, to hopefully minimize sweeping-induced ripple. Here, the concept of an RC time constant appears again, but now with the goal of making RC as long as possible. When filtering a DC voltage supply, using more capacitance is helpful to achieve a smoother voltage with less noise and ripple.

Analyzing in terms of charge is relatively simple. When relays turn on, the ESA sweeping supply's capacitance withdraws charge from the MCP supply's capacitance. This causes the MCP voltage to drop momentarily. Keeping voltage ripple on the MCP supply's large filtering capacitor below 1% requires the MCP supply to have at least 100x more capacitance than the ESA. Conservatively assuming the sweep capacitance is 100 pF (50 pF but effectively 100 because of the voltage doubler), the filtering capacitor would need to be 10 nF. I did manage to fit two parallel 4.7 nF disk capacitors (9.4 nF) on the converter's output, and I know it is possible to add at least one more capacitor in parallel to reach 14 nF. Assuming a worst-case average load of 30 MΩ is seen by the converter, the MCP supply's RC time constant would be 0.28 seconds with 10 nF. That is a cutoff frequency of about 0.6 Hz, so any noise happening at higher frequencies should be mostly filtered out.

Filter Capacitor Selection

For initial testing with the high voltage relays, I used film capacitors rated to 6 kV. Film capacitors are simply alternating layers of metal foil and plastic film packed closely together. They are very stable across temperature and voltage because the plastic film is a relatively ideal dielectric material. Film capacitors did not cause any problems, but they are unfortunately not very volume efficient. Fortunately, I found that ceramic disk capacitors have much higher energy storage density, easily by a factor of 2. However, these ceramic capacitors are less accurately sized and less stable across temperature. This is acceptable though because the exact value of a filtering capacitor is not very sensitive. If the capacitance changes by 10%, that simply means it will be 10% more or less effective at filtering. I do not care about 10% uncertainty, because the filtering capacitors are simply designed with as much capacitance as possible.

Chapter 5. Microcontroller and Associated Electronics

This chapter discusses hardware and software including the microcontroller, input regulator, and digital-to-analog converter as indicated in Figure 5.1 that perform the following functions:

- Sampling voltages and currents.
- Turning on and off features of the circuit.
- Creating analog signals from digital commands.
- Maintaining accurate timing of the sweeps.

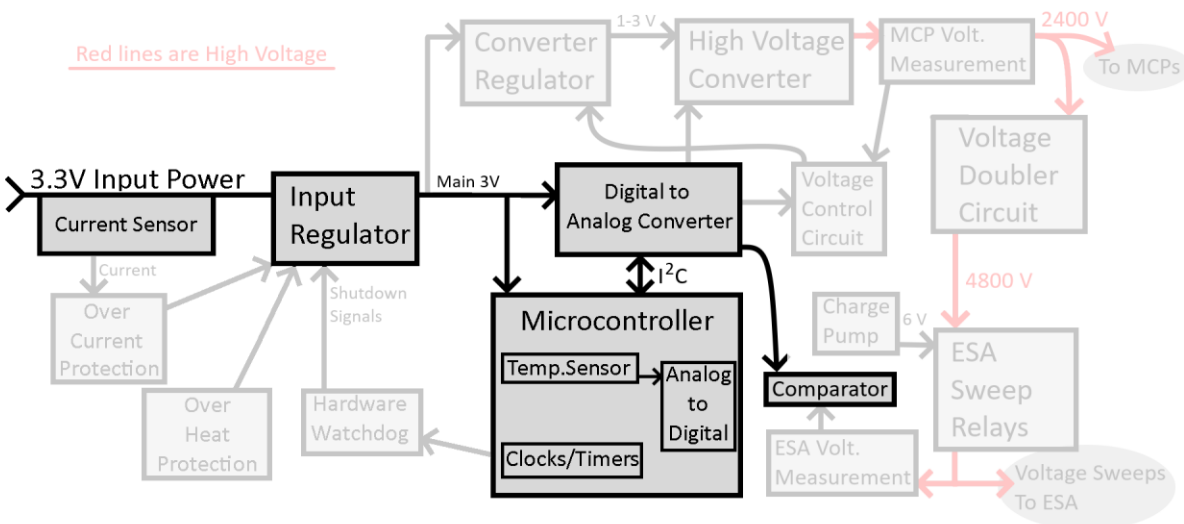

Figure 5.1: Microcontroller, Input Regulator, and Digital-to-Analog Converter Portions of HVPS

Microcontroller Selection

The Space Systems Engineering Lab uses Texas Instruments MSP430 microcontrollers for CubeSat projects. I decided to use those microcontrollers on this thesis for simplicity, and because they are well suited to small, low-power circuits. I gained relevant experience using the MSP430 in a class called *EE608: Power Electronics Design*, taught by Dr. Richard Wies at UAF. That class was extremely helpful in preparing me to design and test the HVPS in this thesis, even though the class and laboratory were not geared toward high voltage applications.

MSP430s are similar to the Atmel chips used on popular Arduino circuit boards. MSP430s are all designed for very low current operation, around 1 mA at 2 or 3 V when fully operating and just a few µA in certain low power modes. This makes them popular for spacecraft and other very demanding low-power situations. The MSP430 has been used for non-educational flight missions, including the Jet Propulsion Laboratory's MarCO CubeSats which flew through interplanetary space to Mars [24]. The MSP430's 20-year-old technology has gained the trust of many in the CubeSat community. They are not explicitly designed for radiation tolerance, but certain FRAM-equipped models have performed well in radiation tests [25].

What can the microcontroller do?

All of the MSP430 microcontrollers have far more than enough processing power even at 1 MHz clock speed to manage a power supply. It is a 16-bit CPU with a clock speed of up to 16-24 MHz. In addition to the CPU, the chip has a few built-in peripherals which were very helpful for this HVPS. The MSP430 microcontroller includes:

- A 10- or 12-bit **analog-to-digital converter** (ADC), which is critical for digitizing any measurements. This could be a separate IC, but it is more convenient to use the built-in functionality.
- Multiple **serial communication** modules. These communications are critical for interacting with other ICs on the HVPS board as well as the outside world.
- A **clock system** with built-in oscillators for high and low frequencies, as well as the ability to use a 32 kHz "watch crystal" type oscillator for better accuracy.
- Multiple **timers** which are very useful for automatically producing waveforms based on clock cycles without needing direct software control.

MSP430 Selection

The standard microcontroller used on most of the SSEP lab's projects is the MSP430F6779A. It is a 100-pin configuration, which consumes roughly 18 mm x 18 mm of circuit board surface area. For controlling just a power supply this microcontroller is excessive in terms of the size and number of pins, so I chose a much smaller microcontroller. At first, for ease of prototyping, I used the MSP430G2553, an older, simpler microcontroller with only 20 pins. It uses the DIP package size which is easy to work with on the LaunchPad development boards, but it is too large to be practical on the real board design. On the first board revision, I switched to using the TSSOP packaged MSP430G2553, which is far smaller than the DIP package, yet allows 38 pins instead of 20 because they are more closely spaced.

During the first board revision, I developed an interest in radiation effects on microcontrollers. It is a huge concern because a single corrupted line of code in memory could render the entire power supply useless. I found a very appealing way to reduce risk by using a MSP430 with Ferroelectric Random Access Memory (FRAM). FRAM is expensive and only practical to use in small quantities, but it is exceptionally resilient and stable. Compared to FLASH and ordinary RAM, FRAM is far less prone to bits being flipped (corrupted) by high energy radiation.

Unlike the more common SRAM or DRAM, FRAM is non-volatile, meaning it retains data when power is cut off. FRAM has practically infinite (>10^{12}) read and write cycles before it may be degraded. This means FRAM can effectively take the place of all types of memory, but only in quantities of a few megabytes or less, because it is very low-density compared to FLASH. Fortunately, the code for a power supply does not need very much memory, so FRAM was easily usable.

I learned from Texas Instruments that FRAM-equipped MSP430s use built-in error correcting codes. This means that a single bit flip is not immediately a problem. This correcting code is more advanced than standard error correction because it is double-error correcting, triple-error detecting (DECTED) [25]. This applies to each 64-bit section of memory. In other words, two bits out of every 64 can become corrupted by radiation without any real impact.

I could not see any disadvantage to using a FRAM-equipped microcontroller, so I used one for the second and third board revisions. I selected the MSP430FR5967. The main reason for selecting this exact model is that Texas Instruments produces a space-qualified variant called MSP430FR5969-SP. The normal chip costs $10, while the space-qualified chip costs around $1000. I wanted to allow for drop-in replacement with the space-qualified model whenever it was needed.

Figure 5.2 shows the overall arrangement of microcontroller and digital to analog converter (DAC).

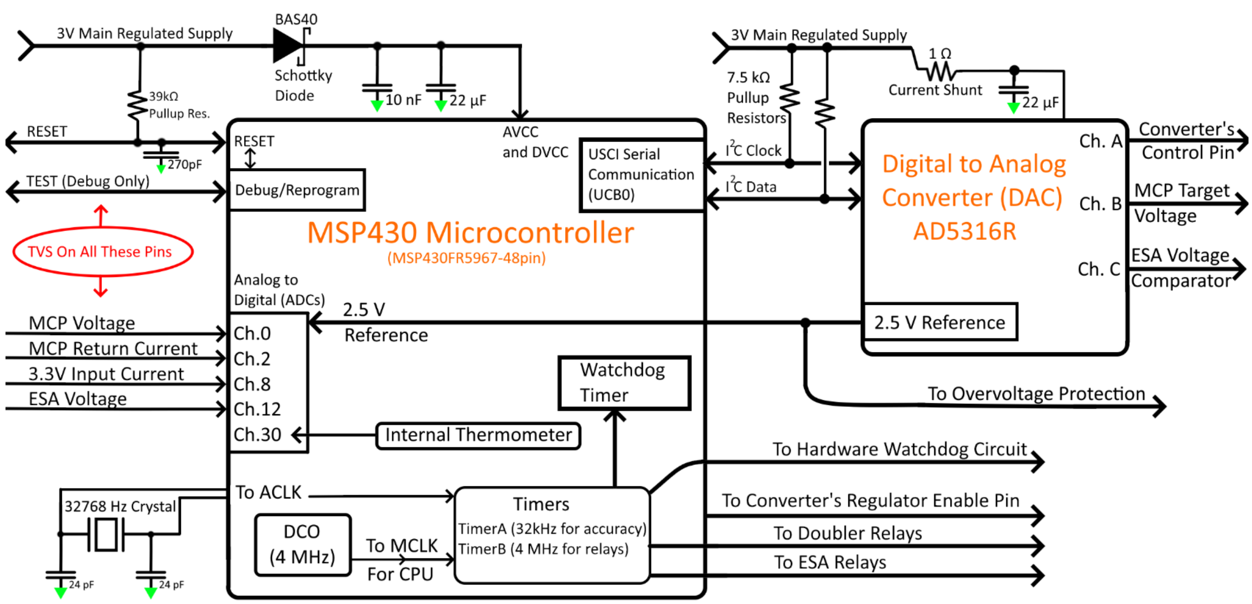

Figure 5.2: Microcontroller and DAC schematic

What is the Schottky diode for?

The Schottky diode in the path of the 3 V regulated supply (upper left in Figure 5.2) allows the MSP430 to have an optional independent power source for a short time (<10 ms). A Schottky diode is an electrical one-way valve just like a regular diode, except the Schottky diode has only about 0.2 V of voltage drop. This means the MSP430 still sees about 2.8 V, well above its 2.5 V minimum for full ADC functionality. This power input is also filtered by separate capacitors. Also importantly, it allows the MSP430 to be powered without powering the entire HVPS board. For reprogramming and debugging, I was able to provide power downstream of the Schottky diode for only the MSP430 without letting power flow to any other parts of the circuit board.

Why use a large and small capacitor together?

While it might seem like the 3-V 10-nF input capacitor is pointless when it is in parallel with a 22-µF capacitor that is 2200x larger, they serve different purposes. The 10-nF capacitor is much better at filtering higher-frequency noise. That is because a large 22-µF ceramic capacitor will have a relatively low resonant frequency, which means it actually fails to short high frequency noise to ground. This filtering technique is recommended by many datasheets.

Microcontroller Clocks

There are several oscillators built in to the MSP430. The fastest one is the Digitally Controlled Oscillator (DCO), and it can run at speeds between 1 and 24 MHz [26]. This DCO runs the main CPU clock, and I used it for timing the relays and nearly everything else that happens during a voltage sweep. I decided to operate it at 4 MHz because that worked well for relay timing. Having a very fast clock speed is not particularly important since even at 4 MHz there is a lot of "dead time" where the CPU has nothing to do. Operating at 4 MHz means it gives 0.25 µs resolution per clock cycle. In other words, each 0.01 sec sweep is about 40,000 DCO clock cycles long. The DCO is easy to work with because it is entirely within the microcontroller and does not require much configuration. However, the DCO is not very accurate or stable. Its frequency is only specified as accurate within ±3.5%, and it has 100 ppm/°C frequency drift with temperature. [26]

I wanted a more stable oscillator so that a 10 ms sweep would always be really 10 ms long. Eventually, accurate sweep timing will be important for synchronizing with the detector electronics. For this need, I took advantage of the fact that the MSP430 is designed to use an external crystal oscillator. An external crystal is many times more accurate than the DCO. On the final board revision, I measured around 100 ppm of frequency error across the entire temperature range (about 100x better than the DCO).

The crystal is very simple to implement, basically a watch crystal soldered to the circuit board with two capacitors, shown in the lower left of Figure 5.2. This crystal is much slower than the DCO at only 32.768 kHz, which means each crystal clock cycle is about 30 µs long. That is too coarse for relay timing, so the crystal cannot be used directly. Instead, I adjusted the number of DCO clock cycles in software, based on the crystal. The code adjusts the number of DCO cycles which defines the length of a sweep. It makes timing adjustments every 100 sweeps (1 sec), and those adjustments are precise to within 1/32768 = 30 ppm. This adjustment method allows the sweep timing to take advantage of the long-term accuracy of the crystal, combined with the short-time resolution of the DCO. Because of the 1 sec interval between adjustments, relay timing takes time after startup to become accurate. Within the first minute after startup, the timing may be relatively bad, especially at extreme temperatures.

Relay Timing

At first, I controlled the ESA sweep relays directly from a purpose-built timer in the MSP430. Once the timer was configured, it would endlessly switch the relays on and off, regardless of whether the CPU was busy or not. I used direct output from the timer because I was worried about software "lagging" and causing an inconsistent delay before the relay would switch. The time delay for software was a few µs at worst, so it turned out not to be a problem. Later, I decided to use software control with interrupts from the MSP430’s timers. The method using interrupts was more flexible in terms of which pins on the MSP430 microcontroller that I used for each of the timing signals. Interrupts allow the CPU to be involved, which is not necessary for relay switching, but the CPU is necessary for other things like ADC sampling.

Analog to Digital Converter (ADC)

The MSP430’s built-in 12-bit ADC works in four steps:

1. The ADC must be configured to measure voltage from a certain pin on the microcontroller. It can only measure one voltage at a time. The sampling time and sampling method are also configured.
2. The ADC is enabled, and sampling begins. I used the longest available sampling time of 150 µs for optimal accuracy. There is still plenty of time available to make five measurements during a sweep at this speed. Because the ADC is separate from the CPU core of the microcontroller, the ADC can be sampling while the CPU does other calculations.
3. When the ADC has finished sampling, it saves the raw ADC result (a 12-bit integer) to memory. This triggers an interrupt for the CPU.

4. Finally, the CPU does calculations on the raw ADC result, converting it to a meaningful number like a voltage, temperature, or current. After this, the CPU will go back to being idle until another interrupt arrives.

Figure 5.3 shows an oscilloscope measurement of the time taken for this whole ADC process. The green trace shows a debugging voltage, which toggles high or low whenever each step of the process is completed.

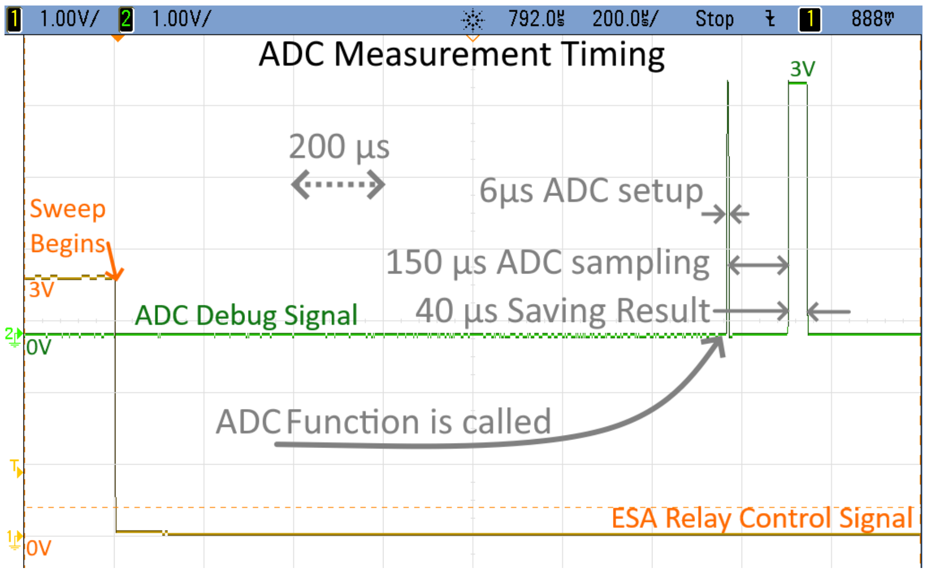

Figure 5.3: ADC timing oscilloscope plot

Jitter in this case means the variations in total time required to do an ADC measurement. The jitter of the ADC is only about ±1 µs. That is only ± 4 CPU clock cycles, or 1/10,000th of a full ESA voltage sweep.

Digital to Analog Converter (DAC)

Digital to analog conversion is necessary for controlling the high voltage converter. In this design, it is needed for direct input to the converter's control pin, as well as setting a target MCP voltage for the converter control feedback loop. Later, I added a third function for the DAC, which is setting a threshold for the ESA voltage comparator.

While an ADC is built in to the MSP430, a DAC is not. I briefly shopped for various DACs and chose the Analog Devices AD5316R (see Figure 5.4 below) because of four main features.

- It has a built-in 2.5-V voltage reference that is very stable. This means I don't have to use a separate voltage reference component. This AD5316R's temperature coefficient is typically only 2 ppm/°C and the voltage drift after 1000+ hours of operation is around 100 ppm [27]. Note that the MSP430 also has a built-in reference, but it is not sufficiently stable.

- It has 10 bits of resolution on 4 separate channels in a single small package.
- It communicates using I^2C, a protocol that I already understood from prior work.
- It has a built-in output amplifier with a gain of 1 or 2.

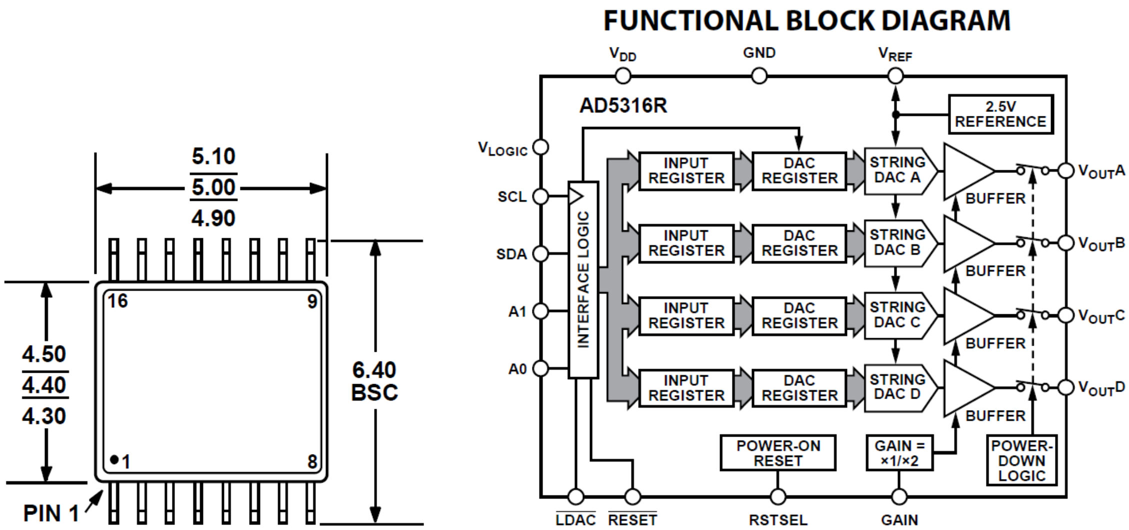

Figure 5.4: AD5316 DAC Package Layout (left) and Pinout with Block Diagram (right) [27]

Like any DAC, the AD5316R takes in a series of bits and converts it into a voltage, in this case between 0 and 2.5 V. With 10 bits of resolution, that range of 2.5 V is divided into 2^{10} or 1024 steps. This means each bit of DAC input value is worth 2.44 mV. In hindsight, I probably should have used a 12- or 14-bit DAC to get finer precision. A 2.44 mV step might sound small enough, but that corresponds to about 8 V on the high side of a 3000:1 voltage divider.

The difficult problem with using the DAC was digital communication. It uses the Inter-Integrated Circuit (I^2C) protocol, which has one clock wire and one data wire. I used the MSP430's Unified Serial Communication Interface (USCI) to send voltage commands to the DAC. The communication is one-way. However, the DAC does need to acknowledge every byte that the MSP430 sends. Figure 5.5 below shows a single I^2C transaction of 4 bytes lasting about 500 µs. I used the "standard mode" I^2C clock frequency of about 100 kHz on the final design.

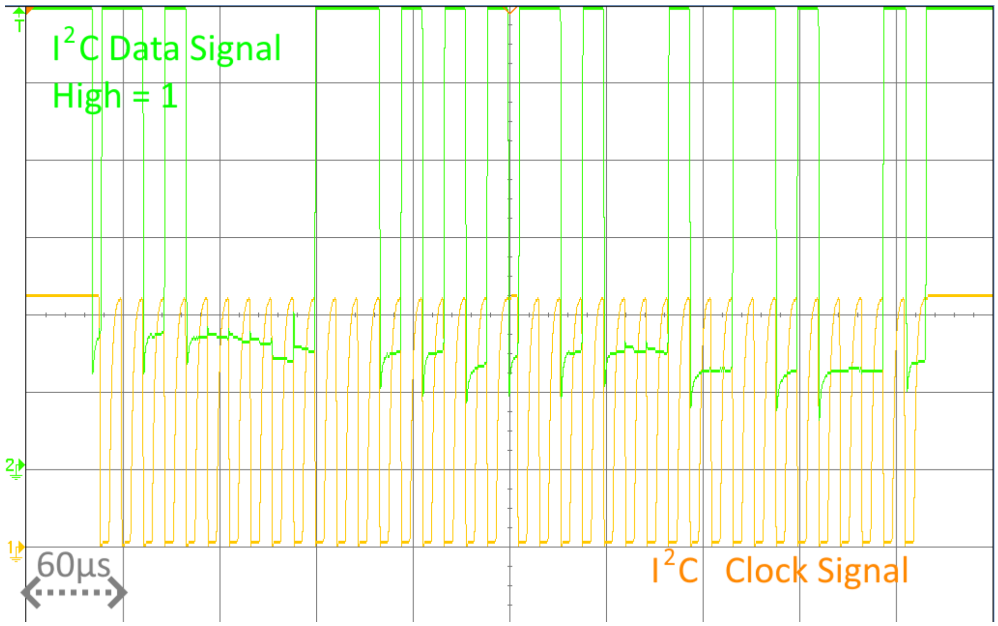

Figure 5.5: I²C Data and Clock Signal Illustrating a Transaction

Software Requirements

The main requirements of the software are:

1. The code must not be able to get stuck in a loop forever.
2. Complex operations like the I²C conversation are managed by a state machine so that the microcontroller always remembers where it is in the process.
3. Minimal time will be spent inside of interrupts, therefore there can be no nested interrupts. The CPU should never stay busy continuously for more than ~100 µs.
4. Use integers for everything because floating-point math is computationally slow and confusing.

The Software

I programmed the MSP430 microcontroller using the C programming language in Code Composer Studio from Texas Instruments. I wrote all of the code as a single .c file, without any libraries. The entire code is found in Appendix 4.

An entire voltage sweep is 10 ms long, which is about 40,000 clock cycles of the 4 MHz main CPU clock. Table 5.1 below explains the sequence of events that the code follows.

Table 5.1: Main Software Loop Timeline for a single sweep

Time (ms)	Events (in order)	Cause/Conditions
TB0CCR0 CYCLE STARTS 0 ms	ESA Relay MOSFET turns off. Doubler relay MOSFET turns off. Watchdog Timer is kicked. Sweep Counter goes up by 1.	Defined by MSP430's Timer B CCR0.
	Clock crystal adjustment code is run. Relay timing adjustment code is run. ESA voltage ADC sampling begins.	Relay timing is software-adjusted based on temperature and the voltage needed. Currently it is tuned so that the ESA and Doubler relays are synchronized, but it does not have to be this way.
	High voltage is enabled or not.	Will turn off high voltage if too hot or cold.
	Target voltage DAC value is calculated. Control voltage DAC value is calculated.	These haven't been sent to DAC yet.
0.03-0.15	Both sets of relays gradually turn off. ESA voltage starts to decrease.	Relay turn off time is variable, depending on temperature.
~0.20	ESA voltage ADC is measured and saved. MCP voltage ADC sampling begins.	ADC sampling takes about 0.2 ms.
TB0CCR2 0.25	DAC set channel B function is called.	This sets the target voltage for the MCP supply
~0.40	MCP voltage ADC is measured and saved. MCP current ADC sampling begins.	
0.5?	Comparator goes from low to high voltage because the ESA voltage dropped enough.	This depends on where exactly the comparator threshold was set.
~0.60	MCP current ADC is measured and saved.	
TB0CCR4 1.0	Temperature ADC sampling begins.	
~1.2	Temperature ADC is measured and saved. 3.3 V current ADC sampling begins.	ADC sampling takes about 0.2 ms.
~1.4	3.3 V current ADC is measured and saved.	ADC sampling takes about 0.2 ms.
TB0CCR5 2.5	DAC set channel A function is called.	I2C transaction timing: Unknown
TB0CCR1 9.5	The sweep has ended. ESA relays MOSFET turns on. Doubler relay MOSFET turns on.	
	DAC set channel C function is called.	This sets the ESA voltage comparator threshold.
9.6-9.8	Relays gradually turn on. ESA voltage rapidly rises.	Relay turn on time is somewhat unpredictable, depending on temperature.
TB0CCR0 10.0	The cycle repeats. Beginning of a new sweep. 0.01 seconds have passed.	

Microcontroller Pins

Table 5.2 below lists the MSP430FR5967 microcontroller pins function and purpose. Green pins are related to analog measurements, red pins are digital outputs, and blue pins are for communications.

Table 5.2: MSP430FR5967 Microcontroller Pins Function, Signal Type, and Purpose

Pin	Selected Pin Function	Type of Signal	Purpose
1	Analog 0	Analog DC input	Detector voltage: divided by 3000 and buffered
2	V_eRef+	Analog DC input	Stable +2.5 V reference from AD5316R Digital-Analog Converter
3	Analog 2	Analog input	Detector current shunt – from INA180 amp.
4	Analog 12	Analog input	ESA voltage – divided by 2652 and buffered
5	P3.1		Unused
6	P3.2		Unused
7	P3.3		Unused
8	P4.7	Digital Output	Enable pin for high voltage. 1 means high voltage on.
9	P1.3	Digital Output	ESA relay control signal. 1 means relays turn on.
10	P1.4		Unused
11	P1.5		Unused
12	PJ.0		Unused
13	PJ.1	Digital Clock	Unused, but I^2C SCL (Clock) passes through this pin
14	PJ.2	Digital Serial Data	Unused, but I^2C SDA (Data) passes through this pin
15	PJ.3		Unused
16	P4.0	Analog Input	Board input current shunt measurement
17	P4.1	Digital Output	Watchdog timer-toggles whenever the timer is reset.
18	P4.2		Unused
19	P4.3		Unused
20	UART Transmit	Digital Output	Communication from MSP to outside
21	UART Receive	Digital Input	Communication to MSP from outside
22	SBW TEST	Digital Debugging	Used for reprogramming/debugging
23	RESET	Digital Debugging	When low, this resets the chip. Also used for reprogramming
24	P2.0		Unused
25	P2.1		Unused
26	P2.2		Unused
27	P3.4		Unused
28	P3.5		Unused
29	P3.6		Unused
30	P3.7		Unused
31	I2C Data (UCB0SDA)	Digital Serial Data	For communication to the AD5316 DAC
32	I2C Clock (UCB0SCL)	Digital Clock	For communication to the AD5316 DAC
33	P4.4	Digital Output	Unused, but ESA Relay control signal passes through this.
34	P4.5		Unused
35	P4.6		Unused, but ESA voltage divider signal passes through this.
36	**GROUND**	**GROUND**	
37	**MSP SUPPLY**	**POWER**	**Approximately 3 V**
38	P2.7		Unused
39	P2.3		Unused
40	P2.4	Digital Output	Doubler relay control signal. 1 means relays turn on.
41	**GROUND**	**GROUND**	
42	PJ.6		Unused
43	PJ.7		Unused
44	**GROUND**	**GROUND**	
45	LFXIN	Analog Oscillator	Input to the 32.768 kHz crystal
46	LFXOUT	Analog Oscillator	Output of the 32.768 kHz crystal
47	**GROUND**	**GROUND**	
48	**MSP SUPPLY**	**POWER**	**Approximately 3 V**

Input Voltages

During the early prototyping process all the way through the first circuit board prototype, I used both 3.3 V and 5.0 V as inputs to the board. Having both voltages is common on CubeSat hardware and on computers in general. In some ways, two voltages were convenient, in other ways it was dangerous and frustrating.

- Two input voltages means two sets of components for the inputs for the following functions: current limiting, over/under voltage protection, power cycling, and filtering. That takes up more board space and increases the risks of more failures.
- Needing to route both 3.3 V and 5 V traces on a circuit board creates extra complexity. Sometimes it is even necessary for the same IC to use both 3.3 V and 5 V. The traces often have to cross paths.
- Two supplies are more complex, both in design and testing. There are many opportunities for error of connecting the wrong supply or forgetting a supply. For example, I have burned up at least two microcontrollers by accidentally applying 5 V to a pin that should only see 3.3 V because of a wiring error.
- In cases of component failures or shorts on the board, 5 V can potentially leak into the 3.3 V area.

For simplicity and safety, after the second board revision in September 2022 I chose to operate the entire board from 3.3 V only. As explained in Chapter 3, this had one complication. I had to generate 6 V with a charge pump IC, the TC1240A. That was only necessary for the ESA relays. Everything else on the board works at 3 V.

Input Voltage Regulation

The main supply voltages of a spacecraft are susceptible to variable loads and noise, which can cause voltage fluctuations of around ±1%. A typical mitigation strategy is to regulate the voltage at the board level. I have chosen the ADP171 voltage regulator for this purpose, the identical regulator that I already used to manage the high-voltage converter’s input.

None of the ICs on this board will be damaged by slightly higher input voltages like 3.40 V, but slightly low voltages are a concern. 3.0 V is the minimum for the most sensitive components like the OP747 amplifiers. Assuming a minimum input of 3.2 V, there is only 0.2 V of margin available for the main voltage regulator to work within. This is called dropout voltage, or the minimum voltage drop from the regulator’s input to output. Thankfully, modern voltage regulators have impressively low dropout voltages, sometimes below 0.1 V. Figure 5.6 below shows the ADP171 voltage regulator's dropout voltage versus load current.

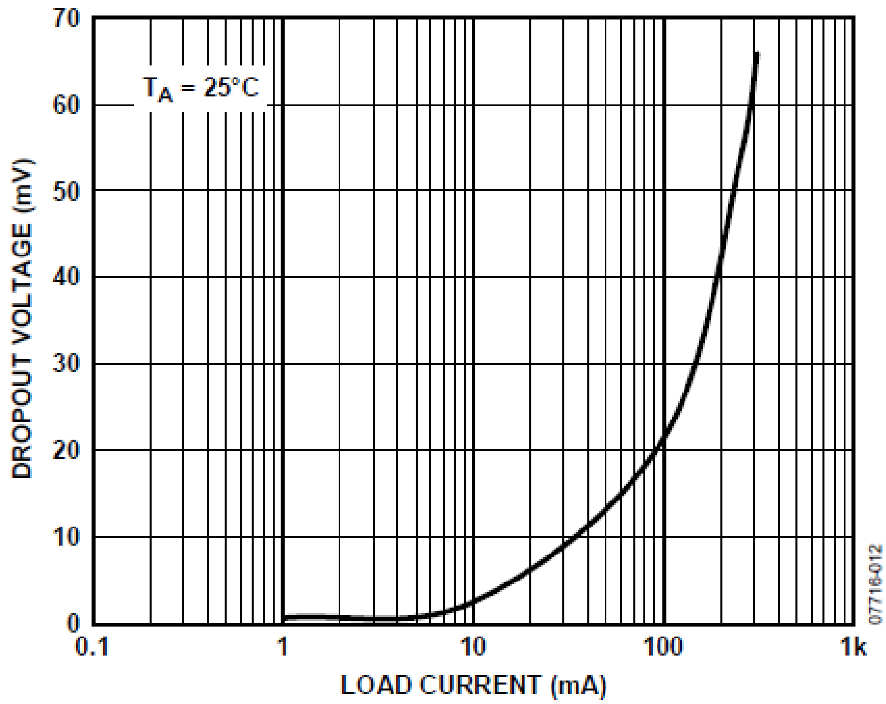

Figure 5.6: ADP171 Dropout Voltage Versus Load Current [22]

The ADP171 regulator has nominally a 66-mV dropout voltage and 130 mV in the worst case of extreme temperature at maximum load. 130 mV worst case voltage drop meant the 3.2 V initial supply could easily still be above 3.0 V after going through the regulator.

As explained in Chapter 4 with the converter regulator, the ADP171 regulator's output is adjustable with the ADJ pin. However, I did not dynamically tune the ADJ pin in this case. Instead, I used it as the designers intended, with a voltage divider from the regulator's output to ground. This voltage divider is tuned to make the regulator always output 3.1 V. Figure 5.7 below shows the schematic of the input regulator and current measurement circuit.

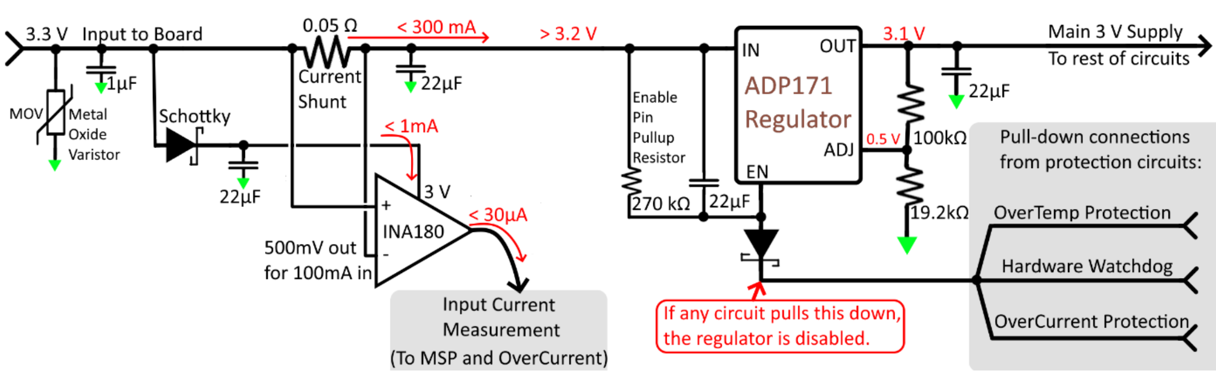

Figure 5.7: Low Voltage Input Circuit Schematic: Current Amplifier and Regulator

Input Current Measurement

I decided to include a basic input current measurement on the power supply for three reasons:

1. Current measurement is necessary as an input for an overcurrent protection circuit.
2. It allows the power supply to measure its own 3 V input current and respond to anomalies.
3. Current measurements can be sent from the spacecraft to the ground as telemetry.

For measuring current, I first investigated Hall Effect current sensors, but didn't find options with sufficient current measurement precision. They are generally aimed at much higher currents, several amperes up to hundreds of amperes. Just like operational amplifiers (opamps), Hall Effect current sensors have input offset errors, but Hall Effect sensors are much worse, on the order of tens of mA even for the most sensitive models.

At first, I planned to build a custom differential amplifier with the usual OP747, but then I learned there are special amplifiers designed for working with high-side current shunts. The amplifier needed to measure up to 300 mA with ±20 mA accuracy, so I picked the first reasonable option I found, the INA180 from Texas Instruments (see Figure 5.8). I picked the variant that has a built-in gain of 100:1 (100 V of output for 1 V of input). Combined with the 0.05 Ω shunt resistor, this means the amplifier will output 5 V for every 1 A of input current from the 3-V supply to the board. Realistically the input current will not go above about 300 mA, because any more would trip the overcurrent protection. Therefore, the highest current amplifier voltage output will be about 1.5 V, which is reasonable for the ADC to measure. The ADC's measurement resolution is about 1 mV, so it does not contribute significant error.

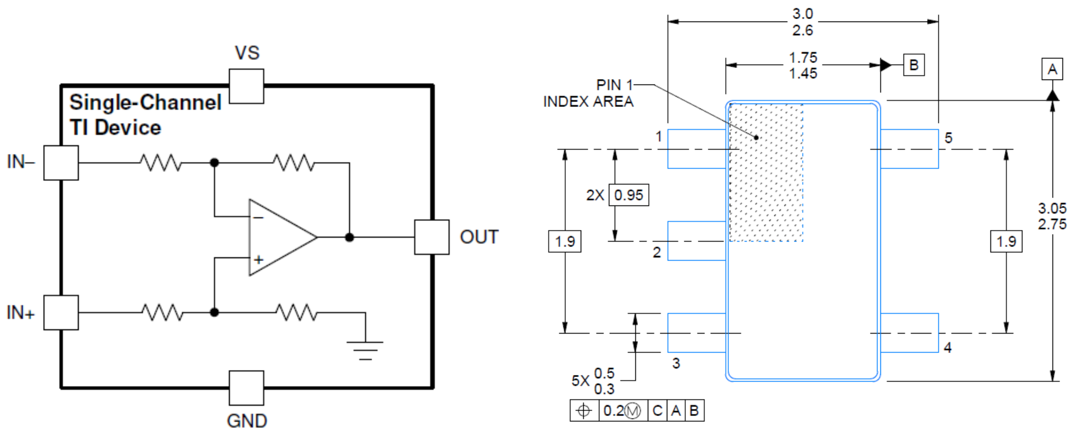

Figure 5.8: INA180 Amplifier Pinout with Schematic (left) and Package Layout in mm (right) [28]

Current Amplifier "Keep-Alive" Circuit

As Figure 5.7 above shows, the INA180 current amplifier draws power through a Schottky diode, just like the MSP430 does. This separates the current amplifier from the main supply and prevents current from flowing back into the rest of the board. In this case, the reason is to make sure the current amplifier is always active, even if the rest of the circuit is disabled. That is important because I found that only 1 or 2 V at the input will cause the INA180 to output an oddly high voltage. That voltage would be enough to erroneously trip overcurrent protection, preventing the board from ever turning on.

As the datasheet diagram in Figure 5.8 above shows, the INA180 has resistors built in, which makes it a differential amplifier with an internal gain of 100. In testing, a 3.21 V input to the board caused the INA180 to receive 3.05 V. That means only 0.16 V was lost across the Schottky diode. The INA180 will still work reliably down to at least 2.7 V, which means there is at least 0.35 V of margin available. With a capacitance of 22 µF, that is 7.7 µC of charge available to run the INA180. It consumes 260 µA quiescent current, plus up to 30 µA of maximum output, for a total of 290 µA. 7.7 µC of charge can supply 290 µA for 26 ms. That means the INA180 will be able to operate normally for at least 26 ms without any new input power coming through the Schottky diode. This ensures the current amplifier will not erroneously trip overcurrent protection because of a momentary loss of 3V supply voltage.

Op-Amps Summary

Shown in Table 5.3 on the following page is a list of all the amplifiers used in HVPS for this thesis. The OP747s are combined into groups of four amplifiers in one IC to save space. The purpose for each amplifier consists of different levels of input voltage and input impedance. Situations demanding a high input impedance were the most troublesome because the amplifier's input bias current causes more voltage error than in low input impedance cases. No application on this board required high output currents.

Why use the OP747 so often?

The reasons for using the OP747 were:

- It has low offset voltage and bias current, meaning it can be used to buffer high-impedance voltage measurements accurately.
- It can output "rail-to-rail", meaning very close to the supply voltages, in this case zero volts and 3 V. Some other opamps I used (like the TLV2264) were unable to even reach within 0.5 V of their supply voltages.
- Power consumption is low, under 300 µA per amplifier. That is about 1 mW per amplifier.
- It comes in a single IC package with four amplifier circuits, which is convenient for board design.
- It is specified over the wide temperature range of -40 °C to +85 °C.

With a gain-bandwidth product of 700 kHz, the OP747 is a relatively slow amplifier by modern standards. This low voltage slew rate was inadequate for the ESA voltage comparator, as seen in Chapter 10.

Table 5.3: Amplifiers used in HVPS Design with Purpose and Characteristics

Purpose	**Input Voltages**	**Input Resist ance**	**Bias Current Required**	**Voltage Offset Required**	**Output Requirements**	**Model Chosen**
Input current sensing	3-4 V common mode, up to 0.15 V differ-ential	0.05 Ω	< 1 mA (not demanding)	< 1 mV (corresponds to 20 mA)	Gain > 30? Must drive ~10 kΩ overcurrent protection	INA180
ESA Voltage Buffer	Input+ 0-2 V Input-Output Feedback	19.8 kΩ	< 100 nA	< 0.1 mV	Gain = 1 (buffer)	OP747
MCP Voltage Buffer	Input+ 0-1 V Input-Output Feedback	100 kΩ	< 10 nA	< 0.1 mV	Gain = 1 (buffer)	OP747
MCP return current buffer	Input+ 0-2.5 V Input-Output Feedback	25 kΩ	< 100nA	< 0.1 mV	Gain = 1 (buffer)	OP747
ESA Relay Current Control	Input+ 0-0.3 V Input- 0-0.3 V	> 2kΩ	< 1 mA	< 10 mV	Output at least 1.5 V for MOSFET gate. Needs high gain.	OP747
Overvoltage Protection	Input+ 0-1 V Input- ~0.9 V DC	110 kΩ	< 10 nA	< 1mV	At least 1 V output for driving BJT base in OVP circuit. Needs high gain.	OP747
Converter Input Voltage Targeting	Input+ 0-1 V Input- 0-1 V	54 kΩ	< 100 nA	< 0.1 mV	At least 0.5 V output. Set up as differential amp with gain of 12.	OP747
ESA Voltage Comparator	Input+ 0-2 V Input- 0-2 V	~100 Ω	< 1 mA	< 0.1 mV	At least 2 V output. Working as a comparator, needs very high gain and high slew rate.	OP747 (bad choice)

ESA Voltage Comparator

The ESA comparator was a last-minute addition in April 2023. It is not a critical part of the power supply, but it does create a potentially useful output signal. I used an available OP747 on the board as a comparator, meaning it operates with very high gain (no feedback loop). The OP747 compares the buffered ESA voltage with an arbitrary threshold set by DAC channel C. Figure 5.9 below shows a schematic of the ESA comparator circuit with input and output voltage signals.

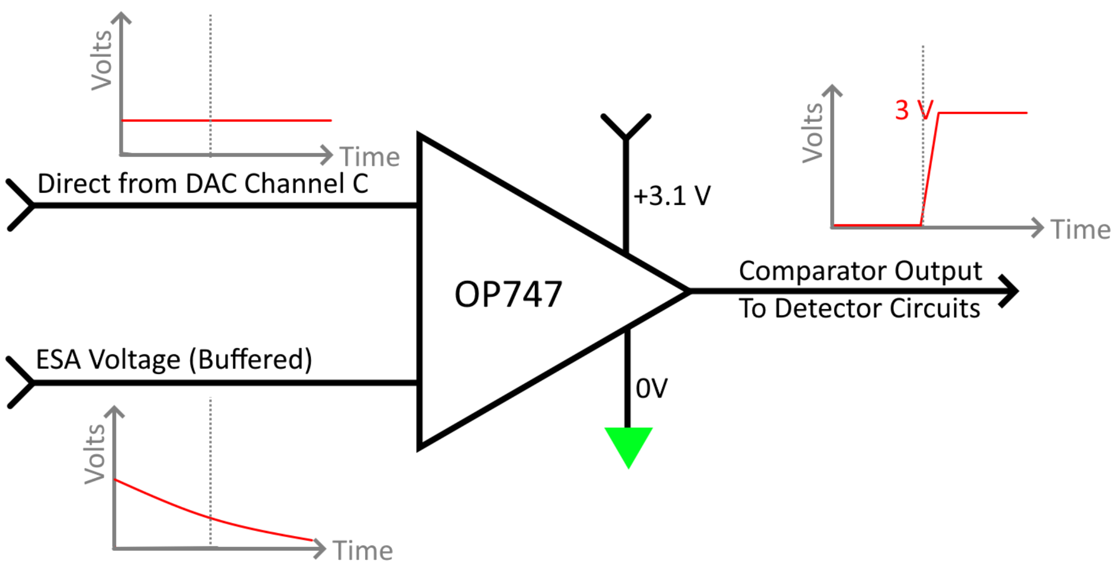

Figure 5.9: ESA Voltage Comparator Schematic and Voltage Signals

The purpose of the ESA voltage comparator is to make a sudden, noise-resistant voltage signal whenever the ESA voltage passes through a certain value. This should be very helpful for interfacing with other electronics. For example, if this signal is created when the ESA voltage passes through 4000 V at the start of every sweep, it provides a reliable timing signal for the detector. This should be much more trustworthy than the relay timing signals because the peak voltage and the relay turn-off time are both variable. Ideally, the comparator should output an instant voltage transition from 0 V to 3 V as the ESA voltage gradually changes. To practically achieve this, I should have used a purpose-built comparator.

MPS430 Temperature Sensor

For estimating the temperature of the board overall, I used the built-in temperature sensor in the MSP430. The temperature sensor produces a higher voltage at higher temperatures. The MSP430's ADC measures this voltage. It is not very accurate and requires calibration, but that is fine because highly accurate temperatures are not required. I calibrated the sensor readings based on known temperatures of the thermal vacuum chamber's platen. It is a linear calibration, so I settled on using the slope and the zero offset.

Figure 5.10 below shows the graph I used for calibration. The blue dots are ADC values I measured from the temperature sensor at known board temperatures. The blue line is the best fit line I calculated using linear regression for converting ADC values to °C. This equation for the best fit line is programmed into the MSP430, so it can calculate its own temperature in °C.

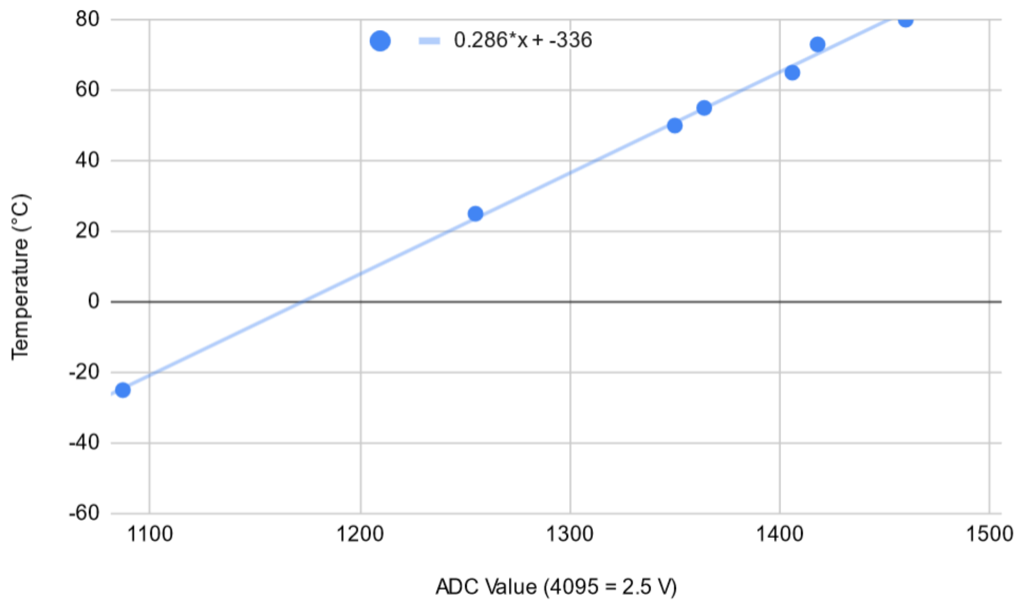

Figure 5.10: Temperature Sensor Calibration on MSP430 using Linear Regression for Best Fit

Chapter 6. Circuit Protections

This chapter describes three circuit protection systems that can reset the entire HVPS board.

- **Overcurrent** protection for the 3.3-V input supply.
- **Overheat** protection.
- **Watchdog** timer reset circuit.

There is also a fourth circuit, MCP **Overvoltage** protection, which does not reset the board.

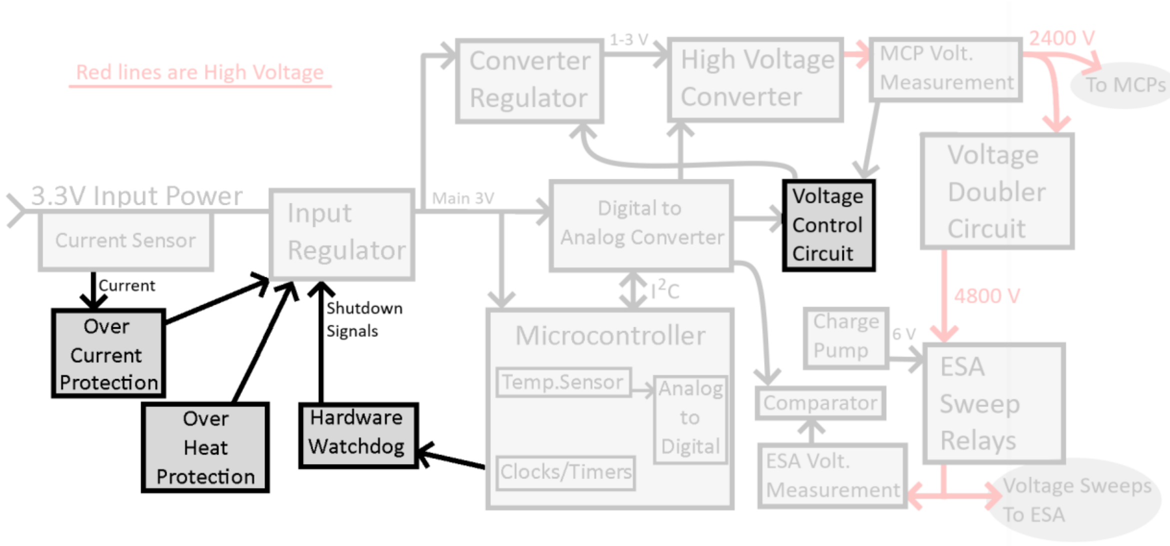

Figure 6.1: Circuit Protection Portions of HVPS Block Diagram

How do the protections reset the board?

The main power regulator is used as a reset switch for the whole board. As seen back in Figure 5.7 the main 3-V regulator's enable pin is normally held at 3 V by a 270-kΩ pullup resistor. Combined with a 22- µF capacitor, this gives a very long RC time constant (several seconds), preventing that regulator from being disabled or enabled immediately upon startup. Figure 6.2 below shows how the enable pin works. As long as the enable pin is kept above about 1 V, it will turn on the regulator. The enable pin has built-in hysteresis, meaning the threshold for turning off is higher than the threshold for turning on. Hysteresis is important because it prevents the regulator from rapidly "flip-flopping" between on and off.

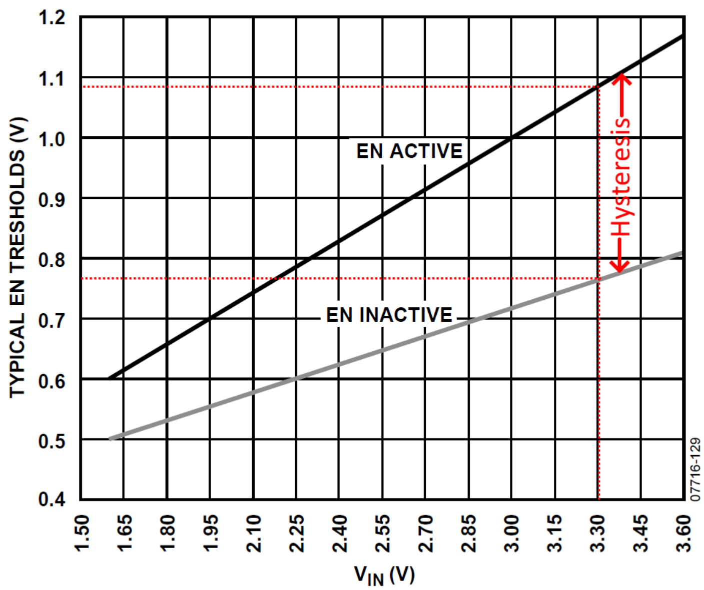

Figure 6.2: ADP171 Regulator Hysteresis [22]

Figure 6.3 below is a repeat of the schematic showing the input regulator and its enable pin.

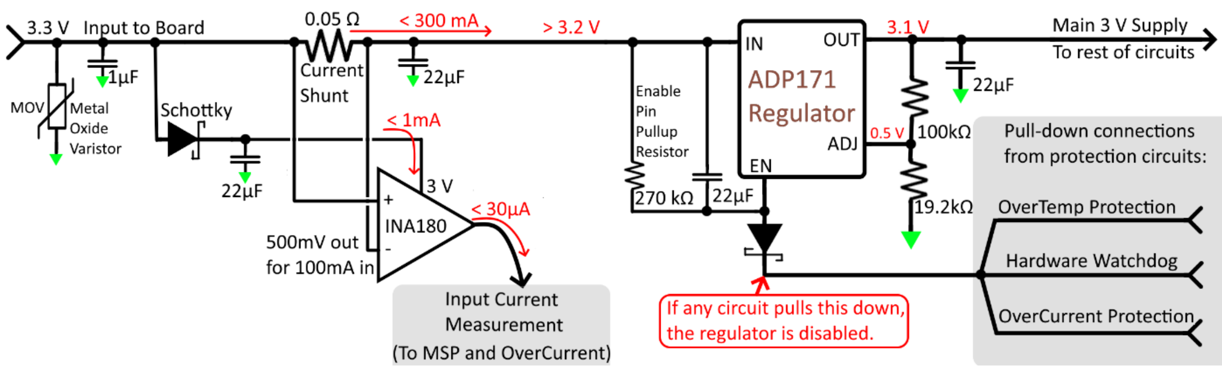

Figure 6.3: Input Regulator Schematic Repeated

A Schottky diode in Figure 6.3 connects three protection circuits to the main regulator's enable pin. Normally, all three protection circuits output 3 V, so they do not pull the enable pin down, and the Schottky diode carries no current. However, if any protection circuit wire does go to 0 V, it will pull the enable pin down below the roughly 0.8-V threshold (see Figure 6.2), and the regulator will be disabled. Once the regulator is disabled, the main 3-V supply will lose power. This will cause nearly everything to shut off: the MSP430 and DAC will reset, all OP747 amplifiers will lose power, and all high voltage components will lose power. Without 3-V main power, all the protection circuits are disabled, meaning they can no longer pull the enable pin down. However, the 3.3-V input to the regulator is not turned off, so the 270 kΩ resistor is still pulling the enable pin up. Once the protection circuits have lost power and are no longer holding the enable pin low, the 270 kΩ pullup resistor gradually brings the enable pin voltage back up above ~1.1 V, making the regulator turn on again. Once the regulator is turned on again, there are two possibilities. If a protection circuit is still triggering, the board will keep resetting infinitely until the problem goes away. Otherwise, the board will start up normally, with no memory of the reset which just happened.

Overheat Protection

Overheat protection is important because it is unsafe to operate a high voltage circuit like this at extreme temperatures. Even if the circuit does appear to function at 80 or 90 °C, high temperatures will degrade the lifespan of components and increase outgassing, which may cause arcs. I was not so worried about low temperatures, because I have not found any reason to believe that cold temperature operation would cause permanent harm.

After researching the various types of temperature sensors, I decided that a Positive Temperature Coefficient (PTC) thermistor with a highly nonlinear temperature response is perfect for triggering an overheat protection circuit. PTC thermistors can have a very sudden increase in resistance beyond a certain temperature, so they are close to being ideal on-off devices. I selected the B59641A0075A062 from TDK Electronics. It is only 500-1000 Ω in the normal range of temperatures, but the resistance begins ramping up suddenly around 60-70 °C. Once it reaches 90 °C, the resistance will be over 30 kΩ. Figure 6.4 shows the PTC thermistor's resistance versus temperature range.

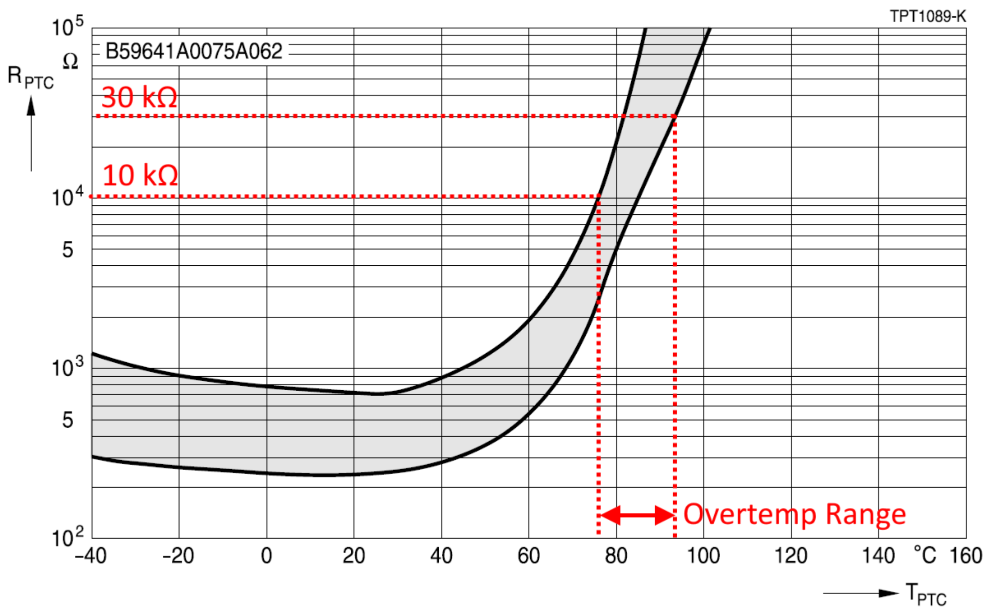

Figure 6.4: Overheat Protection PTC Thermistor Resistance versus Temperature [29]

Figure 6.5 shows the overheat protection circuit I designed, which turns on an n-channel MOSFET once the thermistor's resistance exceeds about 10-20 kΩ.

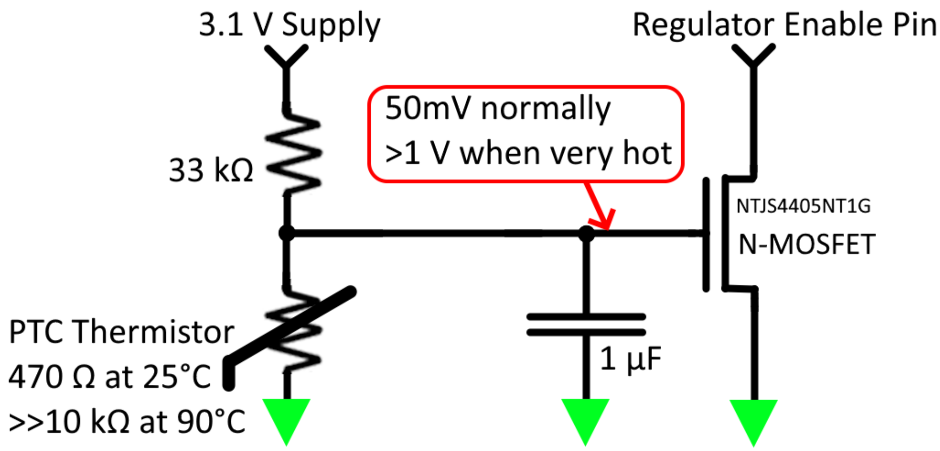

Figure 6.5: Overheat Protection Circuit Schematic with PTC Thermistor and N-Channel MOSFET

When the PTC thermistor is below about 60 °C, its resistance is only around 500 Ω (as shown in Figure 6.4), causing the MOSFET's gate voltage to be low. However, when the thermistor increases in temperature, its resistance will rise enough to make the MOSFET's gate voltage increase. When the gate exceeds about 1 V, the MOSFET will turn on, pulling the regulator enable pin to ground and triggering a board reset as explained previously. The 1-µF capacitor is for filtering noise to prevent momentary accidental resets.

Overcurrent Protection

Overcurrent protection is necessary to make sure this board cannot completely short its 3 V power source. Also, it makes the board more self-sufficient, able to recover from problems by resetting itself, instead of waiting to be reset by an outside signal or source.

Why not use the ADP171 regulator's built-in current limiting?

The ADP171 regulator does have a built-in 450 mA current limiting circuit. There are two reasons I did not want to use it. First, 450 mA is too high of a threshold because I wanted to keep current below 250 or 300 mA. Also, I wanted overcurrent protection to actually cut off current completely, but the ADP171's current limiting will still allow some current to flow. Simply limiting the maximum current will not trigger a reset of the whole power supply, so it will not necessarily stop a fault.

Overcurrent Protection Circuit Design Requirements

The design principles I followed while designing the overcurrent protection circuit are as follows:

- Overcurrent protection should be active as soon as the board turns on, because latchup or other shorts can happen anywhere, anytime.

- It needs to measure current with a minimal voltage drop, so the voltage at the load remains consistent and close to 3.1 V. This is achieved by the INA180 current shunt amplifier.
- The protection must have hysteresis, so it does not get stuck oscillating rapidly on-off-on-off.
- Overcurrent must not be tripped by very short pulses of high current, because that happens in normal conditions whenever a capacitor gets charged up as the board turns on. In other words, the current must remain high for multiple milliseconds, just like with a fuse or a circuit breaker.
- There should ideally be no software involved.

The current measurement is done on the high voltage side, which ensures reliable grounding. The nominal current through the main 3.3-V input is designed to be 100-200 mA, with anything over about 250 mA indicating a fault which should trip the protection circuit.

Figure 6.6 shows the schematic for the overcurrent protection circuit.

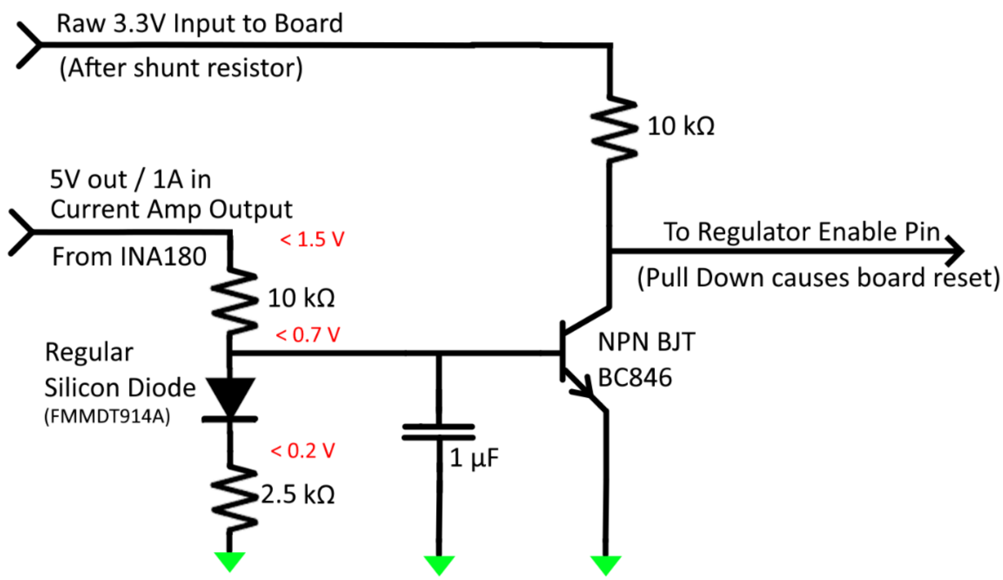

Figure 6.6: Overcurrent Protection Circuit Schematic

The INA180 input current amplifier will output about 1.5 V when the current is high enough that overcurrent protection is triggered (around 300 mA). To achieve this threshold, I made a voltage divider combined with a diode, driving a BJT transistor. When the current amplifier output is above about 1.5 V, the voltage at the base of the BJT exceeds about 0.7 V, causing it to turn on. This pulls the regulator enable pin toward ground, triggering a reset of the board. During normal conditions, the current amplifier output is not high enough to turn the BJT on.

This combination of diodes and resistances does not need to be very accurate, since it is acceptable for the exact overcurrent trigger level to be imprecise. There are two important points, listed on the next page.

1. I made sure to use a BJT and not a MOSFET. BJTs have a predictable turn-on voltage of around 0.5-0.7 V, while the gate threshold voltage of MOSFETs is much less certain, especially if they are radiation damaged.
2. The diode in the low side of the voltage divider is important for cancelling out the thermal changes of the BJT's turn-on voltage. The base of a BJT takes less voltage to turn on when hot, just like a diode. With a diode built into the voltage divider, it automatically compensates for temperature changes, to keep the overcurrent threshold relatively constant.

In testing, I found that the circuit would not draw more than about 180 mA at 3.3 V under normal operation. Therefore, I wanted the overcurrent trigger level to be somewhat higher, such as 250 mA. To test the overcurrent circuit, I soldered on various low-ohm resistors between the main 3-V supply and ground. First, I tried 20 Ω (150 mA), which didn't trigger overcurrent at all, as expected. 15 Ω (200 mA) was barely enough to trigger overcurrent, so it took 450 ms. Meanwhile, a 10 Ω (300 mA) load triggered overcurrent in only 8 ms as shown in Figure 6.7 below.

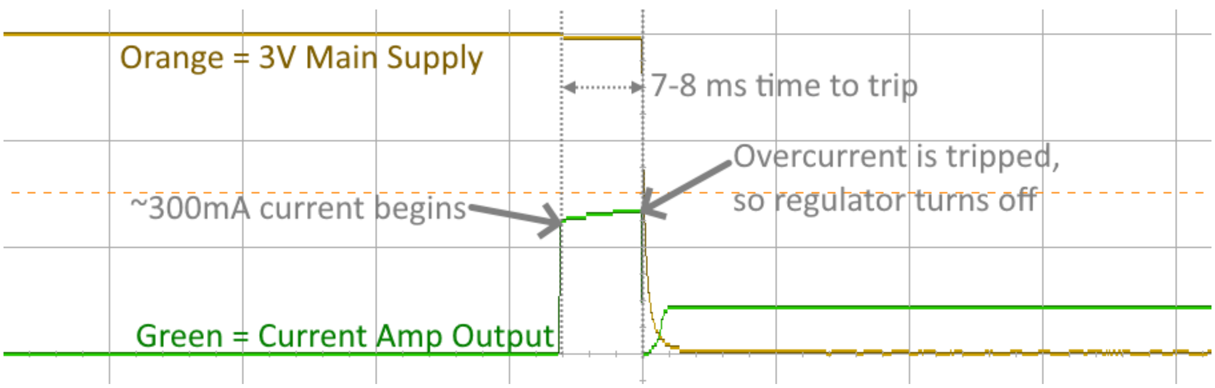

Figure 6.7: Overcurrent Protection Triggered at 300 mA (10 Ω) Load

Also, I tested overcurrent with the high voltage output shorted directly to ground. This tripped overcurrent even faster, taking only about 3 ms, as shown in Figure 6.8. The current amplifiers output reached a peak of 3 V, which corresponds to > 600 mA of input current, but only for a very short time.

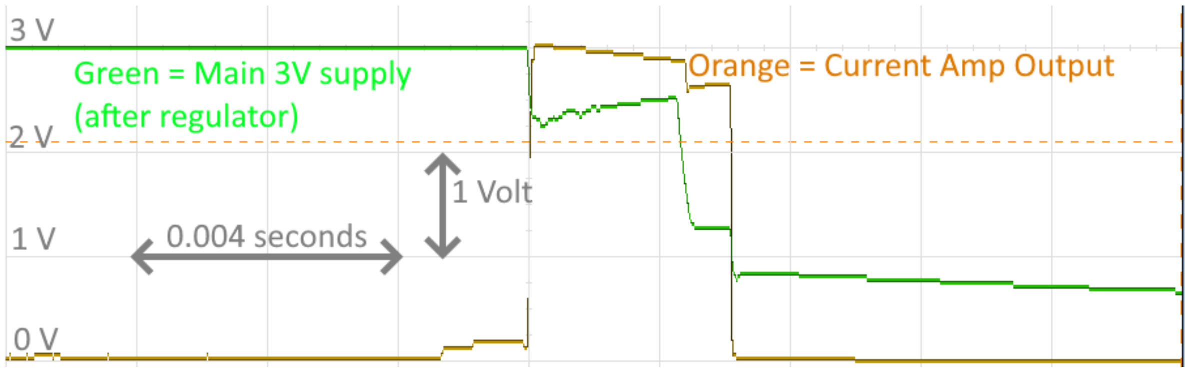

Figure 6.8: Overcurrent Protection Triggered with Direct Short to Ground

Figure 6.9 below shows the relationship between current and the time required to trigger overcurrent protection.

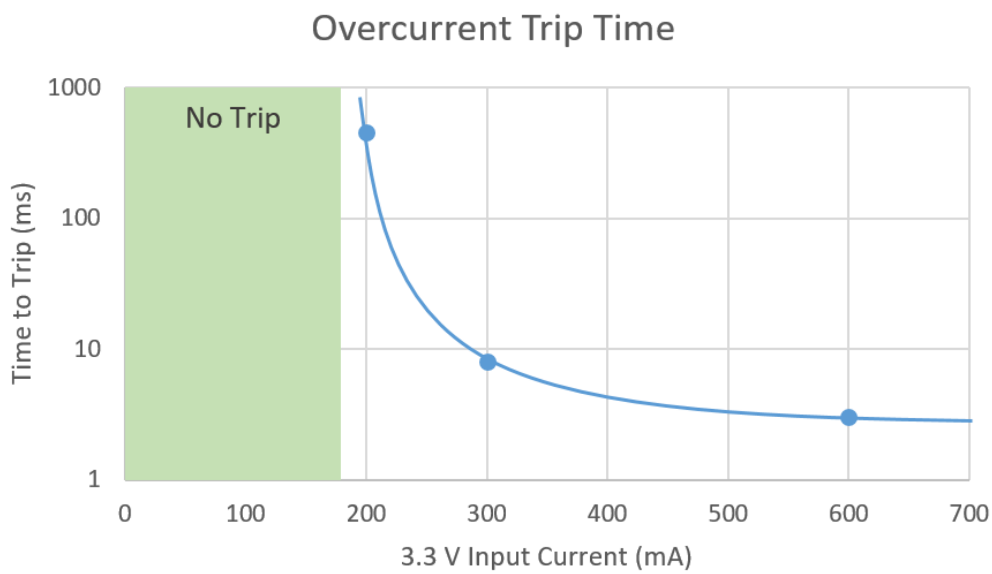

Figure 6.9: Overcurrent Trip Time Versus 3.3-V Supply Input Current

The shape of the curve in Figure 6.9 is similar to that for a fuse or circuit breaker. It is desirable that an extreme overcurrent causes a trip more quickly than a slight overcurrent.

Watchdog Reset Circuit

A watchdog is a monitoring circuit that constantly checks a signal to make sure it is being updated ("kicked"). If the watchdog is not kicked for some chosen amount of time, it forces a reset. When thinking about watchdogs, I was particularly concerned about radiation. High energy particles penetrating an IC or microcontroller can cause the software to hang or even worse, create a short circuit. This is called Latchup. Fortunately, these radiation effects are usually temporary, and can be cleared by simply cutting power to the circuit. Completely cutting power can also be an effective method of extinguishing a high voltage arc.

There is a watchdog timer built entirely into the MSP430 microcontroller. Without any extra design work, it will reset the MSP430 if the timer is not "kicked" frequently enough. This is very helpful for resetting the software if it hangs, but I also wanted a circuit that would be able to reset the entire power supply, not just the microcontroller.

The choice of what signal to monitor was important. For example, it would be a bad idea to use the main clock signal of the microcontroller, because it might still produce a clock even if the CPU was not executing code correctly. I chose to use a pin on the microcontroller which alternates between high and low voltage every time the watchdog timer is "kicked". The kicking only happens when the ESA relays are switched. This means if the ESA relay signal stops for whatever reason, the board will soon reset.

Figure 6.10 shows the watchdog reset circuit.

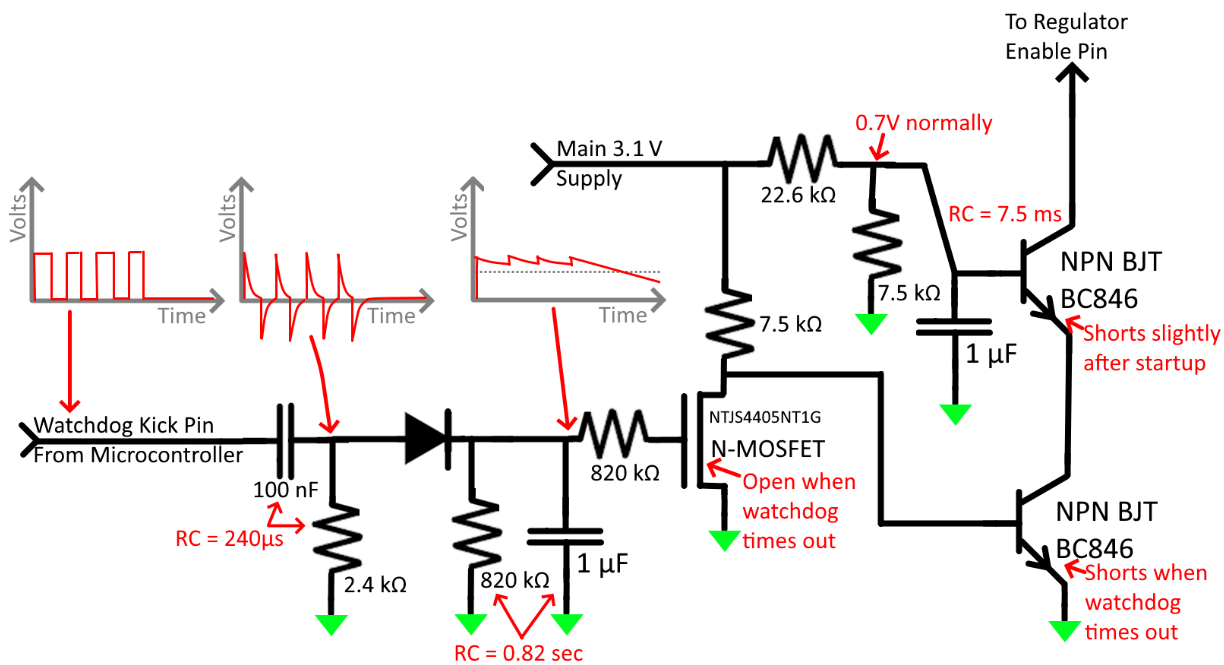

Figure 6.10: Hardware Watchdog Circuit

The watchdog circuit has many steps. Here we follow the signal from left to right in Figure 6.10 for the case when the microcontroller is working normally.

1. The relay control signal from the microcontroller toggles on and off, creating a square wave voltage at all times.
2. This square wave goes through a <u>series</u> capacitor and resistor, turning it into AC which is centered at 0 V.
3. This AC is rectified with a diode and filtered by another resistor and capacitor. Now it is DC with voltage ripple.

4. This DC passes through another resistor to the gate of an n-channel MOSFET, turning it on.
5. This MOSFET being turned on causes the lower BJT to turn off.
6. The BJT turning off allows the regulator enable pin to remain at 3 V, so the board remains powered.

This may sound overcomplicated, but it is necessary to distinguish between a pin with a changing voltage and a pin stuck at high voltage. It is very possible for software glitches or radiation to make a pin get stuck forever high. In contrast, a pin repeatedly jumping between high and low voltage does not happen by accident. If the watchdog kick pin gets stuck either high or low, it will make a pure DC voltage level instead of a square wave. DC does not pass through capacitors, so no signal will even reach step 2 in the above process.

There is yet another complication. Just as the circuit starts up, the watchdog circuit's capacitors take time to charge, so it will briefly act as if the watchdog timer was not working and try to reset. This would cause the board to get stuck infinitely resetting. To prevent this, another RC circuit connected to a the upper BJT in Figure 6.10 causes an approximately 7.5 ms delay before the watchdog circuit is able to reset the board.

As a test, I intentionally shorted the watchdog kick pin, so that it could not transmit any AC through the 100-nF capacitor. This caused the board to repeatedly reset. Figure 6.11 below shows the watchdog circuit in action.

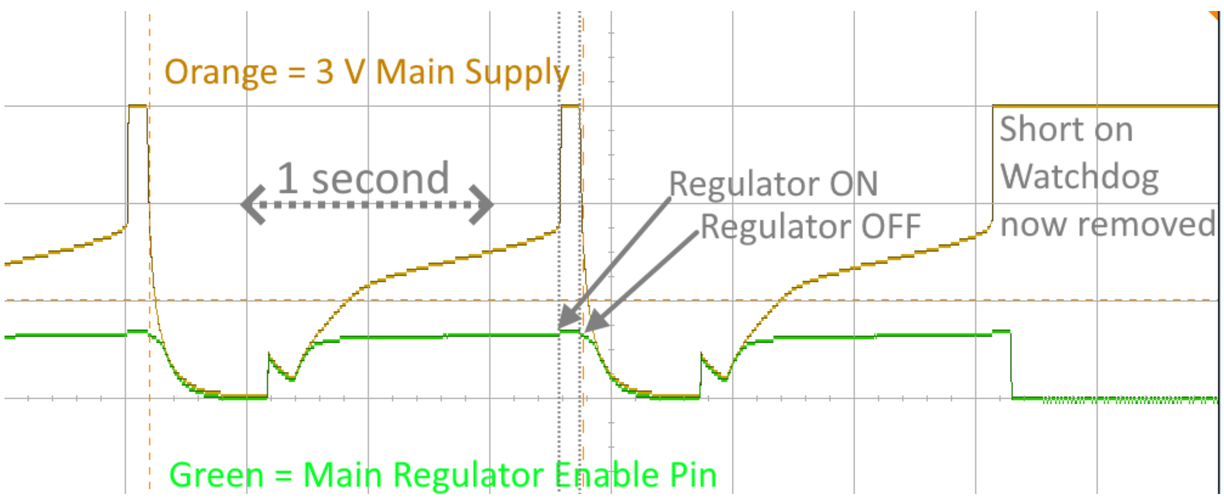

Figure 6.11: Watchdog circuit in action

This same kind of watchdog circuit can be applied to other signals. I think it would be a good idea to use this design for the I^2C clock and/or data signals, to reset the board if the I^2C communication gets stuck. However, the series capacitor would have to be reduced from 100 nF to probably about 1 nF to allow higher-frequency signals to pass the watchdog circuit without disrupting the I^2C signal timing.

Overvoltage Protection

It is very important to establish a solid upper limit on detector voltage. Even a very short overvoltage could be catastrophic. MCPs can be damaged or destroyed by excess voltage, because it will trigger a high voltage arc through the channels of the MCP. Also, overvoltage can cause MCPs to enter "thermal runaway", where their resistance decreases as they heat up, so power dissipation increases even faster, and they generate even more heat. This creates a positive feedback loop which can destroy the MCP or the power supply, or both.

The first line of defense is passive design of the circuit. As shown in Figure 4.6, the low side of the MCP output voltage divider has extra capacitance of 10 nF along with a 100 kΩ resistance, giving it a RC-time constant of 1 ms. This is long enough to prevent a spike of interference or noise from confusing the control circuitry but is still short enough to cause a minimal phase shift of the control loop.

The two ADP171 linear voltage regulators already used in this design provide isolation between the power input to the board and the converter. Without these regulators, the high voltage converter might be hit by a sudden voltage increase, causing an overvoltage. Each regulator provides a Power Supply Rejection Ratio (PSRR) of better than 60 dB at frequencies up to 1 kHz, meaning incoming voltage noise will be dropped by a factor of more than 1000.

The large RC time constant of roughly > 100 ms at the converter output limits how quickly the converter can change its voltage. If there is a malfunction which makes the converter begin to charge up those MCP voltage filter capacitors too quickly, it might draw too much current and conveniently trip the board's overcurrent protection.

However, I decided those passive design features explained above were not enough, considering the severe consequences of an overvoltage event. I designed a dedicated overvoltage protection circuit, shown in Figure 6.12 below.

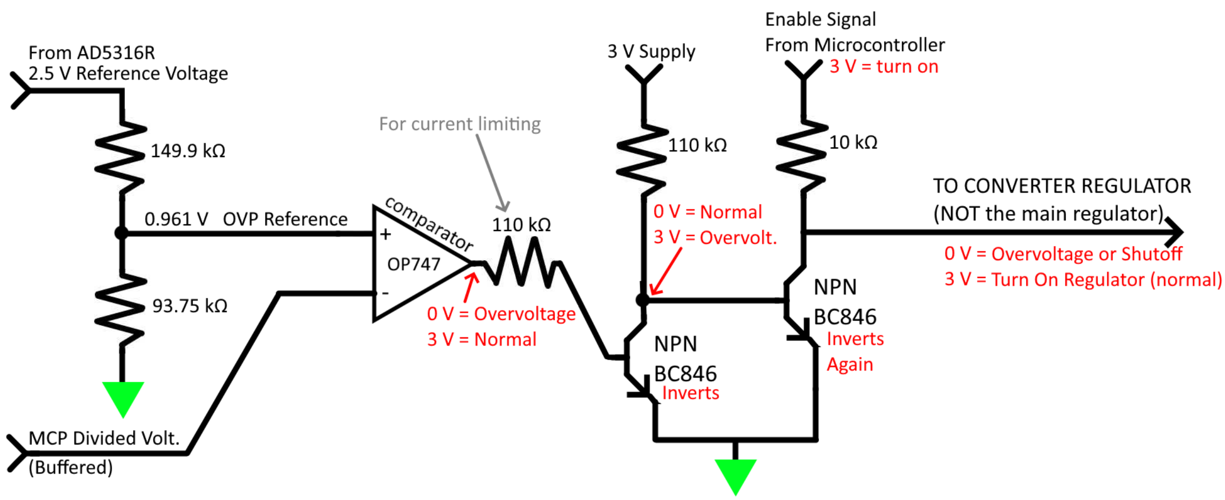

Figure 6.12: Overvoltage Protection Schematic

The overvoltage protection is a hardware-based feedback loop that can cut power to the converter. This protection is all built into the circuit board and cannot be easily bypassed or changed without a soldering iron, which is helpful for safety. The overvoltage cutoff level is defined by a voltage divider starting at 2.5 V. Currently, it is designed to limit the voltage at the reference divider to 0.961 V, or 2880 V for the actual HVPS. This limit voltage can only be changed by altering the resistors in the divider on the left side of Figure 6.12. The only critical components for overvoltage are amplifiers, transistors, resistors, and the ADP171 regulator's enable pin. Failsafe operation is achieved by inverting the signal twice, see Figure 6.12.

If the microcontroller does not supply a 3-V "turn on" signal to the overvoltage protection circuit, the converter regulator's enable pin will stay low (off), regardless of any other signals. Even if the microcontroller does supply the "turn on" signal, the overvoltage reference signal must also be present, and the 3-V supply must be present, and the OP747 must not be shorted. It fails safe in every way I can think of, except the very strange case of resistors failing shorted.

In testing, overvoltage protection does work at clamping the voltage. However, I have occasionally seen small overshoots, if the voltage is rising rapidly as it crosses the overvoltage threshold. The way to minimize this risk is to use a very slow voltage ramp-up when the power supply starts.

Unlike many types of protection circuits, this is voltage limiting, but not resetting or "latching". That means it will block the MCP voltage from going too high, and that is all. It will not force a restart or anything else, and it does not have any hysteresis. If resetting from overvoltage is desired, software on the MSP430 can detect the excessive voltage and directly trigger a reset.

Why not short the high voltage output directly for overvoltage protection?

While directly shorting the converter's output is the most certain way of ensuring no dangerous voltage can remain, it is troublesome. It would require another high voltage tolerant switch, such as another string of relays, which is a lot of complication. Also, suddenly shorting the MCP supply would trigger very high currents through the converter and might introduce electrical noise problems from the voltage transient, or even damage the converter. Cutting power input to the converter instead is much simpler and almost as direct, because the MCP voltage cannot possibly go any higher if no new energy is added to the converter.

Chapter 7. HVPS Board Design

This chapter describes the development history and process of prototyping the HVPS board. At first, I was not building traditional circuit boards at all, but eventually I did graduate into the world of computer-aided electrical design. I designed the later board revisions in Cadence/Allegro. For the first two actual-size boards, I machined them myself in the Space Systems Engineering Lab on a CNC milling machine. For the latest board revision, I ordered it etched from Advanced Circuits instead of milling. In all cases, the components were hand-placed and soldered with hot air and a conventional soldering iron.

Table 7.1 below is a summary of the HVPS board development and prototyping history.

Table 7.1: Summary of HVPS Board Development and Prototyping History

<table>
<tr><th>Name / Picture</th><th>Purpose / Results</th><th>Fate</th></tr>
<tr><td>

Spring 2020: First converter testing

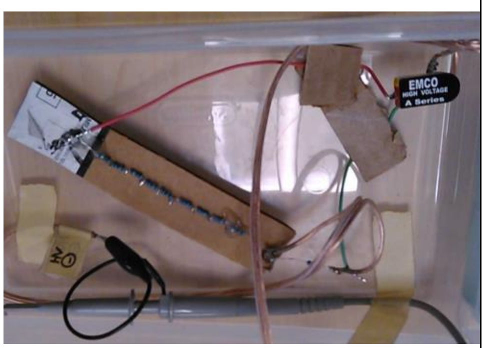

</td><td>

Exploring the potential use of EMCO (XP Power) A-series converters.

Developing basic skills with high voltage design and measurement.

Discovered the slow response time of the converters. Learned how to use control pin and how to vary the input.

Learned how to make voltage dividers and sample them with an oscilloscope.

</td><td>

High Voltage Converter was used for other vacuum experiments.

Plastic tub with resistors still exists.

</td></tr>
<tr><td>

Summer 2021: First perf-board designs

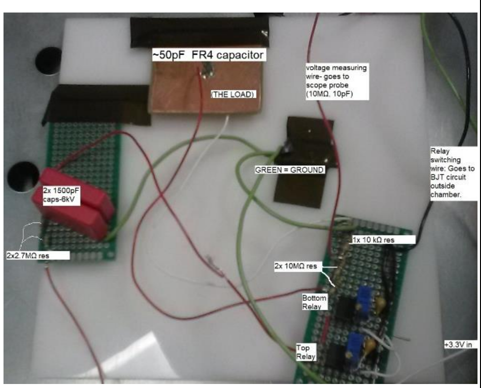

</td><td>

First time operating in vacuum.

Learned how to work with solid-state relays to generate sweeps.

Explored power consumption and optimizing the relay current.

Pushed limits of voltage on the relays and capacitors.

Explored how to vary the ESA sweep resistance.

Design review was in Dec.2021.

</td><td>

Some pieces still exist. Capacitors exist, may be useful.

Some relays were reused, others were fried.

</td></tr>
</table>

Spring 2022: Copper on fiberglass "board" 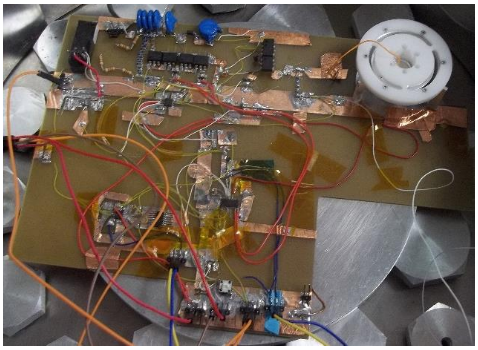	This was the first mostly complete HVPS I made. Still missing protection circuitry. Learned how to use amplifiers and DAC. Learned how to set up I2C communication. First combined MCP and ESA supply. First thermal tests using a heat gun. First design in Cadence OrCAD. Developed voltage doubler idea.	Board fell apart and most parts were reused. Had corrosion issues and copper tape didn't stick well enough.
July 2022: First milled board ("R0") 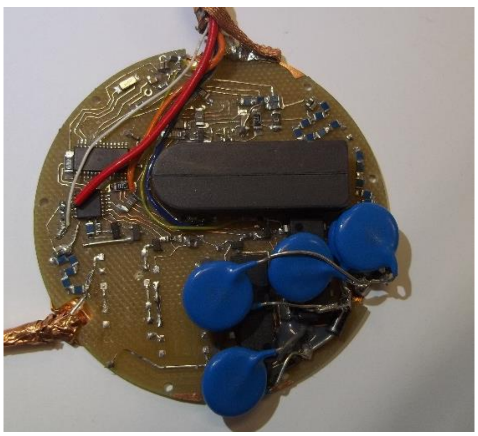	First time packing everything into a tight space. Learned how to design a circuit board and mill it. First time using the power supply to actually measure an electron beam. First major issues with arcing. First time trying to write a proper control algorithm. First time working with epoxy for high voltages.	Problems after electron beam exposure. Cause unknown. Board still exists but many pieces were reused.
September 2022: Second milled board ("R1") 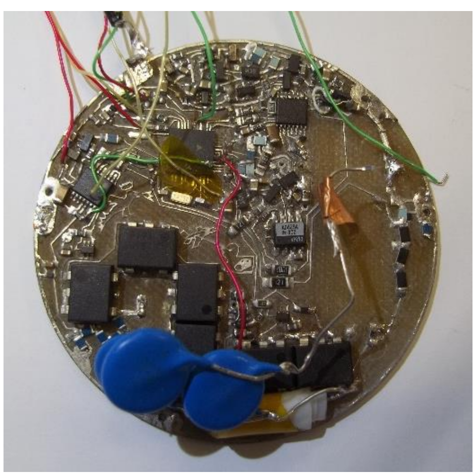	Adding protection circuitry. First thermal vacuum testing. First current measurements. Upgrading to a better microcontroller. First time working with PEEK film for high voltage insulation.	Severe arcing incident in Jan. 2023. Board still exists but some parts we cannibalized for R2. Several questionable solder joints.

March 2023: First etched board ("R2") 	Moved high voltage components farther apart and added more spacing to the board edges. Mostly just minor changes. Added a few convenience features. Improved opamp circuits, buffered all signals, added comparator. Learned about proper epoxying. Long-duration testing. Tuning for reduced power consumption Finalized protection circuits	This is the current design. Still works.

During early prototyping, I did not attempt to make an actual-size board right away. Instead, I spent several months using a much larger and more flexible sheet of one-sided FR4 circuit board material. The bottom side was grounded, and I used pieces of 3M 1181 copper tape to form crude traces on the top side. That copper tape has been absolutely fantastic for shielding and making connections, especially in a vacuum. Figure 7.1 shows the prototype circuit in the SSEP lab's vacuum chamber just before designing R0 of the board.

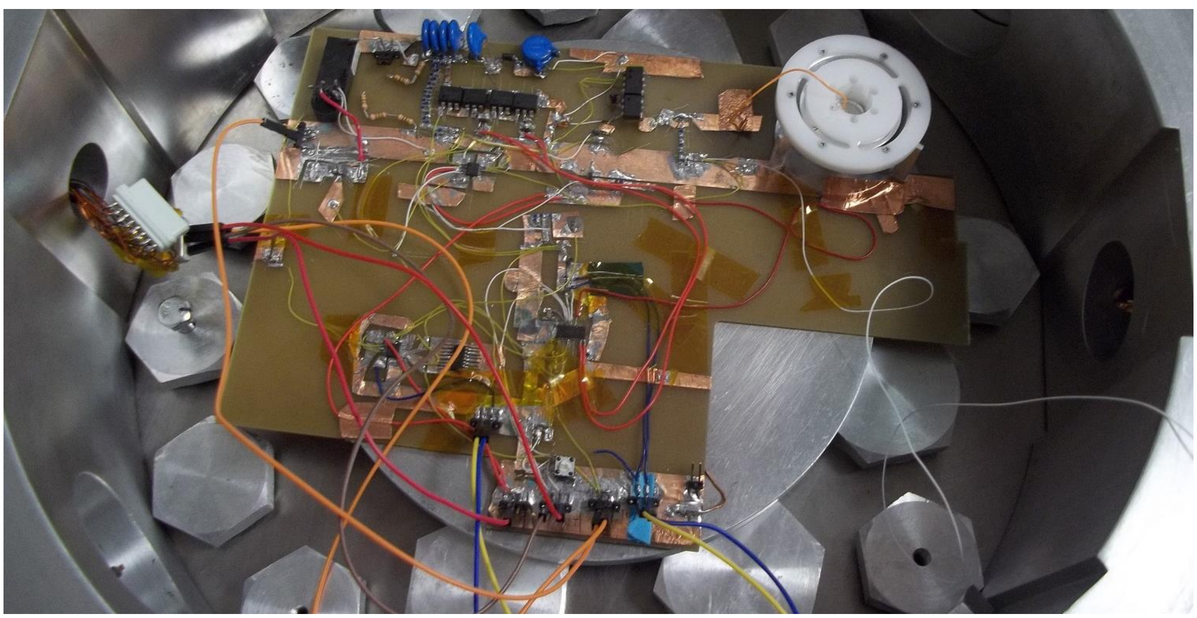

Figure 7.1: Prototype in old vacuum chamber (field of view is 450mm wide)

This method of prototyping helped me develop soldering skills and board layout intuition, without the commitment of an etched or milled circuit. The copper tape was very easy to change as soon as I found a mistake. By the time I started the first actual-size board revision, I was already completely confident that the design could work. Practically all of the design features I have described so far were already developed before I made an actual-size board, except the protection circuits.

The HVPS Board

The board is designed to sit on top of the ACAPS ESA. This defines its maximum diameter as 63 mm. I defined a circle of nine #0-80 threaded screws for holding the board around its edge. Three of the screw locations are missing because other screws for the ESA are in the way. I already had the ESA designed in Solidworks, so I defined the board outline there as well as shown in Figure 7.2. Then, I exported the board outline from Solidworks as a .dxf file.

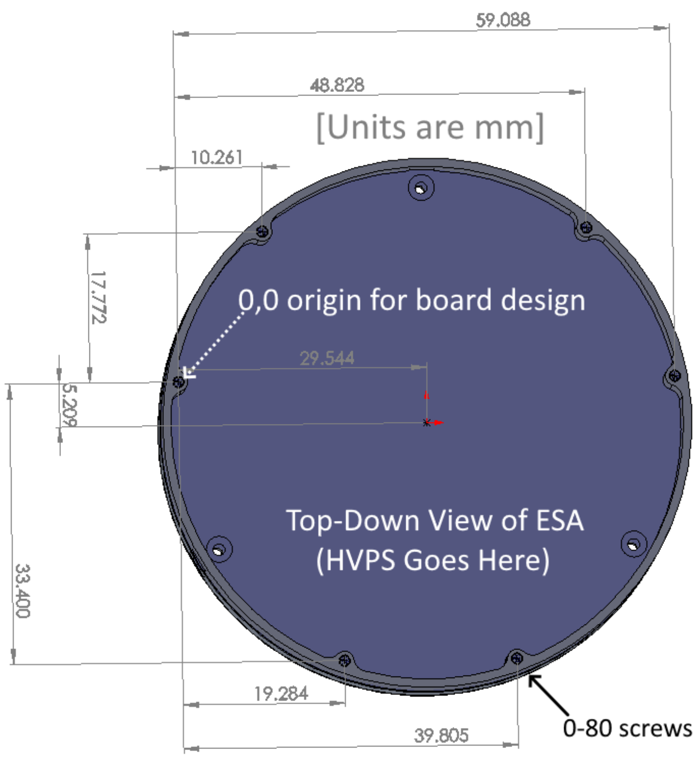

Figure 7.2: Board dimensions in Solidworks

Board Layers

Even though 4-layer boards are completely normal in the SSEP lab, I did not use one for several reasons. First, I thought it was unnecessary. Based on a preliminary "board area budget", I believed that it was possible to lay out all components on a single layer. I considered it a fun and interesting design challenge even if it was not truly necessary. A 4-layer board also would need to have thicker layers to avoid arcing through the insulator material between each layer. This would cause the entire HVPS to be 1 mm thicker or more.

Finally, a 4-layer board is much harder to manufacture on a CNC milling machine, because it requires attaching layers to each other. Practically, 4-layer boards would require ordering the board from a company and waiting for shipping to Alaska, while 2-layer boards are reasonable to manufacture in the SSEP lab. Eventually I did order a 2-layer board, but only after the design was nearly finalized.

The final design is a 2-layer board, with components and traces only on the top side, while the bottom side is only ground. The grounded bottom side of the board is designed to be directly exposed to space. The bottom side must be only a ground plane to maintain a continuous wall of shielding around the whole circuit. The only connections through the board are used for connecting to ground pins on the top of the board. Keeping all signals on the top side is helpful for troubleshooting because it means a signal will never be inconvenient to reach during testing.

Board Area Budgeting

In late 2021, I decided that it was possible to fit the ESA in a 60-mm diameter circle, based on laying out paper mockups of the components. Later I made a 3D mockup to develop additional confidence that it could all fit (Figure 7.3). This was at the time when I still expected to have two different high voltage converters.

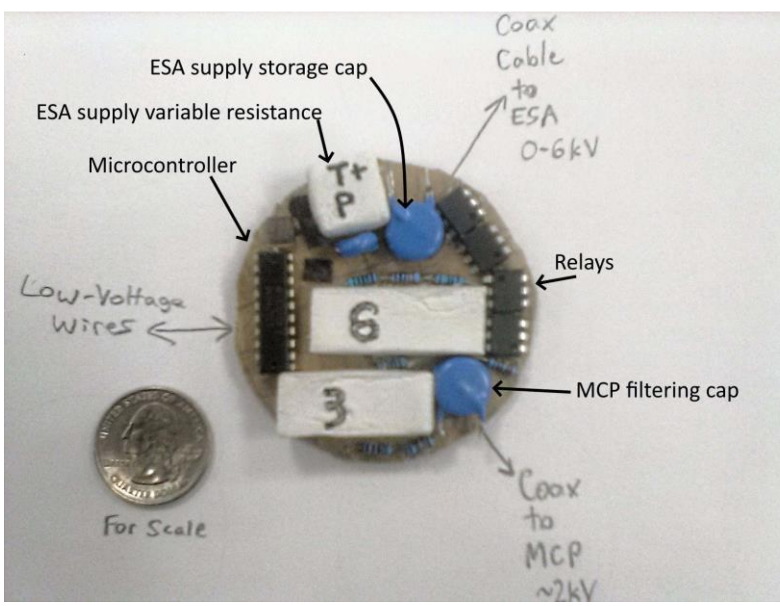

Figure 7.3: Early 60-mm diamter board area mockup

Board Milling

For board revisions 0 and 1, I used FlatCAM to convert Gerber and Drill files exported from Allegro into CNC mill toolpaths. Figure 7.4 shows milling in progress on the SSEP lab's Tormach PCNC 440 milling machine. In this picture, a 3.175 mm (1/8 in) endmill is cutting the outline of the Revision 1 board.

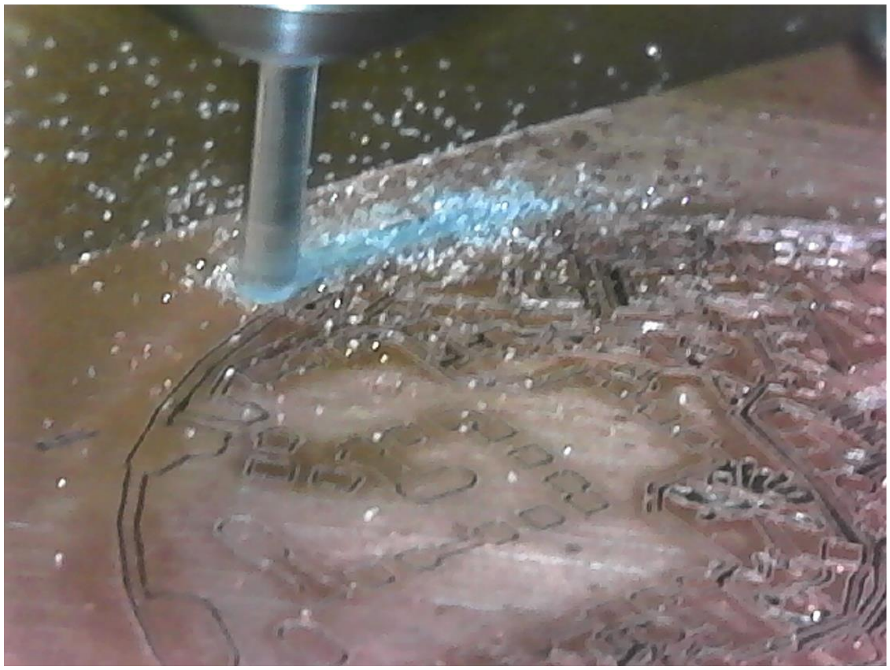

Figure 7.4: Circuit Board Milling

The R0 Board

The first actual-size board (R0) is shown on the following page in Figure 7.5 at four stages of construction. This R0 board used 0.25 mm (0.010 in) minimum trace width, and 0.38 mm (0.015 in) minimum spacing between traces. I routed traces in a quite inefficient way, so a lot of the board area is wasted.

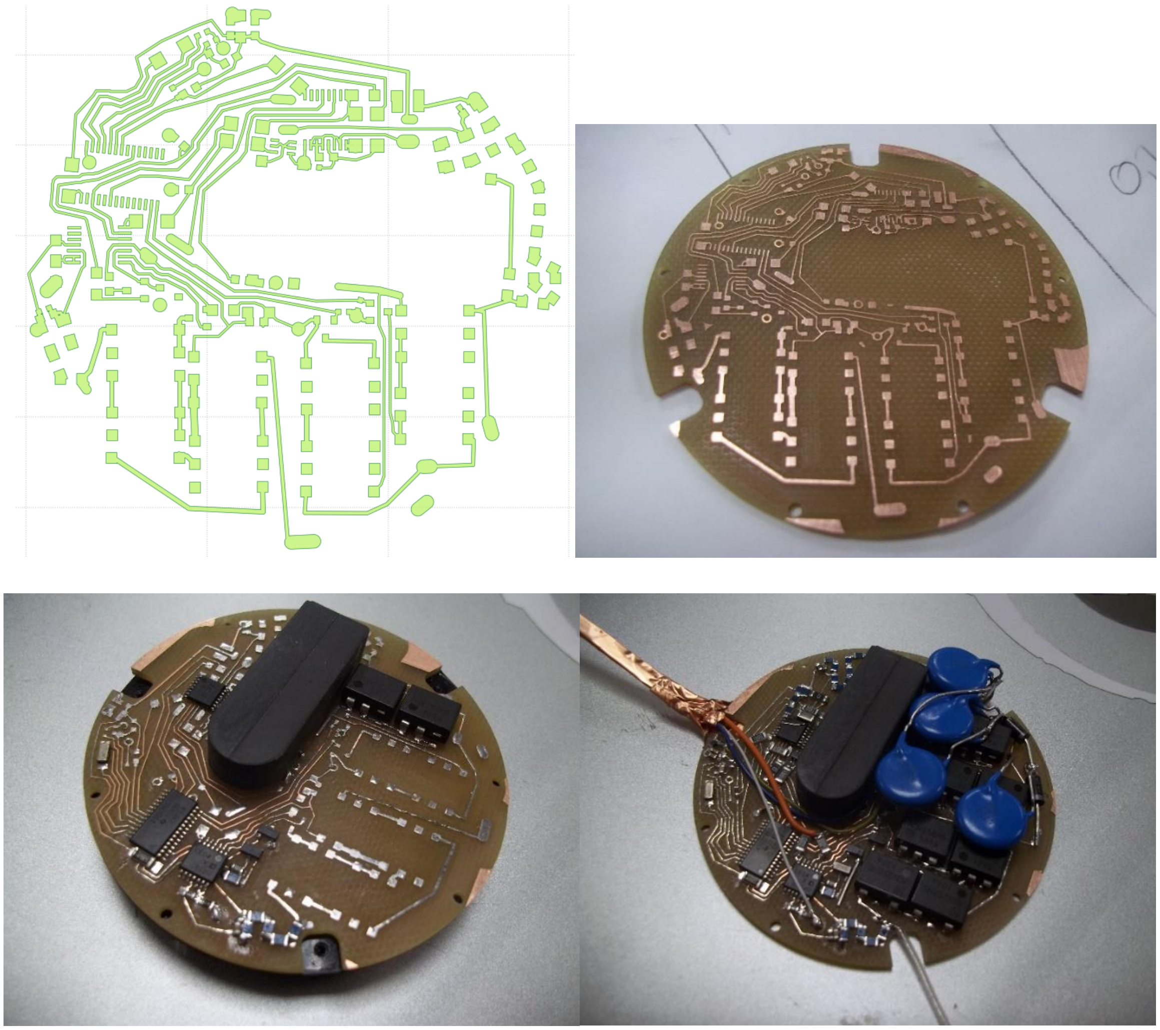

Figure 7.5: Four Stages of R0 Board Construction

R1 Board

For the R1 board (see Figure 7.6 on the following page), I used 0.15 mm (0.060 in) minimum trace width and 0.25 mm (0.010 in) minimum spacing. This was difficult because it required using a smaller 0.127 mm (0.005 in) engraving bit in the CNC mill, which was frustratingly fragile. Such a dramatic reduction in trace width and spacing liberated a large area of board space, so the struggle was worthwhile. Also, I switched to using a microcontroller with tighter pin spacing and removed several relays and resistors. On the R0 board, the doubler's low voltage side had been two relays, but on R1 I replaced them with a string of resistors for simplicity and space savings.

I made the footprint for the MSP430 have long solder pads to accommodate the radiation-hardened version of the microcontroller, which has pins that extend outward farther. On this revision I switched to using the OP747 for most things, except I had an Analog Devices AD623 for driving the converter regulator. This version also included some “artwork” cut into the board. Most importantly, this is when I implemented all of the protection circuitry.

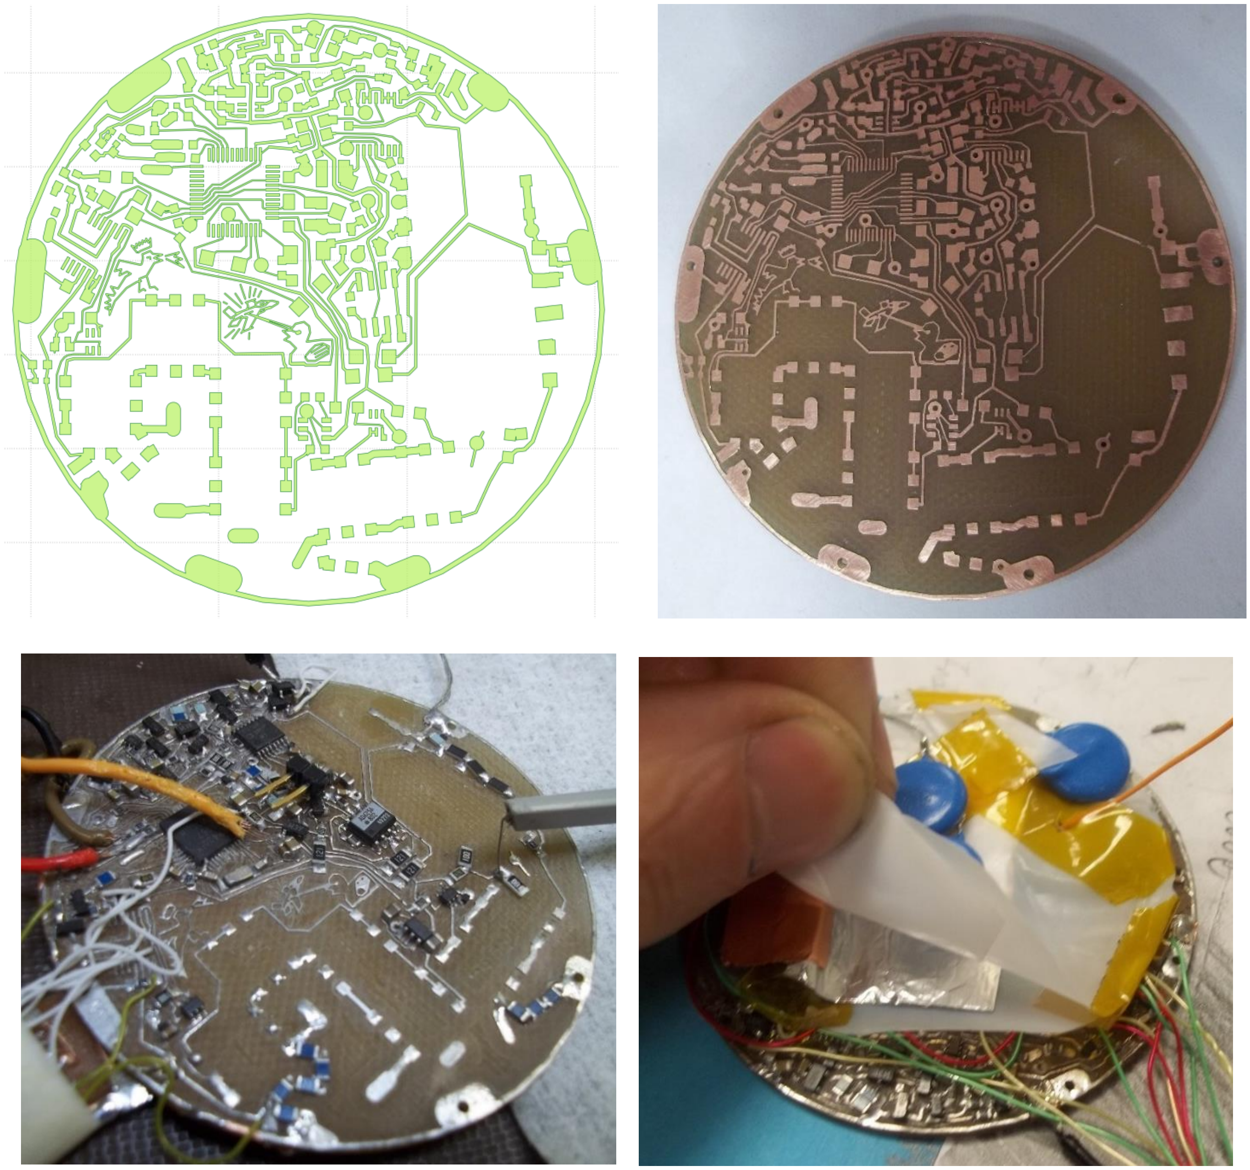

Figure 7.6: Four stages of R1 Board Construction

The bottom left image of Figure 7.6 shows the testing process before I had added any high voltage components. For safety, I generally tried to test everything possible before attaching the converter. The bottom right image shows the completed board, mostly covered with PEEK plastic film and aluminum foil for shielding. This protects the low-voltage electronics and allows the large blue high voltage capacitors to safely lay flat over the board.

R2 Board (Latest Revision)

R2 is very similar to R1 (see Figure 7.7), with a few small improvements listed in Table 7.1. However, this board was etched by Advanced Circuits. The end result was much higher quality than my milled boards. Also, on this revision, I moved all of the relays farther inward toward the center of the board.

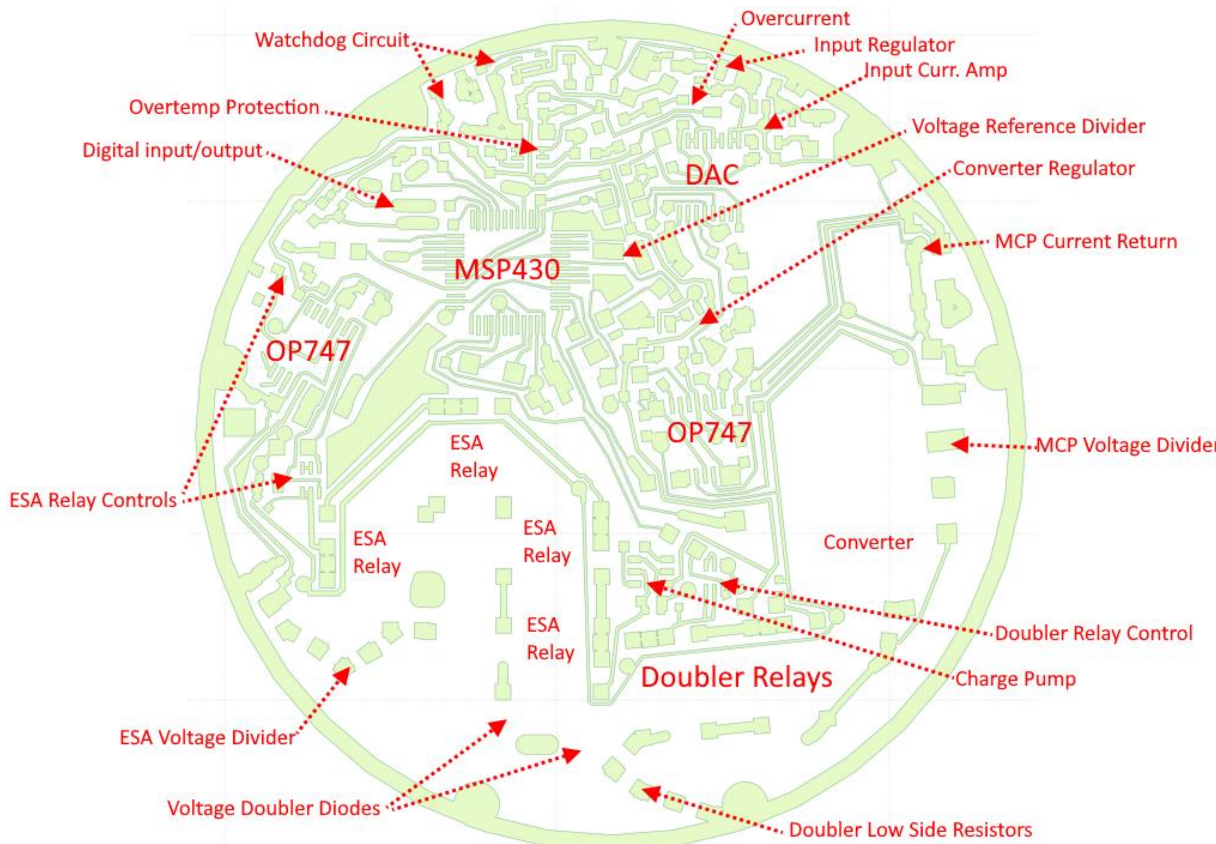

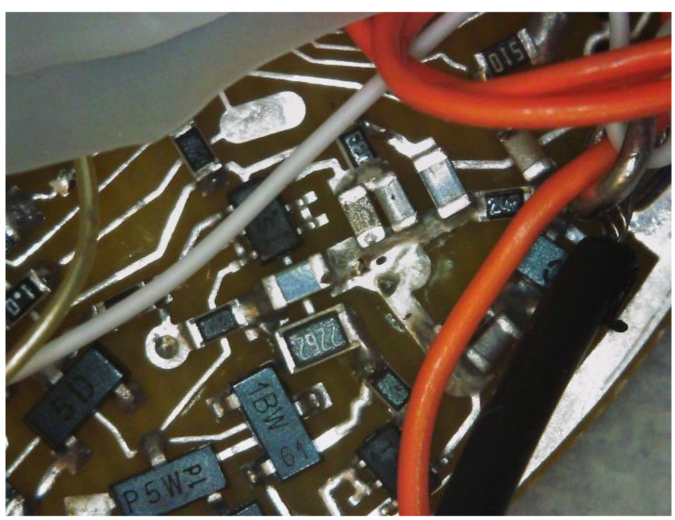

Figure 7.7: Two Views of the R2 Board

Figure 7.8 below shows a picture of the R2 board with the major components labelled. Additional components, insulation, and epoxy were added after this picture was taken. The high voltage converter and MCP capacitors are still missing from the picture, as well as insulating epoxy and all the wiring connections.

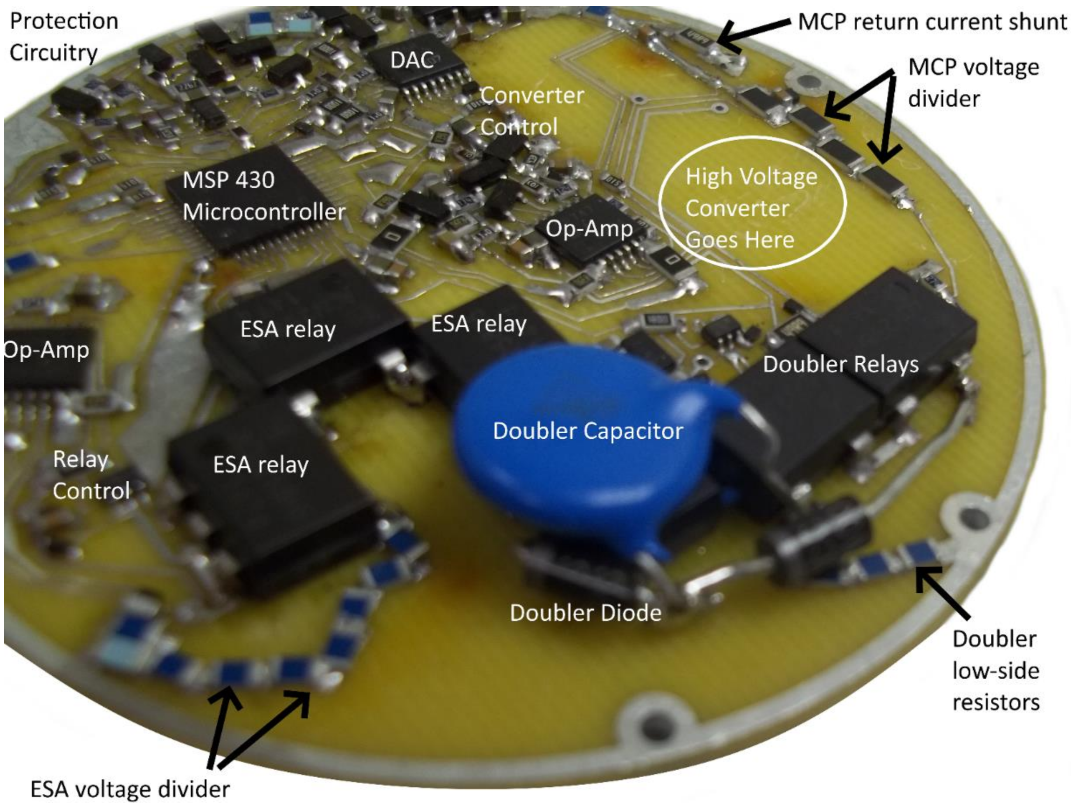

Figure 7.8: Final Revision Board with Major Components Labeled

Trace Shielding

To protect sensitive low-voltage signals, I surrounded many of them with strips (traces) of electrically grounded metal. This has two separate purposes:

- If a high voltage arc forms on the board surface, it will likely jump to ground instead of frying whatever electronics connect to the traces.
- Having ground between traces reduces coupling between them. This effectively makes measurements of analog voltage cleaner because sudden voltage spikes are less likely to appear.

Figure 7.9 shows an Allegro screenshot with an arrangement of ground traces surrounding sensitive analog signals. This area is especially important because the traces go directly under the high voltage converter, which is not perfectly shielded.

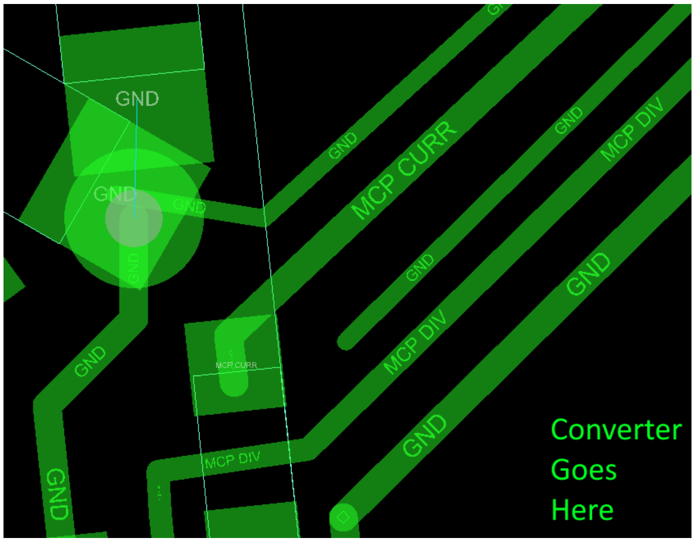

Figure 7.9: Shielding around board traces.

High Voltage Standoff Distances

Figure 7.10 shows the highest electric field strength areas on the R2 circuit board design.

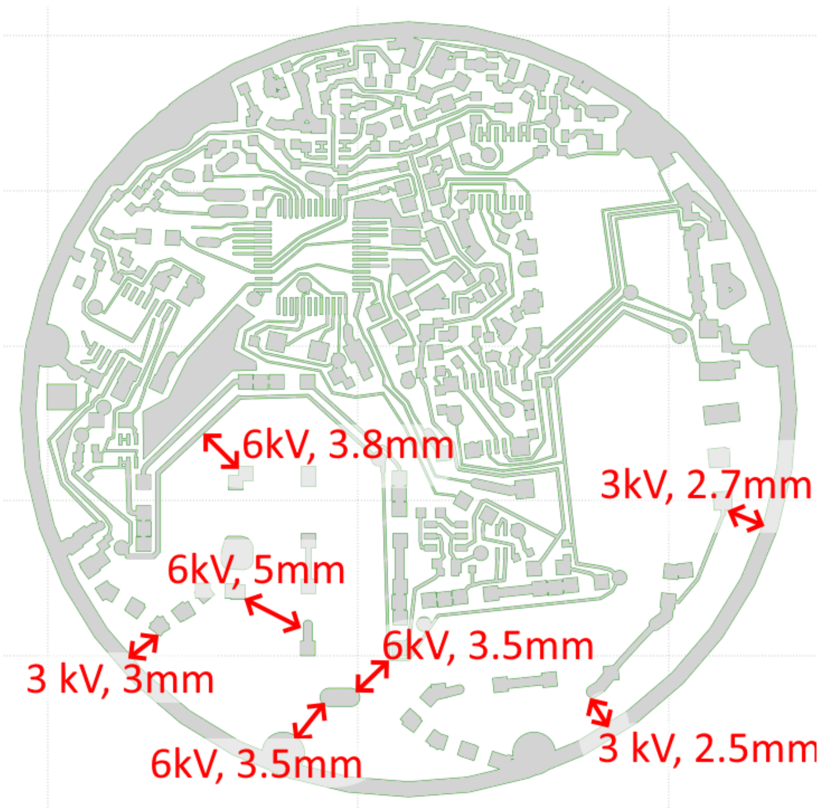

Figure 7.10: High Voltage Standoff Distances on R2 Board

If the board is actually operated at maximum design voltage, the worst offenders in terms of voltage per mm are about 6000 V across 3.5 mm. That is an electric field strength of 1700 V/mm, well above the usual rule of thumb, which is 1000 V/mm. However, I did not allow these to be simple gaps across the surface of the board. I flooded every one of these sensitive high voltage gaps with large amounts of epoxy. Also, I did not actually operate quite at 6000 V, because of high voltage fatigue.

High Voltage Fatigue

Just like mechanical components, electrical insulators wear out over time. Failures are often not immediate. High voltage fatigue is often not noticeable because engineers usually can afford to operate with huge safety margins relative to the insulator material's limits. However, packing high voltages into a small space makes the problem harder to avoid. Even if no arcs are visible, insulating materials will gradually degrade when exposed to voltages near their limits. Eventually they will degrade so badly that conductive "tracks" can form which arcs eventually follow. Once tracking happens, the circuit is basically ruined.

I learned about the empirical study of high voltage fatigue from a set of slides by Steve Battel [30]. As a rule of thumb, the lifetime of an insulator decreases in proportion to voltage increases to the 9th power. That means increasing the voltage by 25% will cut the estimated operating life by a factor of about 7. The exact numbers depend on the insulator in question, and the rule may not apply in all cases, but all materials do seem to have problems with long-term high voltage exposure. Conservative rules of thumb for exposure to high voltage are around 2000 V/mm within insulator materials, and below 1000 V/mm along surfaces. It is possible to achieve much better insulation to high voltages than this, but reliable operation over months or years is certainly not guaranteed.

The circuit board itself is a potential breakdown risk. The fiberglass between top and bottom of the board is about 1.5 mm thick, so even going beyond 3000 V appears questionable using those conservative rules of thumb. Only long-term testing will tell whether a thicker board is needed. I have operated circuit boards of this same thickness at 5000 V for hours in the past, but not weeks or months.

I have experienced high voltage fatigue with several prototypes which worked fine for a few hours or days but then began arcing. This is counterintuitive; usually a device operating for several days at a range of temperatures is seen as having proven a decent level of reliability. With high voltage fatigue, the circuit may have been rapidly breaking down during that time without showing any visible signs. This is why it is important to test for weeks, if not months, at a higher voltage than is actually expected for operation in space.

Although this power supply is certainly capable of producing 3000 V for the MCP supply, I chose to operate at only 2400 V maximum (4800 V with doubler). Based on all the arc problems I was seeing in vacuum, I believe that using the full 3000 V (6000 V with doubler) would not give a reasonable operating lifespan, even though it has worked for a few hours so far. This is not a terrible loss, since 2400 V will likely still be enough for any MCP detector stack.

Why are some areas of the board covered with epoxy?

First, epoxy (and plastics in general) is used to achieve closer spacing between voltages without arcing. It is much more difficult for an arc to propagate through a material, compared to on the surface of a material. Epoxy also helps strengthen connections on the board. This is especially important for the largest components, such as the high voltage converter. Solder joints alone are mechanically questionable, as solder is soft and prone to cracking. Also epoxy helps hold wires in place. Finally, the particular epoxy I used has high thermal conductivity, so it helps pull away heat generated inside some components.

Figure 7.11 below highlights all the areas I coated with epoxy on the R2 board. It also shows three designated wiring notches for wiring connections.

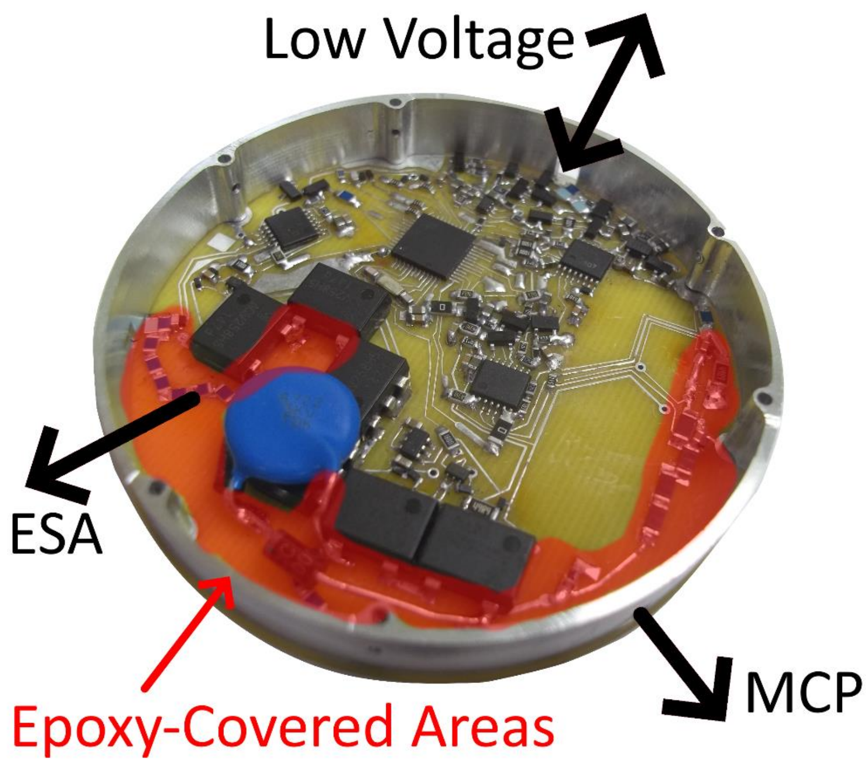

Figure 7.11: R2 Board with Outside Ring Shield and Epoxy Coated High-Voltage Areas

The epoxy I used was Duralco 4525, a black and relatively hard thermally conductive epoxy. I felt comfortable using it on this project because I previously had made vacuum feedthroughs with it.

Eventually, the low-voltage areas of the board should also be coated in a thin layer of plastic or epoxy, but I chose to not do that during testing. Once a board is coated, it becomes very difficult to take measurements or make any modifications to it.

Figure 7.12 shows the R2 board after all thermal vacuum testing was complete.

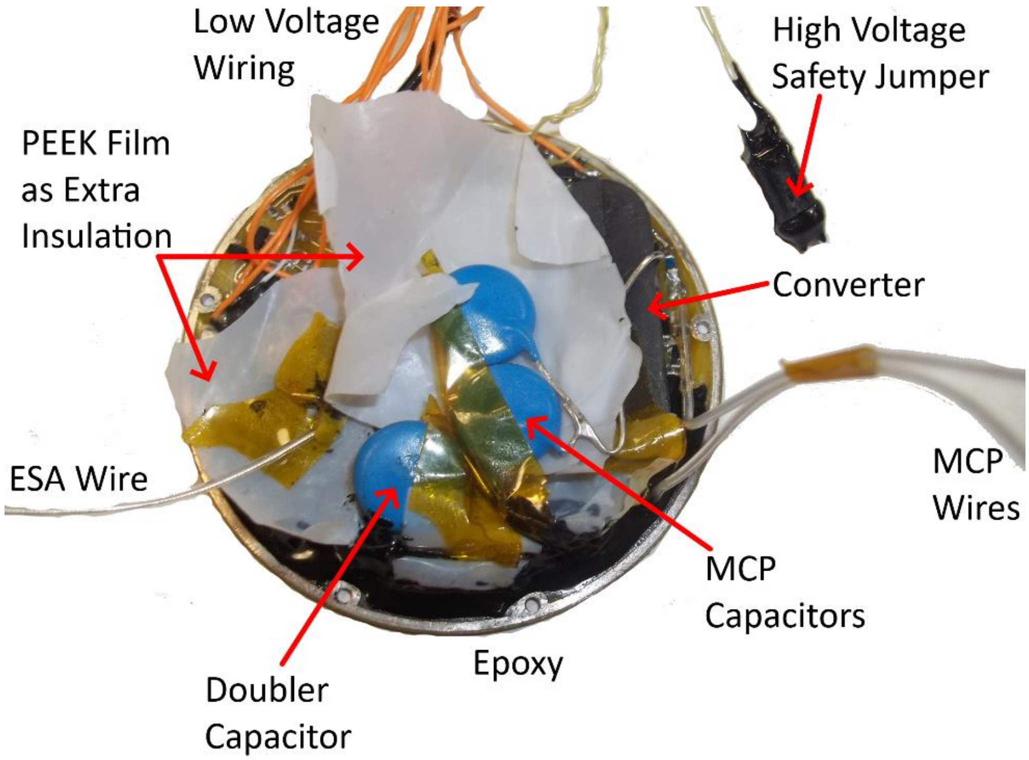

Figure 7.12: Complete R2 Board

How can the board be made safe?

There is a safety jumper which can short the high voltage converter's input, completely preventing high voltages from being generated. Figure 7.12 shows how the safety jumper hangs off of the board on its own wires. If the converter tries to turn on with the safety jumper in place, it will promptly trip the over-current protection, forcing the board to reset. I chose this method because:

1. Directly shorting the high voltage is too dangerous, so I needed to block high voltage indirectly.
2. Any software-based control is not trustworthy enough, because software cannot directly block voltage from being generated.
3. Disabling the converter's control pin would theoretically work but is less reliable than shorting its input. At low loads, the converter can still produce output despite the control pin being at 0 V (see Figure 4.5).
4. Disabling the converter’s regulator might also work, but I wanted an even more direct method.

Rules for how to use the Safety Jumper:

1. Always keep the jumper in place unless you need high voltage output. Think of it like the trigger of a gun. You do not touch the trigger unless you are ready to shoot.
2. If you have the jumper installed but you do not want overcurrent to trip from the converter trying to turn on, turn high voltage off using software.
3. The jumper is like a remove-before-flight pin, so it needs to be easily accessible during vacuum experiments. Once the board is covered by a shield on top of the ESA, the jumper must be secured to this shield.

External Wiring:

The HVPS is designed as a stand-alone unit with one required and four optional connections to external hardware.

- (Required) 3.3-V power supply and ground. Ideally with ground and neutral separated, so that neutral carries all of the current. (3 wires)
- (Optional, not yet implemented) UART data interface (similar to USB) for sending commands to the HVPS and gathering telemetry. (2 wires)
- (Optional) Sweep output voltage (divided by ~3000), so another circuit can sample it. (1 wire)
- (Optional) MCP output voltage (also divided by ~3000), so another circuit can sample it. (1 wire)
- (Optional) ESA Comparator Signal which sends a pulse whenever a new sweep begins. (1 wire)

One major disadvantage of attaching the HVPS on the top of the ESA is that all wiring must pass through the ESA's field of view before it can reach the rest of the satellite. This means the wires should be as narrow as possible. The ESA itself already has three structural support pillars which cross the field of view. Ideally all of the wiring should pass along the pillars to avoid creating any new blind spots in the field of view. The pillars are roughly 3 mm at widest point as shown below in Figure 7.13. The figure shows how the board will sit on the shield, on top of the ESA, along with the wiring path.

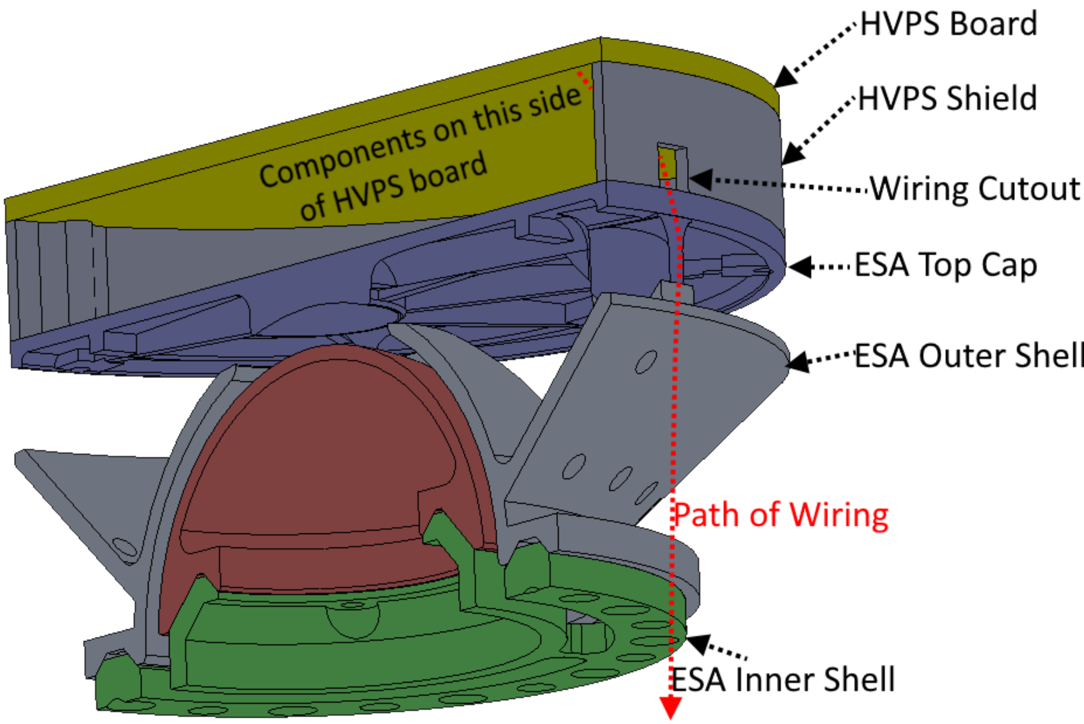

Figure 7.13: ESA, Shield, HVPS, and Wiring in Solidworks

Passing the low voltage wires along the pillars is not a big concern because they can be very thin, such as 0.5 mm in diameter. Low voltage can be tightly clustered together in a single bundle. However, high voltage wires have to be a larger diameter, because their insulation must be thick to avoid arcing. I have not designed these wiring connections on the support pillars yet. I expect the high voltage wires to be only slightly narrower than the pillars. The wires will need some kind of support or channel to hold them in place on the pillars. Maybe the wires should actually be routed through holes instead of going along the very outside of the ESA.

All of these wires on the outer surface of the ESA must be shielded to avoid field leakage. Any electric field leaking into space around the ESA will influence the trajectories of incoming electrons, corrupting the science data. Also, there is a concern about magnetic fields. The 3.3-V supply wire will be carrying a significant current (~150 mA), meaning it will induce a potentially troublesome magnetic field. The best way to avoid this is carrying the return current in another wire in the same bundle as the 3.3-V supply wire. Then, their magnetic fields should almost perfectly cancel out. When the wires are bundled close together, field strength will follow the much more rapid decrease with distance of a dipole. The difficult part is making sure current can only return along that one path, not through the ESA's structure or any other part of the spacecraft.

Connectors

I chose not to use any wiring connectors on the HVPS circuit board itself, because I felt they were too large and inconvenient. All of the standard connectors I normally work with in the SSEP lab would take up an unreasonable amount of board space. There probably are ultra-small connector options available, but instead I soldered wires directly to the board. Those wires then form a bundle at the edge of the board. I made a mounting point for the wire bundle. It is a metal loop soldered near the edge of the board. This metal loop supports most of the force if wires are pulled, so the wires' solder joints are not mechanically as important.

Vacuum Testing Cable

A cable was needed for connecting the HVPS inside a vacuum chamber. The cable ends are 15-pin D-sub connectors as shown in Figure 7.14 with the pinout description listed in Table 7.2.

*Table 7.2: Vacuum Testing Cable Pinout (*NC = No Connection)*

Pin #	Function	Wire color for outside chamber
1	Neutral- current return from power supply (don't confuse this with ground!)	Black
2	Ground-connects to board but should carry no current.	Green
3	+3.3 V main supply	Yellow
4	+3.3 V for the MSP only	Yellow
5	NC*	
6	NC*	
7	TEST- For debug/reprogram	Brown
8	RESET- Also used for debug/reprogram	Red
9	Misc. debugging pin: Connects to P3.4 on the MSP430	Yellow
10	NC*	
11	I2C Data (Not connected on R2 board)	Red
12	I2C Clock (Not connected on R2 board)	Green
13	NC*	
14	MCP Voltage: Divided and Buffered	Green
15	ESA Voltage: Divided and Buffered	Red
Shield	Ground-connects to chamber at the feedthrough port?	Aluminum Foil

Because I was not terribly concerned with outgassing, I made this cable with ordinary PVC-insulated ribbon cable. Kapton-insulated ribbon cable is very expensive. I shielded the cable with grounded aluminum foil wrapping. For clarification, this is definitely not a cable that would be used for spaceflight because it is so large.

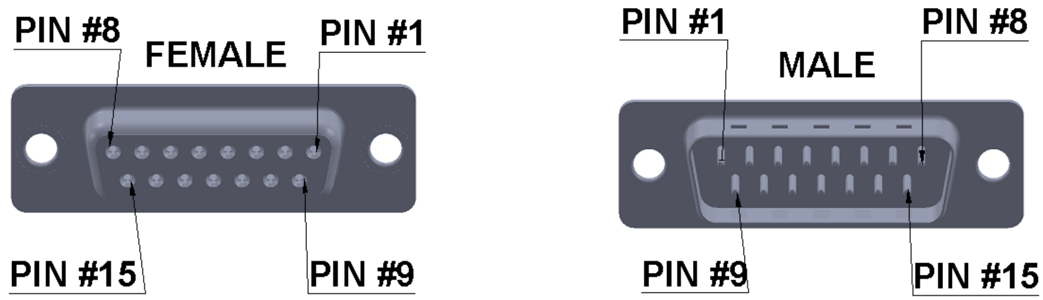

Figure 7.14: 15-pin D-Sub Connector for Vacuum Testing

The Board Shield

For a variety of reasons, I chose to enclose the board with a grounded aluminum ring around the edges as shown in Figure 7.15 below. It is about 0.7 mm thick. This shield should reduce the radiation damage received by the board, especially from electrons. Some higher energy particles will always penetrate, but the overall dose should be reduced compared to an exposed board. It also provides electrical shielding, because a solid wall of metal acts as a faraday cage, and is very good at blocking electric fields from escaping the power supply. This is important because leaking electric fields might influence the trajectories of nearby electrons which ACAPS is trying to measure.

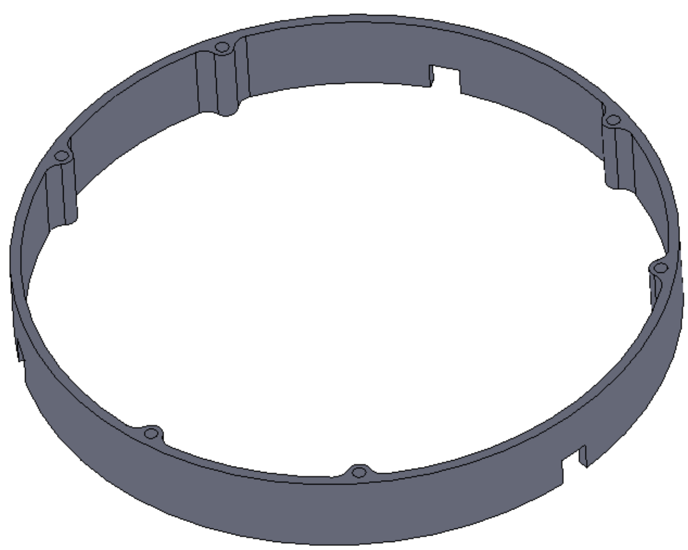

Figure 7.15: Board Shield

The shield also allows more flexible thermal design. The shield can be painted light or dark to optimize the in-orbit temperature range, depending on the results of thermal analysis. The shield will also isolate the HVPS board from direct heating by sunlight, preventing hot spots from appearing on certain components. Finally, the shield stiffens the HVPS board and provides a thick piece of aluminum for the board to be screwed down onto.

Thermal Analysis

A small satellite is usually in an uncontrolled thermal environment, meaning the satellite cannot prevent its temperatures from fluctuating. Temperatures are usually harsh and unpredictable and can change from dangerously hot to dangerously cold within a few hours on a small spacecraft with a low heat capacity. Most CubeSats use completely passive methods to manage their temperatures, except there may be active control for battery heaters. Many small CubeSats have rudimentary attitude control systems which result in wildly unpredictable surface exposure to sunlight and infrared from the Earth. Nothing can be safely assumed about the orientation of different parts of the satellite relative to the sun. Also, throughout a span of several months, a satellite's orbit may change from 60% sun exposure to 100% sun exposure, so even the daily average conditions are variable throughout a mission.

I did a basic thermal analysis assuming the HVPS is thermally insulated from the rest of the satellite. This is a reasonable assumption because the ESA support pillars are made from 3D printed conductive PEEK plastic, not metal. Also, any conductivity they do have will only help moderate extreme temperatures by allowing the HVPS to exchange heat with the satellite. Therefore, a conservative analysis should not consider the pillars at all. Therefore, the only significant heat transfer is by radiation to and from space, the sun, and the Earth.

My conclusions based on a simple spreadsheet analysis (shown in Figure 7.16 on the next page) are:

1. The power consumption by the HVPS is not a large contributor to the overall heat flux. The amount of exposure to sunlight is by far the dominant variable.
2. For absorptivity of 0.7 and emissivity of 0.9, I estimated a maximum temperature of 83 °C and a minimum temperature of -61 °C for the most extreme conditions. I chose the absorptivity of 0.7 to align with the designed range of power supply operating temperatures.
3. The HVPS has such low heat capacity that it can reach extreme temperatures in time scales of around 1 hour. Even at cold temperatures the thermal time constant is only about 90 minutes. For context, a low-earth orbit takes approximately 90 minutes, with up to 40 minutes in the shadow of Earth.

4. I do not believe it is possible to make sure that the HVPS can operate in every possible set of worst-case thermal conditions. However, the extremes I analyzed are unlikely to occur very often. Satellites usually spin, meaning they generally do not stay in the perfect orientation for maximum or minimum sun exposure. If a mission did expect ACAPS to always be facing the sun, the surface absorptivity could be reduced.

Heat Capacities			
HVPS+Topcap Mass	35	g	(Weighed June 2023)
Specific Heat capacity estimate	0.5	J/gK	(Mixture of Aluminum, FR4, and PEEK)
Heat capacity of HVPS	17.5	J/K	
Radiative Time Constant: Hot C	1145	seconds =	19 minutes
Radiative Time Constant: Cold (	5449	seconds =	91 minutes

ESA Top Cap Support Pillars Conductivity			
Width	2.00E-03	m	2mm
Depth	8.00E-03	m	8mm
Height	1.50E-02	m	15mm (conservative)
ESA PEEK Conductivity	0.25	W/mK	
Thermal Conductivity of Pillars	0.00080	W/K	(accounts for three pillars)

Absorptivity	0.7		(moderately reflective)
Emissivity	0.9		(most materials are near 1)
Board Diameter	6.50E-02	m	(65mm)
Board Area	0.00332	m2	(area of circle)
Thermal Emission Area	0.00663	m2	(2 sides)

HOTTEST CASE:			COLDEST CASE:		
(No conduction to the rest of satellite)					
Full sun exposure+Albedo	1700	W/m2	25% Albedo, no direct sun	100	W/m2
Full earth IR	300	W/m2	50% Earth IR	150	W/m2
Solar+Albedo Heating	3.95	W	Solar+Albedo Heating	0.23	W
Earth IR Heating	0.90	W	Earth IR Heating	0.45	W
Electrical Heating (HVPS on)	0.60	W	Electrical Heating (off)	0.00	W
Total Heat Input	5.44	W	Total Heatload	0.68	W
Temperature	356.1	K	Temperature	211.7	K
	83.1	C		-61.3	C
Radiation Conductivity	0.0153	W/K	Radiation Conductivity	0.0032	W/K

Figure 7.16: Basic Thermal Analysis Spreadsheet

Metal Whiskers

Metal whiskers are an insidious kind of needle shaped growth which emerge from a metal surface. Tin-plated electronic components are especially notorious for having whiskers. Thankfully, lead-based solder is mostly immune to whisker formation. Unfortunately, lead is also dangerous to human health and the electronics industry has almost completely moved to lead-free electronics. It is unreasonable to completely avoid tin-plated electronics, so the next best option is to coat the entire board in a thick layer of plastic or epoxy (conformal coating). This is why coating even the low voltage areas is still important, once testing is completed.

Chapter 8. Testing Results

This chapter covers voltage and current output tolerances, vibration, thermal cycling, outgassing, and more. These tests serve the following purposes.

- To verify that the design meets all requirements (Table 2.1 in Chapter 2).
- To give potential users of the HVPS a clear idea of its capabilities.
- To prove that the design "really works".

High Voltage Work Priorities

These are priorities I followed for designing and testing with high voltages.

1. Do not harm people.
2. Do not harm other parts of the spacecraft or test equipment.
3. Do not exceed voltage limits of parts on the board (capacitors, relays, cables, and the PCB).
4. Produce stable, repeatable voltages.
5. Be as power-efficient as possible.

Even though the converter is not likely to output a dangerous level of current continuously, capacitors allow the possibility of a short pulse at high current. The MCP supply has enough capacitance to store a potentially dangerous amount of energy. For safety, I only operated this power supply in vacuum for anything past 1000 V. Using a vacuum chamber not only makes arcing less likely, but it also enforces separation from humans and other equipment. It is difficult to accidentally touch the circuit when a steel-walled vacuum chamber surrounds it on all sides. I also included a safety jumper, which shorts out the power to the high voltage converter, so that it cannot generate high voltage. Finally, the high voltage can be turned off in software, but that is not to be trusted alone.

Testing Environment

Figure 8.1 shows the setup I used for nearly all of this testing, inside of the SSEP lab's thermal vacuum chamber. Heating and cooling happen through the aluminum plate which everything in Figure 8.1 sits on. The picture is taken through a viewport on the front door of the chamber.

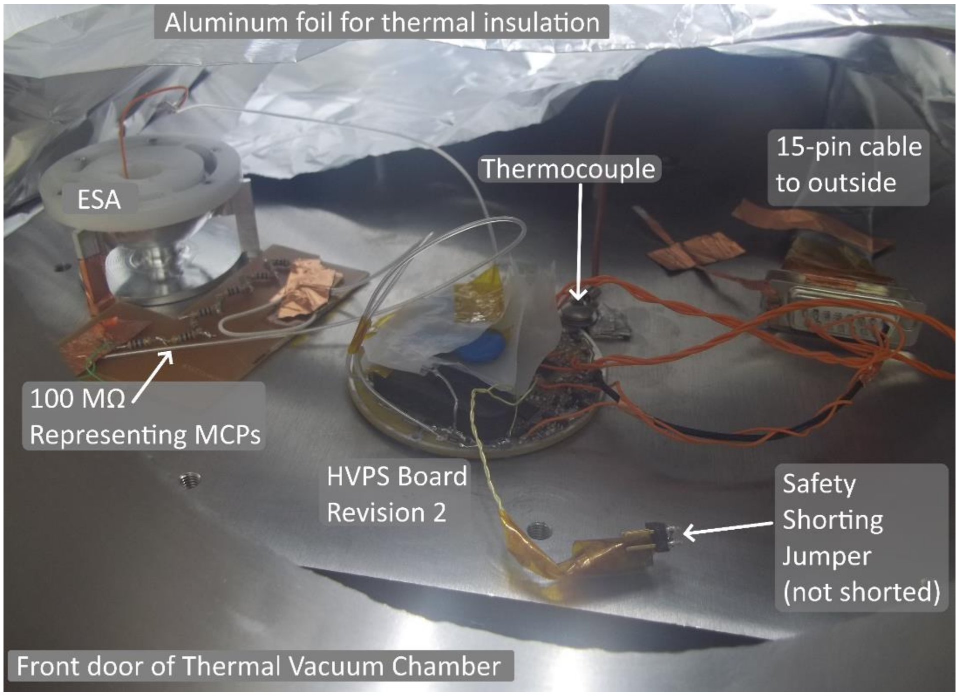

Figure 8.1: Testing Setup inside SSEP Lab's Thermal Vacuum Chamber

Vacuum Baking and Outgassing

This board has had minimal issues with vacuum pressure, probably thanks to good cleaning. Conveniently, good cleaning practices are also helpful for reducing the chance of high voltage arcs. The epoxies of the board and its components are not the main source of vacuum outgassing. Coatings, dirt, condensed vapors, and oils on the surfaces of the board are a much bigger problem. After the board is exposed to a vacuum, the first things to outgas are residual isopropyl alcohol, soldering residues, and water vapor. Whenever a board has been recently soldered or cleaned, the vacuum pressure is elevated for a few hours. I always cleaned the board after soldering because the soldering residues are bad for outgassing and arcing. Heating dramatically accelerates the process of removing contaminants, so I usually baked the board for at least 30 minutes above 100 °C. Then, I let the board cool down to 60 °C before turning on high voltage.

After over two weeks of operating in a vacuum and at least four temperature cycles, the pressure in the SSEP lab's thermal vacuum chamber reached 6 x 10^{-8} torr at room temperature. If we assume the worst case, all of that pressure is caused by the HVPS board outgassing.

The chamber is pumped by a Leybold Turbovac 350i pump, with a pumping speed of 350 L/sec. Therefore, the outgassing rate is at worst (350 x 6 x 10^{-8}) = 2.1 x 10^{-5} torr-L/sec. Since the board has around 100 cm^2 of total surface area, the rate is therefore about 2 x 10^{-7} torr-L / (cm^2-sec). This is a very conservative estimate because it assumes the vacuum chamber does not contribute to outgassing at all. I don't know how to quantify how good this outgassing rate is. However, the fact that the vacuum chamber is still only able to reach about 5 x 10^{-8} torr when empty indicates that the HVPS's contribution to the overall gas load must be small, compared to the chamber's own outgassing and permeation through O-ring seals. In short, I don't believe this HVPS will cause any problems for ordinary high vacuum systems.

Arcing Concerns

There are four main clues I have found relating to arcs on a circuit board.

1. In a vacuum chamber with windows, a short bright flash is often visible. I intentionally darkened the areas where I was looking for arcs, sometimes even using blankets to block out stray light.
2. Arcs almost always contribute sudden electrical noise in any nearby electronics. I have even seen a computer monitor become disabled by arcs from about 1 meter away. An oscilloscope can be set to trigger whenever a voltage exceeds some level. This allowed me to capture the arc waveforms shown in Figure 8.2 below.
3. Sometimes a specific part of the circuit stops working, which provides a clue to where the arc was.
4. Often an arc will trigger a power cycle of the whole circuit board. However, it is hard to distinguish arcing from other issues which might cause a power cycle.

Surprisingly, I have never seen obvious physical damage on a surface from an arc during the design and testing of this power supply. That is probably because the capacitors on this board do not store enough energy for an arc to heat and damage a surface very much. Below in Figure 8.2 is a picture and the corresponding oscilloscope plot from a 2800 V arc.

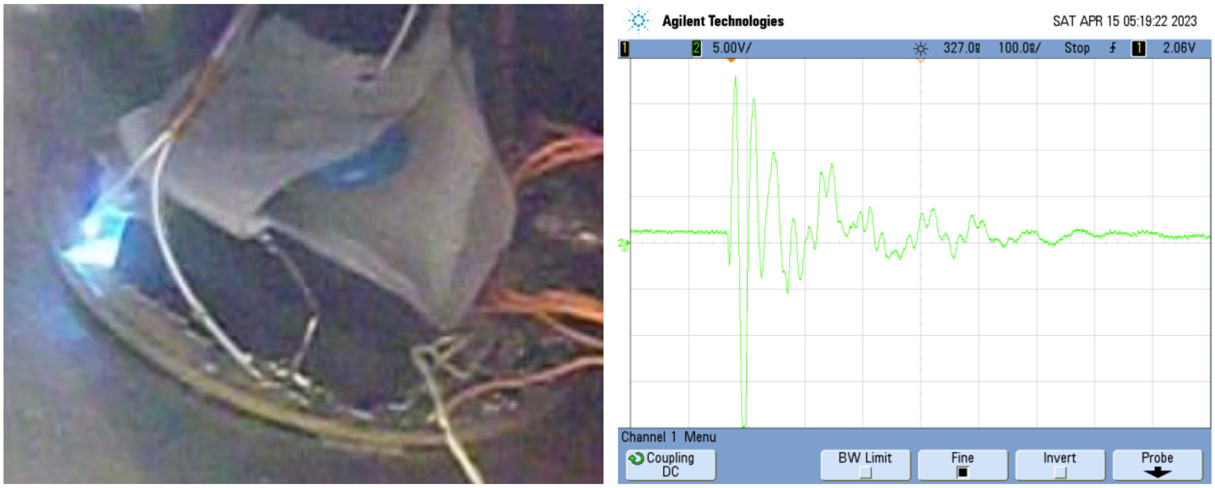

Figure 8.2: Arcing example with picture (left) and oscilloscope plot (right).

The left picture is a single frame of video which captured the arc near the converter's output pin. The scope plot was from the MCP voltage divided and buffered output wire. The energy dissipated in the arc was probably almost the entire energy stored in the 9.4-nF capacitor at the converter output. At 2800 V, that is 0.037 J of energy.

Upon further investigation, I found that the epoxy in that arcing area was damaged and dirty, so I re-epoxied it. Cleanliness is very important for minimizing the chance of an arc. At some point throughout testing, I have fried at least one of every major component on the board. From many arcing incidents, the most commonly fried components have been MOSFETs, especially those MOSFETs which control the high voltage relays. However, the relays themselves are exceptionally robust. I do not think I have ever fried a relay from high voltage. Every time when I have thought a relay failed, it was actually something else, or the relay was damaged on the input side, not the high voltage side. I believe the relays are likely good at surviving because they break down gradually and consistently. I did some experiments with the relays beyond their ratings and found that they just gradually leak more and more current as the voltage passes about 1600-1700 V, but they are not permanently damaged by the gradual breakdown.

Thermal Limits

The final revision (R2) board has been exposed to at least 10 cycles of < -40°C and > +70 °C. This thermal cycling has not yet revealed any mechanical failures. No degradation of electrical properties has been seen yet either. It has spent around one month total in the thermal vacuum chamber, although most of that time has been near room temperature. It has been exposed to extremes of -60 °C and +130 °C.

MCP Voltage Ripple

Figure 8.3 below shows the ripple in MCP voltage that happens every time relays switch on.

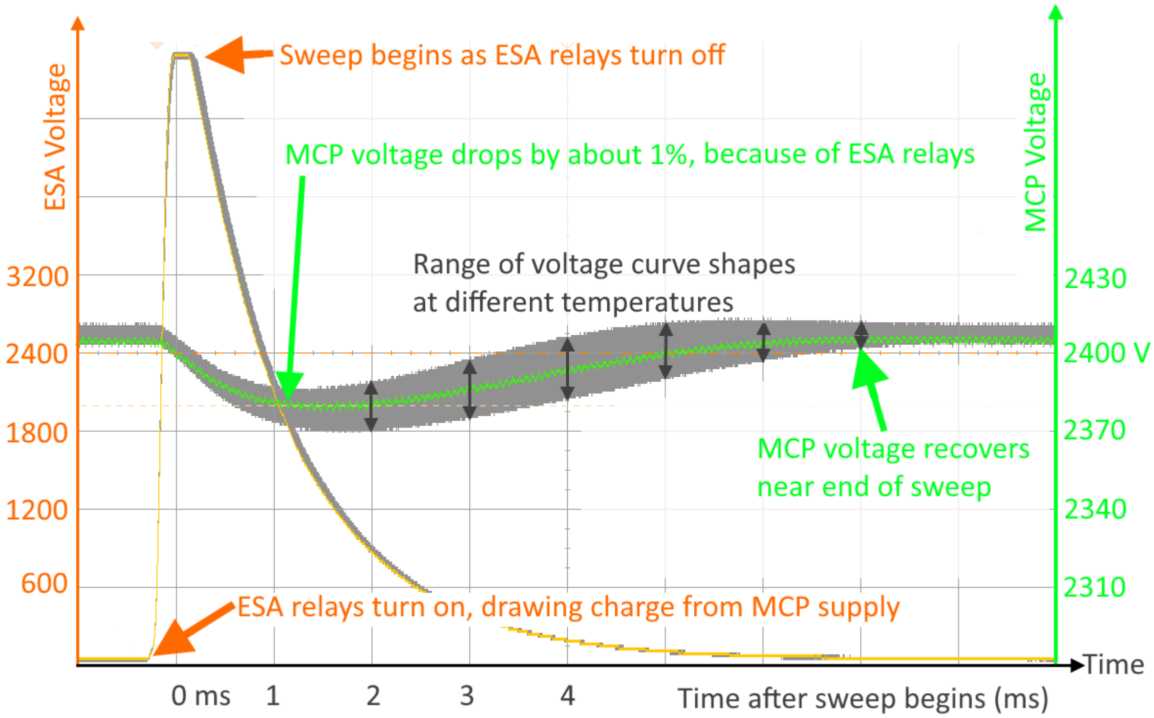

Figure 8.3: Voltage sweep showing MCP voltage ripple range across temperature.

The voltage curves in Figure 8.3 appear noise-free because I was using the oscilloscope's ability to average across 65,536 measurements. Each of those measurements were synchronized with each other by the sharp falling voltage of the ESA sweep. Without synchronizing to the ESA voltage, such averaging is not possible. This sweeping-correlated voltage ripple has a magnitude of about 1% of the total MCP voltage, in the case of 2400 V with the doubler connected to a 14 pF ESA. While it is quite large, this ripple is synchronized with the ESA voltage sweep, so the effect of MCP voltage ripple at any one point in the energy sweep is consistent on a sweep-to-sweep basis. The depth of the ripple and the recovery time definitely depends on temperature. The gray area in Figure 8.3 is the envelope of many MCP ripple waveforms at various temperatures from -40 °C to +70 °C.

MCP voltage ripple which is NOT synchronized with the sweeps is potentially more of a concern, because it will appear like random noise which cannot be calibrated out by correlating it with a certain position in the sweeps. Fortunately, this random component of ripple is much smaller. In fact, it is small enough to be unmeasurable because of measurement noise with the current testing setup. I would need better cable shielding to more reliably measure that part of the ripple which does not synchronize with sweeps.

MCP Supply Drift and Temperature Dependence

This first voltage drift experiment was conducted over only 2 hours with temperature as the only variable. For extra accuracy and better datalogging, MCP voltage data was gathered as .csv files from a Keysight 33470A multimeter instead of the usual oscilloscope. At the start, the board was at room temperature, and I sampled MCP voltage once per second. The 3000:1 divider ratio has been included on the vertical scale of Figure 8.4, which shows a plot of the drift in the MCP voltage offset over time at a constant room temperature. Therefore, 0.1 V on this plot is a variation of about 0.05% of the MCP voltage.

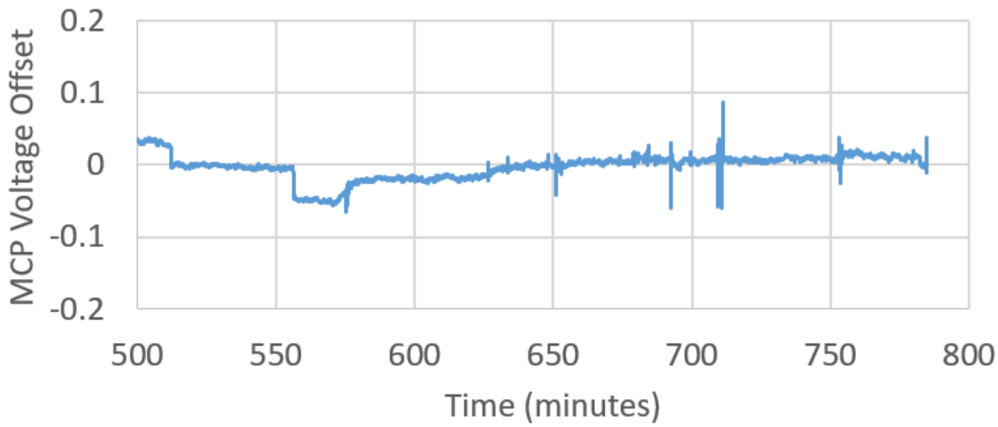

Figure 8.4: MCP Voltage Drift over time at room temperature

This shows that the MCP voltage remained extremely stable (±0.1 V) for five hours at room temperature. This is an example of the very low steady-state voltage drift and shows that there is no significant oscillation on a time scale of seconds to hours. I believe the sharp steps are caused by temperature changes triggering the control voltage (DAC channel A) to change. The average ESA voltage showed a similar level of stability. Next, I heated the board to +60 °C, then cooled it to below -40 °C while measuring MCP voltage. Figure 8.5 below shows changes in the MCP voltage over that 2-hour period of heating and cooling.

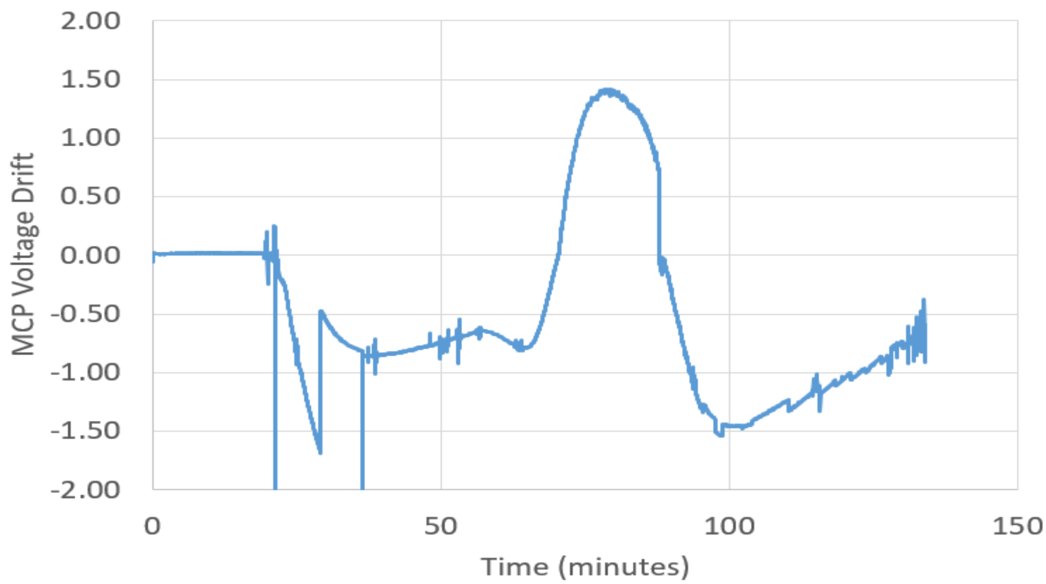

Figure 8.5: MCP Voltage Duirng Heating and Cooling

Once the heating began at about 20 minutes into Figure 8.5, the MCP voltage drifted downward toward a maximum voltage error of -1.7 V from its nominal level of 2400 V. Then, at about 70 minutes, I began cooling the board down toward -40 °C, and the error was similar in magnitude, reaching about +1.4 V around -10 °C, and then the error went negative again toward - 1.5 V at the coldest temperatures. Based on this data, the MCP voltage temperature dependence alone appears to be only about ±2 V or ±0.1%.

In May and June 2023, I left the HVPS board operating for a total of about 4 weeks to test for overall reliability and long-term voltage drift. For 160 hours (the longest period of continuous measurements), I saw less than ±0.5 mV divided voltage variation. This corresponds to < ±2 V MCP voltage drift. The temperature drift is more significant than time drift across hundreds of hours, and both are well below 1% of the 2000 V MCP voltage.

Power Supply Rejection

The two ADP171 regulators perform near-perfectly at isolating the board from noise on the 3.3-V input. There are no graphs to show, because I never found a measurable effect from noise at the input.

Input Voltage Range

I have adjusted the board's nominally 3.3-V input between 2 V and 4.5 V and the board still functions. Of course, at voltages below about 3.1 V, the maximum voltage available to the converter is reduced, so it cannot always reach its full high voltage output. With that exception, the board can function down to around 2.4 V with no apparent issues. High input voltages, like 4.0 V or more, are a bad idea because the input regulator will have to dissipate a lot of power, wasting it as heat. However, it still does manage to regulate the voltage, so the rest of the board can function normally.

Sweep Linearity

Figure 8.6 shows two typical 2400 V sweeps with the doubler turned on. The peak ESA voltage is about 4600 V.

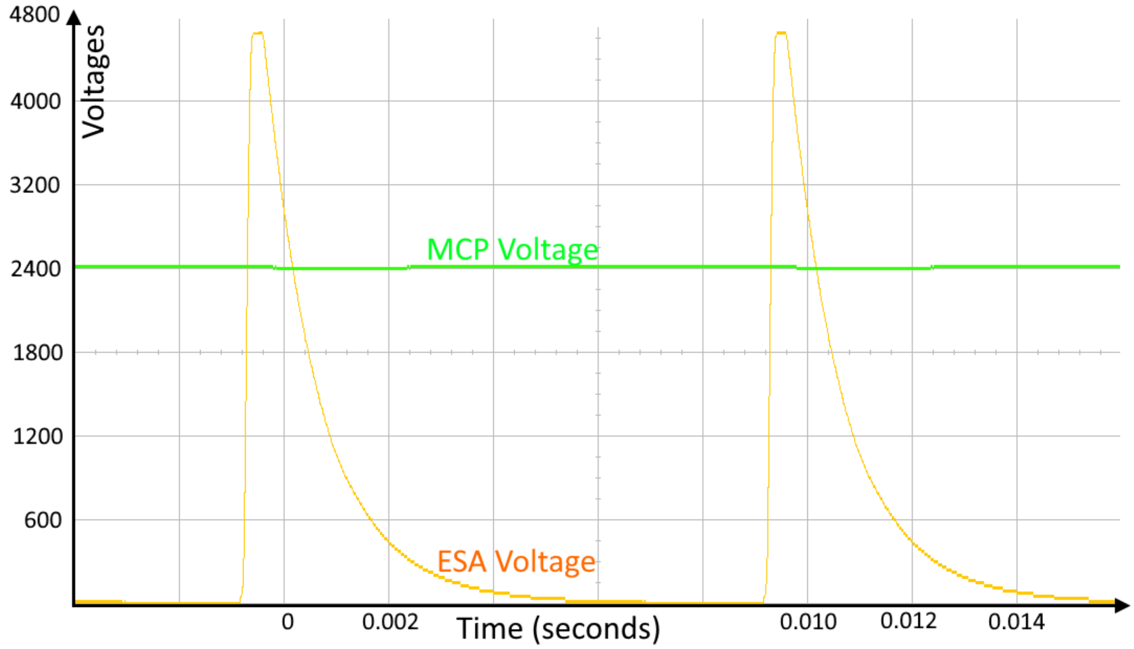

Figure 8.6: Two Typical 2400 V MCP Voltage Sweeps

However, these sweeps are slightly different from a perfect exponential decay. To analyze the sweeps more closely, I exported voltage readings from the oscilloscope as a .csv file. Then, I was able to cancel out the scope's voltage offset error and graph it on a logarithmic scale. An exponential sweep should appear perfectly linear on a log-scaled graph. First, I analyzed with high voltage conditions at high temperature, to get maximum nonlinearity. Figure 8.7 shows a log-scaled view of a sweep with 2400 V and the doubler turned on, at +70 °C.

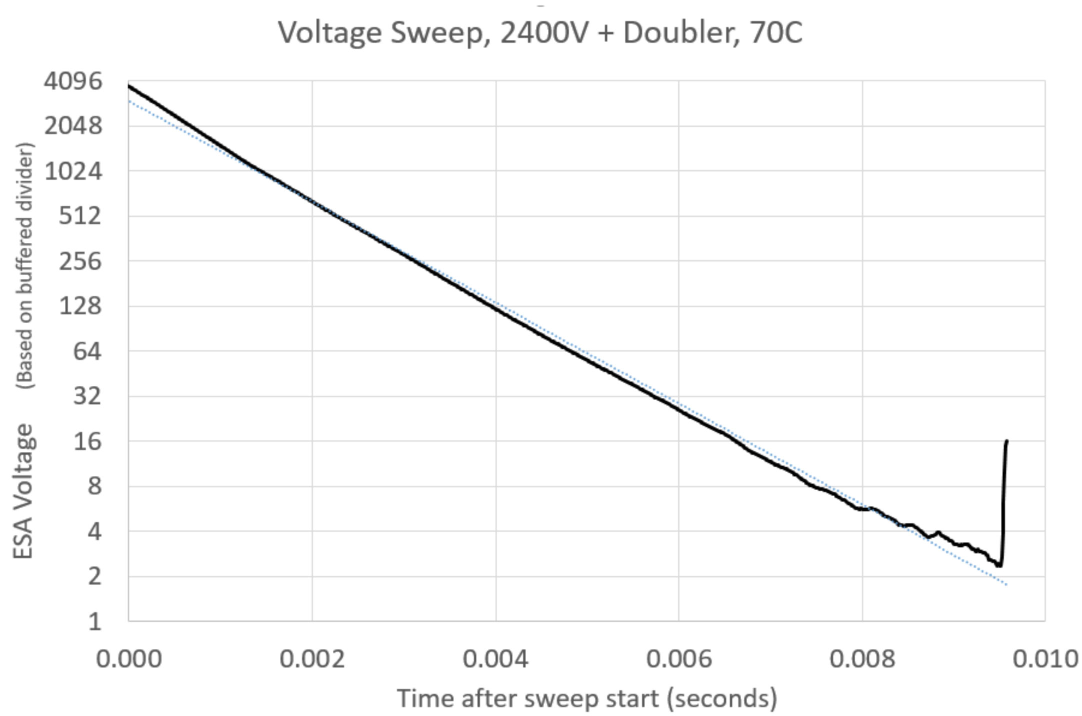

Figure 8.7: RC Sweep Nonlinearity Versus Time (2400 V, Doubler On at +70 °C)

As an alternative way to look at the same data, Figure 8.8 below plots RC decay time constant against ESA voltage.

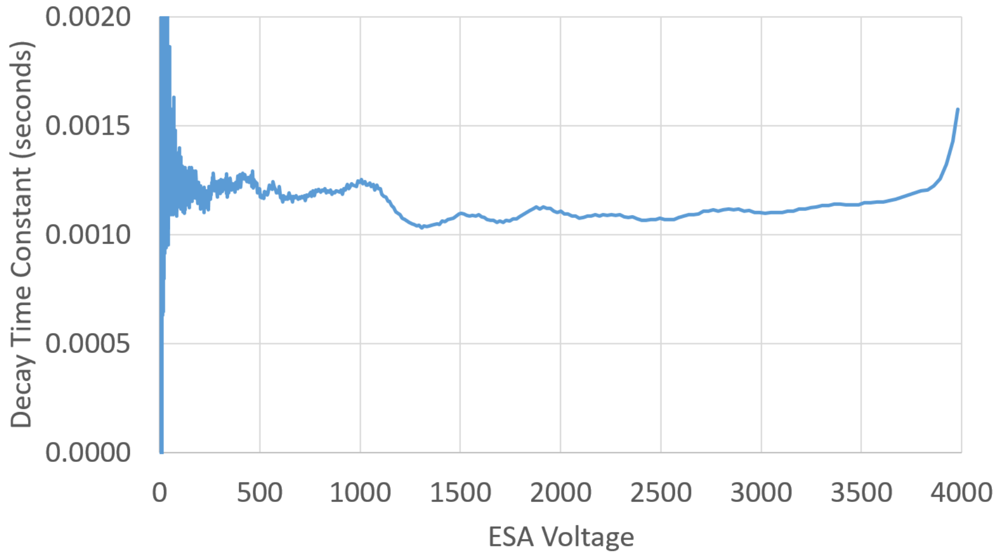

Figure 8.8: RC Sweep Nonlinearity Versus ESA Voltage

At 2400 V with doubler, RC varies by about 10% during the sweep. At the very highest voltages at the start of the sweep, RC is noticeably longer, and I do not have a good explanation for this, except maybe coupling through parasitic capacitance of the relays. At voltages below about 1000 V, the RC is also noticeably longer, probably due to resistor nonlinearity. Finally, at very low voltages, RC increases much higher, but the data also becomes noisier. This increased RC at low voltage is almost certainly due to leakage current through the relays. To investigate this further, I took measurements with the oscilloscope zoomed in on only the low-voltage end of the sweep as shown in Figure 8.9 below.

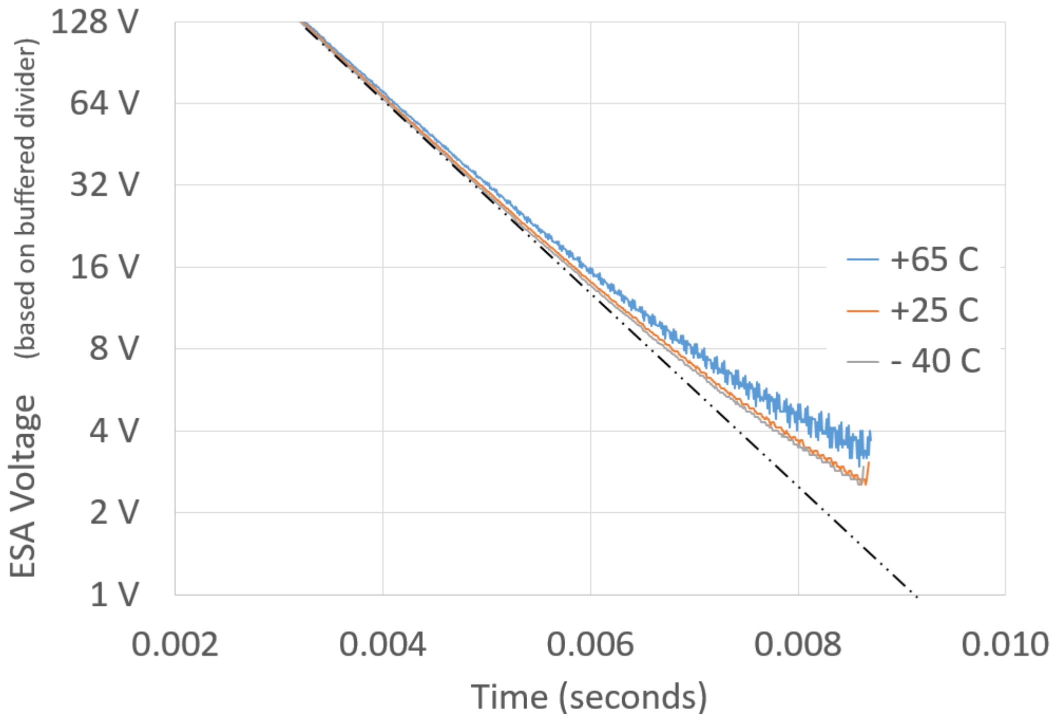

Figure 8.9: Low end of sweep at 2400V with doubler at various temperatures

Figure 8.9 shows that the nonlinearity becomes more severe as voltage decreases below around 10 V. This nonlinearity also shows a temperature dependence, making it more difficult to calibrate for. At +65 °C, the sweep is significantly more nonlinear, probably because of increased leakage current from the relays. Next, I tested with only 800 V and no doubler to see if the leakage would be different as shown in Figure 8.10 below.

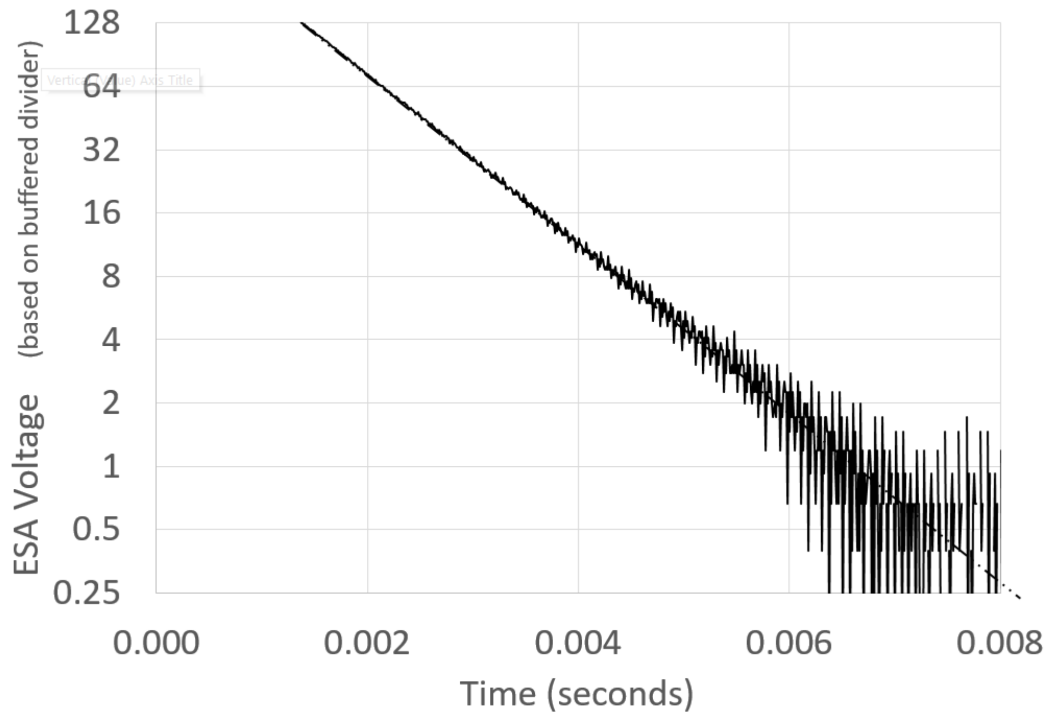

Figure 8.10: Low End of ESA Voltage Sweep at 800 V without doubler, room temperature

Figure 8.10 shows dramatically improved linearity from simply using a lower voltage. This makes me even more confident that ESA relay leakage is the cause of the end-of-sweep nonlinearity.

Sweeping Output Parasitic Capacitance

Based on the measured time constant of 1.1 ms and the sweeping resistance of 52.5 MΩ, I calculated an effective RC sweeping output capacitance of 21 pF. I have measured that the ESA and its wires are 14 pF. This implies the remaining 7 pF comes from the HVPS board itself. This is a combination of parasitic capacitance from resistors, relays, and the circuit board material.

ESA Voltage Comparator Testing

The comparator produces a rising voltage edge when the ESA voltage sweep passes through the comparator's reference voltage (which is defined by DAC channel C). This rising edge for a single sweep is shown in Figure 8.11. Note that the comparator threshold was set just below the ESA peak sweep voltage.

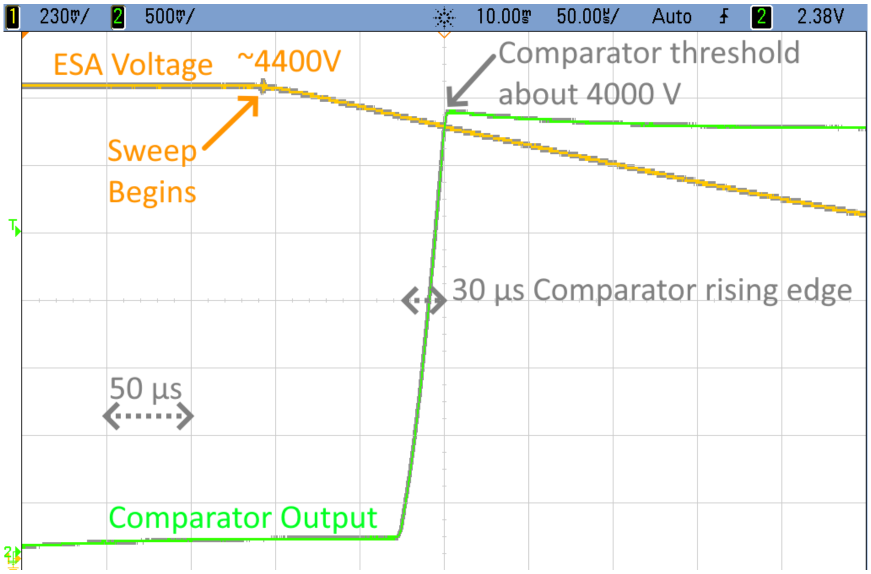

Figure 8.11: Comparator Rising Edge with ESA Voltage Sweep

The slow rising edge shows that unfortunately, this OP747 is doing a poor job of role-playing as a comparator. I should have used a dedicated comparator chip, instead of pushing an instrumentation amplifier into this job.

I also measured the sweep-to-sweep variation in the comparator output. While it is slow to rise, the rising edge is surprisingly consistent, as seen below in Figure 8.12. Note that this oscilloscope plot is much more zoomed in than Figure 8.11.

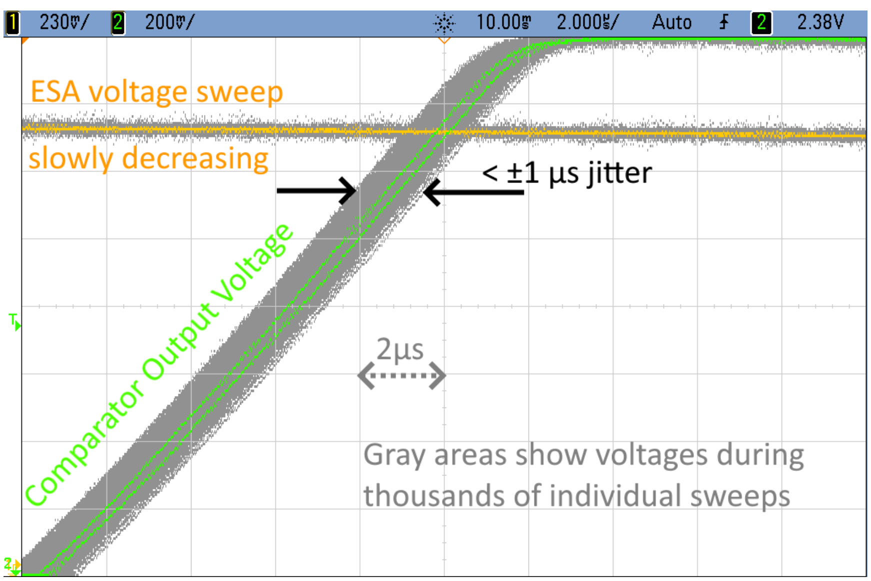

Figure 8.12: Comparator Jitter, Sweep-to-Sweep

Sweep Timing

To measure the consistency of timing across many sweeps, I programmed a debugging signal to go from high to low every time a sweep begins. I set the oscilloscope to trigger off of this signal, and then display the voltage waveform 1 second (100 sweeps) into the future.

Figure 8.13 shows that the duration of any 100 sweeps may vary by about ±70 ppm compared to another set of 100 sweeps. Note that clock jitter across a short time like 1 second is different from long-term clock drift. Across longer periods of time, I measured up to about 300 ppm of total clock error, across a full range of temperatures.

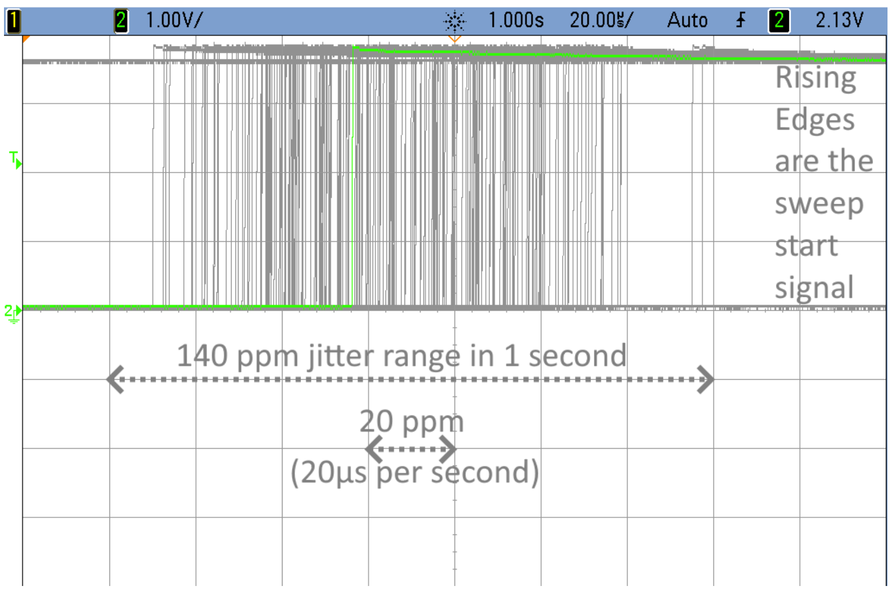

Figure 8.13: Relay Sweep Clock Jitter across 1 second

Power Consumption

When idle with all high voltages turned off, the board consumes about 4 mA (~12 mW). It also consumes a similar current when protection features are forcing the board to repeatedly reset (such as overheat protection when it is too hot). I did not spend much time analyzing idle power because it is so low. The host spacecraft can simply cut power to the HVPS completely if 12 mW is a problem.

Active Power

All of the power consumption data shown in Figure 8.14 was gathered on the same day with the same code loaded onto the MSP430. The only difference between lines of measurements is the voltage target setting.

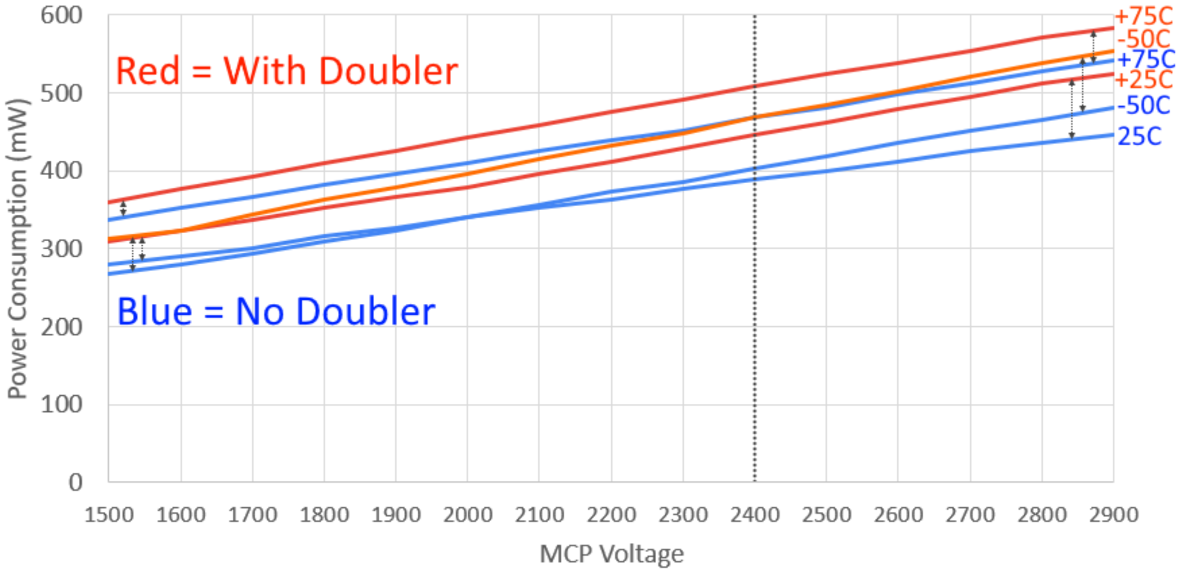

Figure 8.14: Power Consumption Versus MCP Voltage at Various Temperatures

Figure 8.14 shows that power consumption increases very linearly with increased output voltage. Temperature has a noticeably effect, similar in magnitude to the difference in operating with or without the voltage doubler. At the nominal voltage of 2400 V with the doubler, power consumption is around 475 mW. This data point will be analyzed in the next section.

Power Consumption and Efficiency Analysis

A basic power consumption budget and analysis of the HVPS is outlined in Table 8.1 on the next page. It is difficult to define what efficiency means in the context of power consumption, because it requires defining what the useful output is. If we take the strict definition that only power delivered off the board to the MCPs and ESA counts as useful, the efficiency is 16%. That is a measure of the efficiency of the HVPS as a whole. However, if we include all known consumers of high voltage, then efficiency is 25%. That is more a measure of the converter and regulator's efficiency.

The impact of using the voltage doubler at 2400 V is only between an 8 and 16 % increase in power, depending on temperature. I think this is because the doubler relays are already tuned to be on for near the minimum duration possible. This power consumption is not absolutely too high, but lower power would be helpful. What counts as an acceptable level of power consumption depends on how often the HVPS is turned on. A small spacecraft might only be able to spare 0.1 or 0.2 W on average without draining its batteries. However, ACAPS is not expected to operate at all parts of its orbit, only in areas of scientific interest near aurora. Also, ACAPS won't realistically operate constantly on an orbiting satellite because it would generate far too much data to send down to the ground. I estimate a data gathering rate of up to a few megabits per second.

Table 8.1: Basic Power Analysis

Power Type	Power (mW)	Details
Idle Consumption	16	5mA at 3.3V. (Protections, DAC, MSP, opamps, charge pump, etc)
Doubler Relays	3	10% duty cycle at 10 mA, 3 V
ESA Relays	6	10% duty cycle at 10 mA, 6 V
(HV) MCP Load	58	100 MΩ at 2400 V
(HV) MCP Divider	19	300 MΩ at 2400 V
(HV) ESA sweeping	15	15 pF measured at 4500 V, 100 Hz
(HV) ESA Parasitic Capacitance	7	7 pF calculated at 4500 V, 100 Hz
(HV) Doubler Parasitic Cap.	1	4 pF (assumed) at 2400V, 100 Hz
(HV) Doubler Low Side Resistors	13	5% duty cycle of 2400 V on 22.5 MΩ
Explained Total (Low Voltage)	25	First three rows only.
Outputted Total (High Voltage)	73	This only counts power which leaves the board, going to MCPs and ESA
Explained Total (High Voltage)	113	All high voltage power, including doubler, dividers, and parasitics.
Explained Total (All)	138	Power with a well-known destination, although some of it is not useful
Unexplained Power at High Voltage (Goes through converter)	312	475 mW avg. measured – 138 mW explained. Majority of this is probably lost inside the regulator and converter.
High Voltage Efficiency based on only MCP and ESA outputs	16%	73 mW MCP+ESA high voltage delivered / 450 mW high voltage input.
High Voltage Efficiency based on all known HV consumers	25%	113 mW high voltage explained / 450 mW high voltage input

Electron Beam Testing:

Once the R0 board was functional, I tested it with an electron beam on top of an ESA. After all, that is the purpose of this power supply. This ESA is designed to detect electrons, so I set up an electron beam at 2000 eV and aimed it at a 3D printed PEEK plastic ACAPS ESA. Below the ESA was a Faraday cup detector and an amplifier used for measuring the number of electrons detected at any moment in time. Above the ESA was this HVPS board R0. Figure 8.15 shows the HVPS mounted on top of an ESA and exposed to an electron gun in a vacuum chamber.

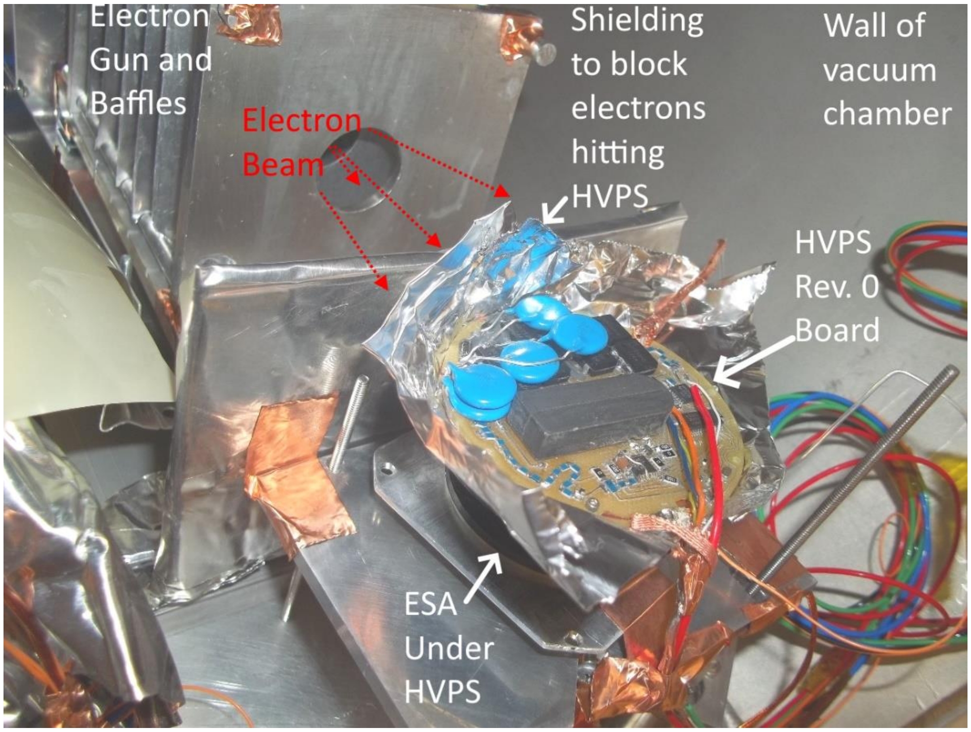

Figure 8.15: Prototype HVPS board with electron beam.

Figure 8.16 shows the resulting signals from the electron beam test setup.

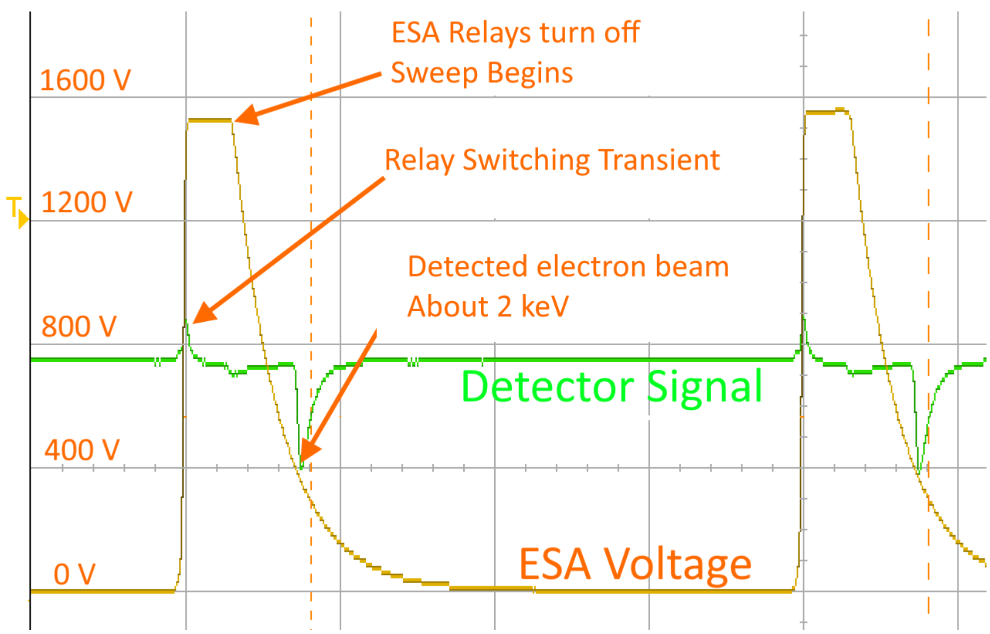

Figure 8.16: Detector Signal with Voltage Sweep

The ESA voltage sweep is the yellow trace. The green trace is the detector output signal. Normally, the detector signal is stable near zero, but there is a repeatable spike whenever the ESA voltage is approximately 350-400 V. The spike is negative because electrons are negatively charged. The location of this signal spike at around 350-400 V on the voltage sweep makes sense, because the ESA's energy/voltage ratio is about 5 or 6. This corresponds with the electron beam energy of 2000 eV. The short rise and fall time of the detector signal shows that the ESA detected electrons only within a narrow range of energies during the sweep, exactly as expected. Figure 8.16 is important, because it proves that this ESA sweeping supply can be used for detecting electrons with a real ESA. It proves that the combination of a voltage sweep and a detector signal does allow the electron energy to be derived.

During electron beam testing, there were occasional interruptions in the sweep voltage, probably caused by arcs on the surface of the circuit board. They only happened when the electron beam was active. I assume the electrons were bombarding the surface, causing plasma to surround the board surface. This plasma would have made it easier for arcs to form. While this beam testing was informative, in the future the HVPS board should be enclosed in metal shielding to block the electron beam from directly harming the circuit board.

Despite the probable arcing, I was able to operate the R0 board under the electron beam at voltages up to around 2000 V for several hours, until it was damaged. Eventually, several MOSFETs which did not have TVS protection were fried. The chamber shown above in Figure 8.16 has no windows, so I do not know exactly where the arcs may have happened. Instead of fixing the R0 board, I moved on to building board revision R1.

Chapter 9. Dead Ends and Unfinished Ideas

Adjustable Sweep Rate

Throughout most of this design process, I intended to allow for tuning the speed of the exponentially decaying output voltage. Using software, it is easy to alter the timing of solid-state relays. However, changing the RC time constant of that exponential sweep would require tunable hardware to alter the capacitance or resistance.

Changing the capacitance or resistance encounters two major problems:

1. The tunable part must operate at an unusually high voltage. For example, to get a ±5% adjustment range means the component (resistor or capacitor) must withstand at least 10% of the full output voltage, or about 500 V.
2. It must be linear across many orders of magnitude. If a tuning component has different resistance or capacitance at 0.5 V compared to 500 V, the contribution to the R-C time constant will change throughout the sweep. That would be very bad because the sweep would not be a consistent exponential decay. It would become harder to predict accurately, and some parts of the voltage range would get too much or too little time.

First, I looked at tuning resistance.

Can the relays be used for resistance tuning?

The AQV series solid state relays are a sort of optocoupler, so I thought maybe they can operate in the linear region to behave as resistors. Unfortunately, they are designed to "snap" very suddenly from off to on. Even in a benchtop experiment at constant room temperature, I was unable to reliably tune their resistance. The difference between on and off was only a few mV on the input side, and that threshold is hard to predict.

Peltier + Thermistor Adventure

Next, I considered adjusting the resistance of thermistors using a thermoelectric cooler/heater (also known as a Peltier module). Software would heat or cool the Peltier module, and the thermistors would sit on the Peltier module, bonded to it with thermally conductive epoxy. I ran several experiments across a wide range of temperatures using actual voltage sweeps from relays, and the results were promising. Ordinary 5-MΩ through-hole thermistors were able to withstand 500 V each, so a stack of 12 gave 60 MΩ and was able to linearly adjust the resistance for a full range voltage sweeps. The oscilloscope plots in Figure 9.1 and Figure 9.2 show that the resistance was tunable across a wide range, and the sweeps it produced were very linear.

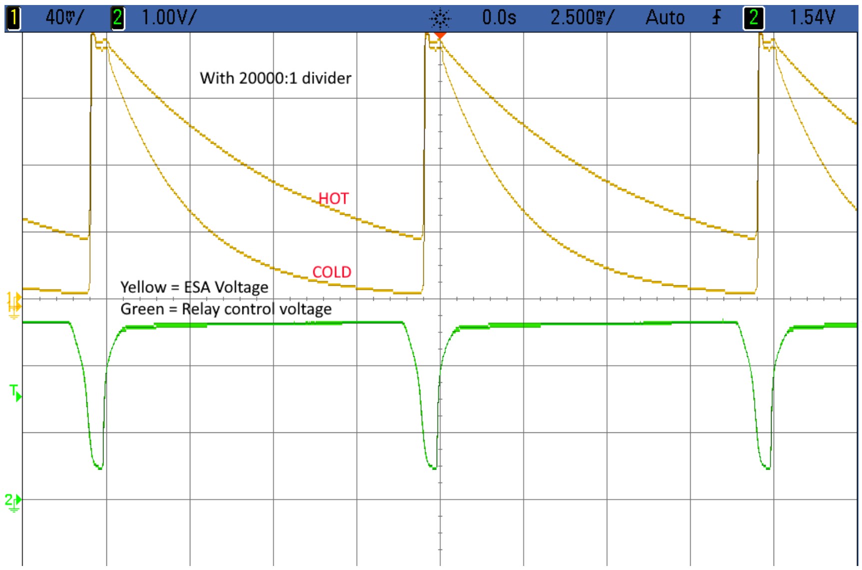

Figure 9.1: Peltier + Thermistor Sweeps of two RCs

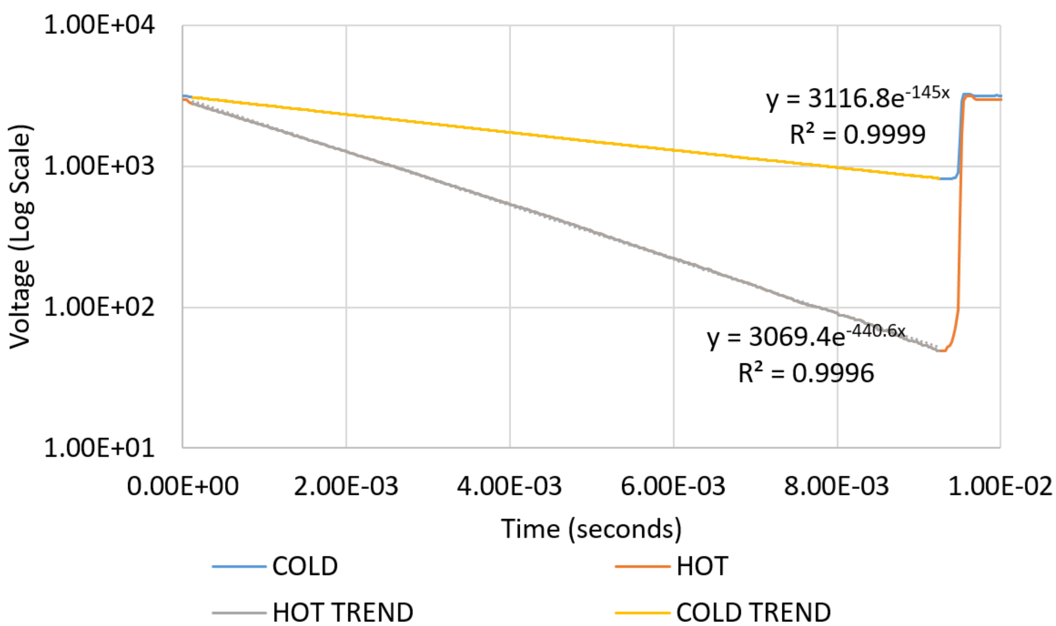

Figure 9.2: Peltier + Thermistor Sweep Linearity

Figure 9.2 shows good linearity from 3000 V down to 50 V. Because thermistors are so sensitive to temperature changes, I was able to tune the sweep's RC by a factor of 5 or even more. This method had great potential because it could allow for massive changes in the sweep rate. Changing RC would allow operators of the spacecraft to flexibly decide what sweep rate they really need at any given time.

One way of thinking about this Peltier and thermistor combination is as a thermal transistor, where the output resistance is related to input current. The output resistance linearity across voltage is far better than any MOSFET or BJT I have worked with.

However, the Peltier + thermistor method had a fatal flaw. Controlling the temperature is expensive in terms of volume, mass, and power. The additional volume (about 2 cm^3) required for the Peltier module and thermistors was a barely manageable disadvantage, but the electrical control requirements were actually even worse. Peltier modules have very low resistance, meaning they demand high currents. It was not practical to drive the Peltier directly off the main 3.3 V supply. Instead, it would have required a high-current switching regulator on the HVPS, including the ability to swap polarity to change between heating and cooling. It is probably possible to fit such a commercially available regulator on the HVPS board, but I decided it wasn't worth the trouble. Even with good insulation around the Peltier module, power consumption would be high because of heat leakage through the module itself.

Why not tune RC with a transistor?

While components that can switch 3000 V or more are generally large and difficult to implement, switching a few hundred volts is easier with a single ordinary transistor. Specifically, I looked at MOSFETs, because their output can behave somewhat like a resistor, and they are easily controlled by DACs.

I found an option at 500 V that seemed reasonable, the FQ1N50C N-type MOSFET. It is physically small, being a TO-92 package that can easily fit into a few mm of vertical space above the board surface. It is designed for much higher currents than this supply will ever require. Being rated for 500 V means it would allow ±5% tuning of the sweep at most, since 10% of the 5000 V total can be across this MOSFET. The rest of the voltage would still be sustained by ordinary resistors.

After experimenting with this MOSFET, I realized some problems and decided against it.

- The MOSFET cannot be in series with the ESA voltage measurement divider, because adjusting the MOSFET's resistance must not adjust the divider ratio. Instead, it would need to be in a separate stack of resistors. Using two stacks of high voltage resistors is possible, but it is definitely a board design complication.
- MOSFETs are not linear enough to provide a reliable resistance adjustment. They do not have consistent resistance as their drain-source voltage changes. RC would not be constant throughout the sweep. This is a similar problem to optocouplers, where their Current Transfer Ratio changes with voltage, so they have to be under closed-loop control.
- MOSFET leakage current is severe enough to cause significant power consumption and limits the sizes of resistors. Special MOSFETs may be needed.

- It is difficult to tune accurately, because a small increase in gate voltage causes a dramatic reduction of resistance. Also, the gate threshold voltage is very sensitive to temperature and radiation damage, so it would have to be dynamically adjusted in space.

Figure 9.3 below shows how the FQN1N50C 500V MOSFET does not behave as a consistent resistor across even one order of voltage magnitude.

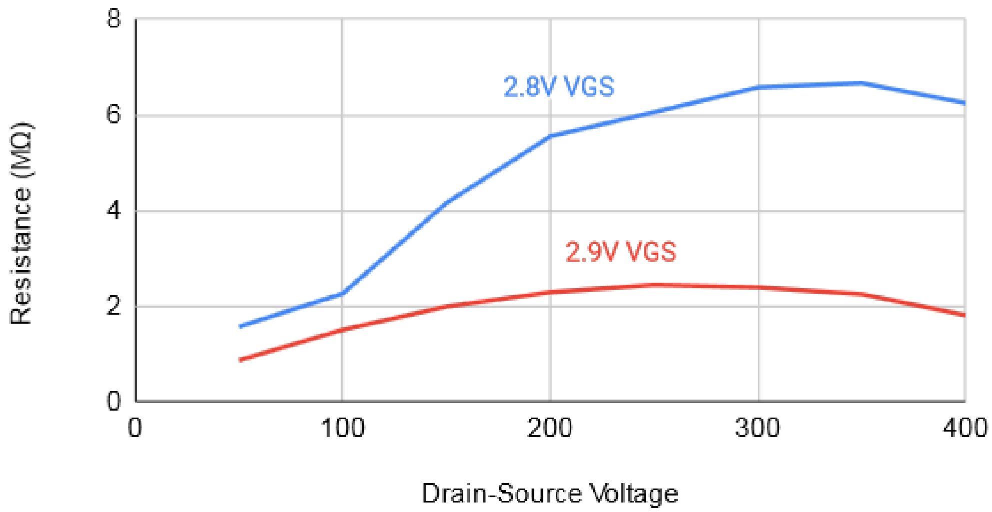

Figure 9.3: MOSFET inconsistent resistance

Although the MOSFET does not work, perhaps there is a technique I am missing. I have heard discussions of JFETs being more linear than MOSFETs.

Capacitance Tuning

Tuning the capacitance is also very difficult. Here on earth, there is an option for very linear, high voltage, finely tunable capacitance, the mechanically adjusted rotary capacitor. These are used for tuning knobs on radios, for example. However, they are completely impractical for use on a small spacecraft, because they require physical rotation, so a motor would be needed. Also, they have very poor density in terms of charge storage per volume, because they use vacuum as the dielectric material.

Another widely used method of tuning capacitance is a **varactor diode**, which is similar to a normal diode but the capacitance depends heavily on how much DC voltage is applied to it. This means a varactor cannot possibly work for a sweeping supply, because the RC time constant is supposed to stay consistent while voltage changes by several orders of magnitude during a sweep. They are great for working with high frequencies superimposed on DC, but not here.

Finally, I realized that tuning the sweep by only a small increment is probably not even helpful. Instead of actually adjusting the sweep, the true sweep voltage can be measured and compensated for using software in the detector electronics.

Active Thermal Control

Early in the design, I believed it would be practical to heat the critical parts of the circuit to a uniform temperature. This was inspired by learning about ovenized crystal oscillators and thermoelectrically cooled optics.

I decided against active thermal control for these reasons:

- It simply isn't necessary. As I learned more, I realized components (especially resistors) are available with very good temperature coefficients, and software can compensate for most of the remaining thermal errors.
- Thermal insulation would become a major concern to keep heat leakage acceptable. Even with excellent insulation, thermal control power might still dominate.
- A thermoelectric (Peltier) module can both heat and cool, but it is more complicated to control, takes more board space, and is only efficient for small temperature differences.
- A heating-only design would be very simple, but the target temperature would have to be very high, perhaps above 60 °C. A lower target temperature would risk sometimes being exceeded by natural heating in orbit, so the temperature would not be constant. A high target temperature would degrade the electronics.

Vibration

I have yet to perform a full vibration test on the board as of July 2023. I plan to do a test including the HVPS board, the shield, and the ESA. The main concern on the board is the reliability of epoxy joints which hold the large 3-kV disk capacitors and the converter itself. The converter is concerning because it is so massive, and it has only epoxy and five 1 mm x 1 mm solder pads holding it to the board. I have intentionally dropped the board on the floor many times, as a way to test for bad solder joints.

Table 10.1 below lists the HVPS requirements, their status as of July 2023, and some details.

Table 10.1: HVPS Requirements, Status, and Details

Requirement	Status	Details
Supply an exponential voltage sweep across the range of < 10 V to > 5000 V, at least 20 times per second, to a load of 10 pF to 50 pF. Sweep must be repeatable within ±5% of nominal voltage at all points throughout the sweep.	Almost Success	Achieves all parts of this requirement except for the uppermost voltage. Current iteration is tested and trusted at only about 4600 V maximum.
Supply a DC voltage which is tunable between 1500 and 2500 V, at up to 40 μA. Must be stable within ±1% voltage across all operating conditions and all current loads.	Almost Success	ESA sweeping causes ripple voltage that is sometimes beyond ±1%. However, it is repeatable on a sweep-to-sweep basis. This design is downrated to 2400 V because of reliability concerns. MCPs most likely will not demand 2500 V anyway.
Consume only standard voltages of 5 V and 3.3 V from the host CubeSat. Consume less than 1.5 W average power during operation.	Success!	Only uses 3.3 V at around 0.5 W.
Do not apply any mechanical demands on the host spacecraft, beyond those already imposed by the ACAPS spectrometer (ESA).	Success!	Attaches to ESA, does not directly interfere with the rest of the spacecraft. However, securing and shielding the wiring is an unresolved challenge.
Survive a range of -60 ° C to +100 ° C. Operate within specifications from -40 ° C to +60 ° C.	Success!	Thoroughly tested in thermal vacuum chamber.
Survive the vibration environment of a rocket launch, based on the NASA GEVS standard.	Untested	
Recover safely from high voltage breakdowns at the outputs.	Success?	Breakdowns (arcs) at the output either trigger a board reset or the voltage simply recovers. No high voltage components have been damaged by an output arc during testing. However, it is difficult to be certain that it will always recover safely. I have not rigorously tested with intentional shorts on the output.

What went well?

Parts of the HVPS design that were successful are listed and described below:

- Fitting the circuit onto the limited board area went far better than expected. I think it was due to combining the two supplies to use a single converter, simplifying the voltage doubler, and using very unorthodox techniques to pack components tightly on the board.
- Testing and prototyping circuits at high voltage inside a vacuum chamber was a great success, both from a safety perspective and in terms of performance:
    - Vacuum dramatically reduces any chance of people touching high voltage. It provides a psychological barrier to avoid blurred lines of safety. I did not have any safety incidents.
    - Vacuum allows higher peak field strengths than open air, so the circuit board design can be denser and higher performance. It is a spacecraft power supply, it should not need to work in air.
    - Operating in a vacuum during early prototyping helped me be more aware of thermal concerns because convection heat transfer was always absent.
    - Testing frequently in a vacuum led to a culture of smoothly transitioning between prototyping and testing, instead of the usual attitude where testing is seen as a separate phase at the end of a design effort. I never went through an explicit "testing phase" or a "designing phase" because those two activities were constantly blended together.
- Prototyping with copper tape on bare circuit board material was very useful. It was much better than having to make a real board during early stages of design, also better than breadboarding, especially at high voltage.
- Manufacturing circuit boards on the SSEP lab's CNC milling machine was also very useful. It allowed me to begin testing a board within one day of designing it.
- Driving both parts of the supply from a single DC-DC converter was successful. The doubler turned out to be reasonably efficient, despite the fact that it uses resistors on the low side instead of relays.
- Widespread use of filtering and TVS diodes was decently effective at minimizing fried electronics.

The relays, diodes, and high voltage capacitors have all impressed me with their reliability.

What did not go as well?

Parts of the HVPS design that did not go well are listed and described below:

- The XP Power (Emco) A-series high voltage converters turned out to be far less ideal than I first thought. They are disappointingly inefficient at low loads. If the converter and regulator combination was 50% efficient instead of 25%, the total power consumption would be 251 mW as opposed to 475 mW at 2500 V with the doubler. I believe that is a reasonable goal to pursue.
- I neglected MCP current limiting. There should probably be hardware based MCP current limiting in this design, just like how there is overvoltage protection. As it is, only software is capable of responding to a high MCP current.
- I spent a lot of time trying to make software solve hardware problems, such as managing the converter's control pin.
- I wasted a lot of time troubleshooting and repairing fried components (especially MOSFETs), where the damage could have been avoided with better use of TVS diodes, better board cleaning, and more disciplined vacuum baking.

How reliable will it be?

The only way to prove reliability is to actually operate the circuit in outer space. Sadly, there is a dilemma: Until a design has been proven in space, people will be afraid to send it to space because it is unproven. This means the path to trusted reliability likely is flight on a single low-budget CubeSat, or as a ride-along on a sounding rocket. Below is a summary of the reliability situation:

- Performance across a wide range of temperatures has been tested well in the lab. The temperatures on a spacecraft near earth should generally be less severe than what I tested with.
- Operating lifetime is uncertain. This power supply has only operated for about 1-month total time (~700 hours). In the future, it will be left running in a vacuum chamber for much longer. There may still be issues with long-term degradation of the insulators, but every day that becomes less of a concern.
- The protection circuits are likely to reduce the risk of catastrophic damage compared to most student-built CubeSat circuit boards, which generally don't have any protections.
- Radiation tolerance is totally uncertain. It is extremely difficult to test, and I don't have access to any high-energy radiation beams. This is actually normal. Most CubeSat payloads are flown into space with minimal regard for radiation. Surprisingly, this strategy of simply ignoring the problem often works just fine in Low-Earth Orbit for small, low-cost missions of 1 year or less. I believe I have designed some circuits that are radiation tolerant by design, but there are surely still numerous weak points that I don't even know about.

Future development of this power supply would ideally include testing the entire circuit board under a high energy radiation beam, ideally using heavy ions instead of just protons or gamma rays. Many components are not radiation-qualified in any way, and there is uncertainty about how the system may recover from faults. It is nearly impossible to simulate radiation faults, because I have no idea how the silicon really works inside the microcontroller, transistors, and ICs on the HVPS board. The best radiation proxy I know of is extreme temperature, but only as a proxy for cumulative radiation damage, not single particle events.

What should be done before it goes to space?

I would not be comfortable launching this HVPS into space without making the following changes.

- **Switching regulator for better efficiency**: I estimate the ADP171 linear regulator accounts for about 1/4 of the total power consumption. Small adjustable switching regulators are available with much higher efficiencies.
- **Longer strings of components**: This means more resistors in series to reduce the nonlinearity seen at the highest voltages. Also, more relays in series would help to reduce the leakage seen at the low end of a sweep. There is a small amount of extra board space available, with some careful rearranging. This will probably make epoxy encapsulation even more critical.
- **More resilient I^2C communication**: There should be a hardware watchdog circuit monitoring the I^2C data and/or clock pins. This would reset the entire board if either the DAC or MSP430 gets stuck and communication between them breaks down.
- **Resolve regulator concerns**: The ADP171 regulators showed a frightening tendency to enter latch up when severely over currented. Considering this, my design's overcurrent protection scheme is probably not enough to assure safety. I would not feel comfortable launching it into space with these regulators, even on a sounding rocket.
- **Adjust overvoltage and overcurrent**: These thresholds should be tuned more precisely. Overcurrent should be just barely above the normal current level. Overvoltage should be something like 2500 V, not 2800 V.
- **UART communication**: Currently the software is not yet written for making this HVPS board communicate with other parts of a spacecraft. The necessary hardware exists on the board.
- **FRAM usage**: For better radiation resistance in space, the software should use FRAM for long-term and short-term memory. The hardware is able to do this.
- **Software resets**: The software should intentionally trigger a full board reset in certain conditions, such as failure to generate voltage or anomalous current readings.

- **Mitigation for tin whiskers**. Even though I used lead solder, whiskers are still a concern because most electronics I used arrive tin-plated from the factory. The only way to be sure is coating everything in a thick layer of plastic. NASA experiments have shown that board coatings of a few microns in thickness are able to block whiskers [31]. Whiskers can actually burst through thinner coatings.

More Exotic Ideas:

- **More efficient High Voltage Converter**: There is a competitor to XP Power/EMCO called HVM Electronics. Their UMHV series of converters are ½" cubes with much higher efficiency at low load compared to the XP Power A-series. In most other ways, they are similar. I believe the HVM series are a better fit for ACAPS's needs in every way except for their height. The height concern would be removed if the power supply was redesigned to fit inside of the ESA's inner shell.
- **Redesign to fit an even tighter space**: It may be possible to pack the HVPS inside of the ESA's inner shell. This would reduce its impact on the overall size of ACAPS to basically nothing, because the inner shell space is wasted otherwise. It would also remove a lot of the wiring concerns associated with putting the HVPS on top of the ESA. Unfortunately, the inside the ESA is limited to less than half as much board space as the top of the ESA.
- **Faster sweeps and higher voltages** will be discussed below:

How fast can it sweep?

I estimate that this design can achieve 200 Hz sweeping with only minor software changes. The relays can certainly cycle quickly enough. Based on actual turn-on times of well below 1ms on the current design, I believe the relays can maybe even achieve over 500 Hz, but it is likely not practical. Optocouplers might allow faster sweeping rates. It is entirely possible to operate an optocoupler as just an on-off switch.

To push toward 300 Hz or more, a few changes can be made:

1. Minimize extra capacitance on the ESA HV output to reduce sweeping power consumption. At higher switching speeds, capacitance has more of an impact on total power consumption.
2. Operate relays at higher current to achieve faster switching. Supplying higher relay current would simply require changing resistors R15 and R47. This would also increase power consumption, but relay current is a small fraction of the total.
3. It may be possible to recover some of the ESA's capacitively stored energy, instead of the current method which burns it all away as heat in a resistor. The issue with any energy recovery method will be achieving linearity. High-quality resistors are supremely linear in terms of voltage versus

current. A less linear method would cause the output voltage sweep to be less consistently exponential and predictable.

Going past 200 Hz is unlikely to be a priority for most science purposes. This is because short sweeps will not collect as many electrons per sweep, so the data will be noisier. To counteract that, the scientists will have to average the data across multiple sweeps, which defeats the purpose of having so many individual sweeps.

The basic 100 Hz square wave relay cycle is still the foundation of the code, and everything is timed relative to the relays. Eventually, this turned out to be a problem since I had made it very difficult to use any sweep rate other than 100 Hz. The code may need to be rewritten to allow a user-selectable sweep rate. I estimate that will take 20 hours of programming and testing. There is no real difficulty, it is just a lot of code to change. One of the reasons I chose 4 MHz as the main clock speed was because it barely allows a 16-bit timer.

How high can the voltage be pushed?

I really don't know, but I guess 10,000 V is near the practical limit even if the available volume could be 100% filled by electrical components. I have not seriously considered going beyond 6,000 V for the following reasons:

1. Power consumption increases dramatically at higher voltages,
2. Many components stop being available beyond 5 or 6 kV.
3. The circuit board itself would have to be thicker.
4. The number of relays stacked in series would become unreasonable, and I would need to provide additional isolation between relay input and output. Also, the required capacitors become very large.
5. The additional energy range provided would not be as useful for studying auroral electrons.

What enables this design to work?

I believe this design was only possible thanks to recent developments of electronic design. I am thankful to be designing space circuit boards in the year 2020, and not the year 2000.

- **Modern thin-film resistors** are so accurate and stable across temperature that I never had to worry about sorting out the bad resistors or calibrating them.
- **High voltage shock to components** is less of a risk because nearly all electronics have some level of ESD protection built in to their input pins.

- **FRAM memory with built-in error correction** means I am able to trust the microcontroller's memory to work in space without spending a huge effort on software redundancy.
- **ADCs and DACs are so precise and accurate** that they are basically plug and play: I never saw a need to do any calibration.
- **Tiny, cheap voltage references** are so good that they can just be trusted, again not needing calibration.
- **Dealing with oscillators and timers** was not much of a concern, because they are mostly handled within the **microcontroller**.
- **Modern ceramic capacitors have very high density**, which let me achieve decent voltage filtering in limited space.
- **Miniaturized high voltage converters** are only possible thanks to improvements in high-frequency magnetic materials. The converter's transformer is only able to be so small because it can operate at a high switching frequency.

Final Conclusion

In summary, this power supply is a proof of concept. With moderate additional work, I am confident that every one of the requirements can be successfully achieved. I do expect to send this circuit to space within the next few years. It has been designed with that intent from the beginning.

References

[1] National Oceanic and Atmospheric Administration, "Space Weather Prediction Center Homepage," 2023. [Online]. Available: https://www.swpc.noaa.gov/. [Accessed 2023].

[2] S. T. Lai, Fundamentals of Spacecraft Charging, Princeton University Press, 2012.

[3] H. Zhang, *Solar Wind Interactions with Planets and Other Solar System Bodies,* Fairbanks, AK: PDF lecture notes from Space Physics course, 2020.

[4] *ACAPS' partially funded proposal to NASA H-TIDeS (ROSES B-8),* Fairbanks, AK: Unpublished, 2019.

[5] C.W. Carlson, et.al., "An Instrument for Rapidly Measuring Plasma Distribution Functions with High Resolution," *Advances in Space Research,* 1983.

[6] H. O. Funsten and D. J. McComas, "Limited Resource Plasma Analyzers: Miniaturization Concepts," *American Geophysical Union,* vol. Geophysical Monograph 102, 1998.

[7] H. Richards, C. Maldonado, J. Harley, R. Balthazor, and M. McHarg., *Low SWAP Laminated Electrostatic Analyzers for CubeSat Applications,* Poster at the 2020 Small Satellite online conference., 2020.

[8] C. J. Pollock, V. N. Coffey, J. D. England, N. G. Martinez, and T. E. Moore., "Thermal Electron Capped Hemisphere Spectrometer (TECHS) for Ionospheric Studies," *Geophysical Monograph 102,* 1998.

[9] E. Kuffel, High Voltage Engineering Fundamentals, 2nd ed., Newnes, 2000.

[10] Microchip Technology Inc., "Microchip TC1240/TC1240A Positive Doubling Charge Pumps with Shutdown in a SOT-23 Package (DS21516D)," 2001-2012.

[11] XP Power, "A Series DC-HVDC Converters Datasheet," 2022.

[12] Southwest Research Institute, "SwRI SW1502A High Voltage Optocoupler Preliminary Data Sheet," 2016.

[13] Amptek, "High Voltage Optocoupler HV801 Datasheet, Rev. D7," Accessed 2021.

[14] S. J. Bame, D. J. McComas, M. F. Thomsen, et al., "Magnetospheric Plasma Analyzer for Spacecraft with Constrained Resources," *Review of Scientific Instruments,* vol. 64, p. 1026, 1993.

[15] D. T. Young, "Space Plasma Particle Instrumentation and the New Paradigm: Faster, Cheaper, Better," *Geophysical Monograph 102,* 1998.

[16] S. J. Bame, et al., "ISEE-1 and ISEE-2 Fast Plasma Experiment and the ISEE-1 Solar Wind Experiment," *IEEE Transactions on Geoscience Electronics,* Vols. GE-16, no. No. 3, July 1978.

[17] *Personal Communication with Craig Pollock,* Early 2023.

[18] Panasonic Corporation, "AQV258 HE 1 Form A Solid-State Relay Datasheet," 2021.

[19] *Personal Communication with Jeff Rothman, UAF Geophysical Institute Electronics Shop,* 2022.

[20] C. Pollock, et al., "Fast Plasma Investigation for Magnetospheric Multiscale," *Space Science Review,* no. 199, pp. 331-406, 2016.

[21] Hamamatsu Photonics, K.K., "MCP Assembly Technical Information," 2006.

[22] Analog Devices, "ADP170/171 300mA, Low Quiescent Current, CMOS Linear Regulators, Rev.C.," 2014.

[23] Analog Devices, "OP777/OP727/OP747 Precision Micropower Single-Supply Operational Amplifiers, Rev. D," 2011.

[24] J. Schoolcraft, A. Klesh, and T. Werne., "MarCO: Interplanetary Mission Development on a CubeSat Scale.," in *AIAA SpaceOps 2016 Conference*, March 2017.

[25] J. D. Crutchfield, "“MSP430FR5969-SP: FRAM ECC Information” Replies #6 and #8 to the thread," Texas Instruments MSP Low-Power Microcontroller Forum, 6 Aug. 2021. [Online]. Available: https://e2e.ti.com/support/microcontrollers/msp-low-power-microcontrollers-group/msp4.

[26] Texas Instruments, "MSP430FR596x, MSP430FR594x Mixed-Signal Microcontrollers Datasheet. Document Number SLAS704G," Revised Oct. 2018.

[27] Analog Devices, "AD5316R Quad, 10-bit nanoDAC with 2ppm/°C Reference, I2C interface, Rev.D," 2020.

[28] Texas Instruments, "INAx180 Low-and High-Side Voltage Output, Current-Sense Amplifiers Datasheet. Document Number SBOS741H," Revised July 2022.

[29] TDK Electronics AG, "PTC thermistors as limit temperature sensors. SMD, EIA case sizes 0402, 0603, and 0805, Superior series. Series B59421, B59641, B59721," August 2019.

[30] S. Battel, "High Voltage Engineering Techniques for Space Applications," in *Presentation slides from NESC-Battel Engineering.*, [8] , April 2-3, 2012.

[31] J. Brusse, H. Leidecker, and L. Panashchenko., "Metal Whiskers Discussion for the UNOVIS Consortium," June 18, 2008.

Appendix 1: Table of Specifications

Items in red are not fully tested.

Specification	**Min**	**Max**	**Notes**
Input Voltage	3.2 V	3.6 V	Between 2.5 V and 3.2 V may function but will not supply the highest voltages. Above 3.6 V will cause excessive heat dissipation.
Input Current - Disabled	1 mA	5 mA	When Overtemperature, Overcurrent, or Watchdog protections are active.
Input Current - Idle	3 mA	6 mA	When no high voltage is being generated.
Input Current - Without Doubler	50 mA	200 mA	
Input Current - With Doubler Activated	70 mA	200 mA	
Overcurrent Protection Trip Level Minimum	180 mA	250 mA	Time to trip depends on how extreme the current is.
Debug/reprogram interface voltage	3.3 V nominal		For use with a Texas Instruments LaunchPad development board, or an MSP-FET.
Operating Pressure	0	10^{-4} Torr	DO NOT OPERATE IN OPEN AIR- HUMAN SAFETY CONCERN! Has been partially tested at 1000V in air: NOT RECOMMENDED.
Detector Output Voltage	1200 V	2400 V	Below 1200 V gives inconsistent performance. Between 2400 V and 2800 V may work for a short time but is not recommended.
Detector Overvoltage Protection Level	2800 V	3000 V	This is voltage limiting, not latching or resetting. No hysteresis.
Detector Voltage Ripple associated with sweeps	0.5%	2%?	Ripple severity depends on capacitance and whether or not doubler is active. Higher ESA capacitance will make this worse.

Detector Voltage Ripple (Not synched with sweeps)	< 1%		Sweep-to-sweep variation on a 10 kHz bandwidth. This is uncertain, difficult to measure accurately because of noise on the measurement probe.
Sweep Output Voltage	1000 V	4800 V	The maximum is 2x the detector output.
Sweep Timing Accuracy		< ±300 ppm	Across time and temperature. Not exhaustively measured. This mostly depends on crystal oscillator.
Detector Load Resistance	60 MΩ	Infinite?	Has not been tuned beyond a 100 MΩ load.
Detector Current Measurement Range	0 µA	100 µA	24.9 kΩ shunt resistor, up to 2.5 V. Shunt resistor can be changed.
Spectrometer Load Resistance	Infinite		Designed and tested for capacitive loads only. Lower resistance diminishes peak output voltage and increases power consumption.
Sweep Load Capacitance	0 pF	50 pF?	Designed for up to 50 pF. Tested mostly at 15 pF.
Parasitic Capacitance of sweeping output	7 pF		Derived from RC time constant.
Survival Temperature- Powered Off	-60 °C	120 °C	Below -60 °C has not been tested. Up to 120 °C baking has been done.
Temperature- Survival when powered on	-60 °C	90 °C	Below -60 °C has not been tested. The microcontroller has operated past 90°C, but the high voltage will not work.
Temperature- To remain within specifications	-40 °C	75 °C	Software will cut off high voltage beyond this range
Over-temperature protection trip level	75 °C	90 °C	Overheat protection will cut power to the board.
Mass of Circuit Board	< 30 g		Depends on how much epoxy and how many MCP filtering capacitors.
Mass of Shielding	< 7 g		Including screws
Vibration Tolerance			Not yet tested
Total Height	< 10 mm		Board when enclosed by shield. The shield is 8 mm tall.

Appendix 2: Entire Microcontroller Code

If you copy this code and use it as main.c in a project in Code Composer Studio with a MSP430FR5967, it should work.

```
//   ACAPS High Voltage Power Supply Main Code

// #TODO: Adjust divider ratio for ESA voltage. Perhaps do some smoothing.

#include <msp430.h>
#include <msp430fr5967.h>       //specific to this chip.
#include <stdint.h>             //integer types

// VERY IMPORTANT VARIABLES
int16_t   targetVoltage = 2400;  //above 3276 will cause errors. But then again you're not supposed to go that high
anyway!
//also under about 1800V will struggle
uint8_t  HighVoltageEnable = 1;
uint8_t  VoltageDoublerEnable = 1;
uint16_t thresholdVoltage = 2300;  //for comparator output


//Declare functions
void HVPS_kick(void);          //kick watchdog
void HVPS_setup(void);
void HVPS_portsZero(void);
void HVPS_MCPvoltSample(void);
void HVPS_ESAvoltSample(void);
void HVPS_MCPcurrSample(void);
void HVPS_DACsetChA(void);
void HVPS_DACsetChB(void);
void HVPS_DACsetChC(void);

//Variables
uint8_t MainTimerFlag = 0x00;
uint8_t DACFlag = 0x00;
uint8_t ADCFlag = 0x00;
uint8_t timerFlag = 0x00;

uint16_t a = 0x00;

uint8_t nextADC = 1;

uint16_t tempSample = 0;
int16_t X = 0; //temporarily used in handling ADC values.
int16_t Y = 0;
int16_t XY = 0;

int16_t X2 = 0; //temporarily used in handling ADC values.
int16_t Y2 = 0;
int16_t XY2 = 0;

int16_t X3 = 0;
int16_t Y3 = 0;
```

```
int16_t XY3 = 0;

int16_t MSPtemp = -1;
uint16_t curr1Sample = 0;
uint16_t Xc = 0;
uint16_t Yc = 0;
uint16_t curr1 = 0x0123;
uint16_t CURR3V = 0; // measured in mA

uint16_t ESAvoltSample = 0; //raw sample 12bit
uint16_t ESAvolt = 0; //converted to volts
uint8_t ESAvoltSampleFlag = 0;
uint16_t MCPvoltSample = 0; // raw sample 12bit
uint16_t MCPvolt = 0; //converted to volts
uint8_t MCPvoltSampleFlag = 0;
uint8_t MCPcurrSampleFlag = 0;
uint16_t MCPcurrSample = 0; //raw sample 12bit
uint8_t MCPcurr = 0;

uint16_t sweepCount = 0;//16bit count of sweeps since startup (counts up to 65535)
uint16_t sweepCountHighWord = 0;//overflow for sweepCount (counts up to 4 billion)


int16_t timeAdj = 0x00;//this will be plus or minus a few cycles added to the sweep duration
uint16_t c = 0;
uint16_t Mcycles = 40140;//will be adjusted based on crystal. theoretically 40,000 exactly at 4MHz for 100Hz
relays
uint16_t refCycles = 32764; //number of cycles of ACLK to make 1 second.
int16_t XtimeAdj = 0;//X,Y,XY used for smoothing time adjustment.
int16_t YtimeAdj = 0;
int16_t XYtimeAdj= 0;

//For the DAC and MCP Divider stuff
uint8_t DAC_Enable = 1;

uint16_t DACval_1   = 11;         //10bits
   uint8_t  DACval_1H = 0;         //upper half
   uint8_t  DACval_1L = 0x00;      //lower half (kinda insignificant, only last 2 bits) Gotta have it though
uint16_t DACval_2   = 12;       //10bits
   uint8_t  DACval_2H = 0;
   uint8_t  DACval_2L = 0x00;
   uint16_t DACval_2_TARGET = 0;
int16_t  targetShift = 0;       //Measured in volts. adjusts the voltage target for DACchannel 2. Shouldn't be needed.
uint16_t DACval_3   = 13;
   uint8_t DACval_3H = 0;
   uint8_t DACval_3L = 0x00;

int16_t MCPvoltError = 0;       //measured in volts. Positive means voltage is too high.
uint8_t startup = 1;            //gets set to 0 when startup is complete.

uint8_t i2cstep   = 0;          // for the i2c state machine
uint8_t DACch = 0;              //Selected channel of DAC (1=A, 2=B)
int16_t ctrlTARGET = 100;       //DAC value for control voltage (DAC channel 1). Will be adjusted later, this is just
a starting value
```

```
uint8_t HVON = 0;               //This is 1 only when the HV converter is powered. this is automatically changed, DO NOT manually set it to 1. Set HighVoltageEnable instead.
uint8_t DOUBON = 0;             //This is 1 only when the doubler is activated. this is automatically changed, DO NOT manually set it to 1. Set VoltageDoublerEnable instead.
uint8_t OverTemp = 0;           //1 if temperature is too hot OR too cold
uint16_t HVONsweeps = 0;        //counts up every sweep while high voltage is active
uint16_t HVONsweepsHighWord = 0; //used when HVON overflows

volatile unsigned int i;        // volatile to prevent optimization-used for busy waiting

void main(void) {

  HVPS_kick();
  HVPS_setup();
  HVPS_kick();

// MAIN LOOP BEGINS HERE!!!
  for(;;) {//this repeats forever, only way out should be power cycling the board.
      __bis_SR_register(GIE);//Re-enable interrupts

      LPM0;//it should normally just sit here until an interrupt happens. Within the interrupt, it will leave LPM0


      __no_operation();
      __no_operation();
      __no_operation();

//
//    IDLE POINT   :   code only advances past here when an interrupt happens. Inside the interrupt, it leaves low power mode

//doubler relays turn on
      if (MainTimerFlag & BIT1){
        MainTimerFlag &= ~BIT1;//first clear the flag
        if ((DOUBON == 1)&&(VoltageDoublerEnable == 1)){
          P2OUT |= BIT4;// doubler relays on
        }
      }
//timer flag 2 (just after the sweep has begun)
      if (MainTimerFlag & BIT2){
        MainTimerFlag &= ~BIT2;


      //Optional: voltage ramp-up
//      if ((sweepCount % 600) == 599)
//      {
//        if ((targetVoltage < 2500) && (startup == 0))
//        {
//          targetVoltage = targetVoltage + 100;
//        }
//      }

      HVPS_DACsetChB();//placeholder location for this
      }
//ESA relays turn on
      if (MainTimerFlag & BIT3){
```

```
            MainTimerFlag &= ~BIT3;
            if (HVON == 1)
            {
               P1OUT |= BIT3;//turn on ESA relay
            }
P3OUT |= BIT4;//debug
        HVPS_DACsetChC();//Set threshold for ESA voltage comparator
        }

//Sweep Start- ESA relays turn off
        if (MainTimerFlag & BIT0){//The start of each new cycle (100Hz nominally)
            MainTimerFlag = 0x00; //clear ALL main timer flags

            //MainTimerFlag &= ~BIT0;//clear flag
            P1OUT &= ~BIT3;  //turn off ESA relays
            P2OUT &= ~BIT4;//doubler relays turn off
P3OUT &= ~BIT4;//debug
            HVPS_kick();//kick the watchdog
//count sweeps
            sweepCount = sweepCount+1;

            if ((sweepCount % 100) == 0)
{//adjust sweep timer based on crystal to maintain accuracy- changes division of DCO (fast clock) based on LFXT
(slow clock). Important because fast clock is not very accurate otherwise
               c = TA0R;//TA0 is running on the slow crystal-ACLK
               if (c != 0)
               {//make sure the long-term timing has begun

                  timeAdj = -1*((int16_t)(c - refCycles));

                  if (timeAdj>5){
                     timeAdj = 5;
                  }
                  else if (timeAdj<-4){
                     timeAdj = -4;//add a little asymmetry
                  }

                  XtimeAdj =  4*timeAdj;
                  YtimeAdj =  7*XYtimeAdj;
                  XYtimeAdj = (XtimeAdj + YtimeAdj)/8;

                  Mcycles = Mcycles + XYtimeAdj;

                  TA0R = 0;
                  TB0CCR0 = Mcycles;
               }
               else //meaning the timer is at 0 (just started)
               {
                  TA0CTL = TASSEL_1 | MC__UP  ;//start timerA so it will be ready to check next time.
               }

               //adjusting relay on-time because they change a lot with temperature
               TB0CCR3 = 38200 + (MSPtemp*4) - (targetVoltage/10);//ESA relay on time
               TB0CCR1 = 38200 + (MSPtemp*4) - (targetVoltage/10);//doubler relay on time
               //TB0CCR1 = 38000 + (MSPtemp*7);//doubler relay on time
            }
```

```
        if (sweepCount==0xFFFF){
           sweepCount = 1;
           sweepCountHighWord = sweepCountHighWord + 1;
        }

        HVPS_ESAvoltSample();//move on to the next thing to sample


   //Extreme temperature shutdown
        if ( (MSPtemp > 80) || (MSPtemp < -45) )
        {
           HVON = 0;
           OverTemp = 1;
        }
        else if ((MSPtemp < 75) && (MSPtemp > -40))
        {
           OverTemp = 0;
        }

        if ((OverTemp == 0) && (HVON==0) && ((sweepCountHighWord>0)||(sweepCount>200)))
{//high voltage decide to enable
           HVON = 1;
           //DACval_1 = ctrlTARGET;
        }

        if ((HVON == 1) && (HighVoltageEnable == 1))
{// HIGH VOLTAGE ON- decide voltage
           P4OUT |= BIT7;// Pull the enable pin of the high voltage regulator high. Note you also have to bring the
control pin high for it to turn on.

           if (  ((sweepCount % 3) == 0)  &&  (  (HVONsweeps>500) || (HVONsweepsHighWord>1)  )  )
           { //number above governs rate of increase
              MCPvoltError = MCPvolt-targetVoltage;

              if (MCPvoltError > -1000) //if within such and such V, we're done with startup
              {
                 startup = 0;
              }

              DACval_2_TARGET = ((targetVoltage + targetShift)*16)/120 ;

              if (DACval_2_TARGET > 400)
              {
                 DACval_2_TARGET = 400; //this limit corresponds to the highest possible converter output
voltages.
              }
              else if (DACval_2_TARGET < 10)
              {
                 DACval_2_TARGET = 10;
              }

              if (DACval_2 < DACval_2_TARGET)//sets a cap on how high it will ramp to.
              {//theoretically above should be target*10/293, but 297 works better.
                 DACval_2 = DACval_2 + 1; //10-bit number
                 DACval_2H = DACval_2>>2;
```

```
                DACval_2L = DACval_2<<6;
            } else if (DACval_2 > DACval_2_TARGET)//ramping down in case the target has decreased.
            {
                DACval_2 = DACval_2 - 2; //10-bit number
                DACval_2H = DACval_2>>2;
                DACval_2L = DACval_2<<6;
            }
        }

//Control Voltage set
        if (sweepCount % 5 == 0)//despite this, a new command gets sent to the DAC every single sweep. This
just decides
        {
        //Decide control voltage DAC value
            ctrlTARGET = 270 + ((targetVoltage-2000)/10) + (110*VoltageDoublerEnable) - (5*(MSPtemp)/2) ;

        //Limit the control value to reasonable
            if (ctrlTARGET > 1000)
            {
                ctrlTARGET = 1000;
            }
            else if (ctrlTARGET < 320)
            {
                ctrlTARGET = 320;//Converter generally needs about 1 diode drop minimum to turn on.
            }

            if ((DACval_1 < ctrlTARGET)) //ramping up
            {//want control around 0.7V, which is 287/1024
                DACval_1 = DACval_1 + 1 ; //control voltage increase
                DACval_1H = DACval_1>>2;
                DACval_1L = DACval_1<<6;
            }
            else if ((DACval_1 > ctrlTARGET))//ramping down-should only happen if temperature has increased
(need less volts when hot)
            {
                DACval_1 = DACval_1 - 3 ; //control voltage decrease
                DACval_1H = DACval_1>>2;
                DACval_1L = DACval_1<<6;
            }
        }
//count HV sweeps
        HVONsweeps = HVONsweeps+1;
        if (HVONsweeps == 0xFFFF){//allows counting past 65536
            HVONsweeps = 1;
            HVONsweepsHighWord = HVONsweepsHighWord + 1;
        }

//Set comparator threshold
        DACval_3 = ((thresholdVoltage * 2)/13);
        DACval_3H = DACval_3>>2;
        DACval_3L = DACval_3<<6;

    }//end of HVON=1 if statement

    if (HighVoltageEnable == 0)
    {
```

```
            HVON = 0;
        }
        if (HVON == 0)
{//KILL HIGH VOLTAGE
            DACval_1  = 0;
            DACval_1H = 0;
            DACval_1L = 0;
            DACval_2  = 0;
            DACval_2H = 0;
            DACval_2L = 0;
            P4OUT &= ~BIT7;
        }
//ACTIVATE DOUBLER
        if (((HVONsweeps>400)||(HVONsweepsHighWord>0)) && (VoltageDoublerEnable == 1) && (DOUBON
== 0)){
            //compensate for the slight drop of MCP voltage?? Or is that a mirage?
            DOUBON = 1; //turn on doubler after giving time to settle down. Good to keep it turned off at first?
        }
        if ((DOUBON == 1) && (VoltageDoublerEnable == 0))
        {
            DOUBON = 0; //turn off the doubler if its not allowed.
        }
    }
    //end of things to do at start of sweep.

//ADC sampling START
    if (MainTimerFlag & BIT4){
        MainTimerFlag &= ~BIT4;
        //begin with temperature sensor
        if (((ADC12CTL1 & ADC12BUSY) == 0)  &&  (nextADC==1)){
            ADC12CTL0  &= ~ADC12ENC;
            ADC12CTL3  = ADC12TCMAP | ADC12CSTARTADD_1;//select temp sensor
            ADC12CTL0  = ADC12ON | ADC12ENC | ADC12SHT1_10 | ADC12SHT0_10;
            ADC12CTL0 |= ADC12SC;
        }
    }

//ADC sampling CONTINUE
    if (nextADC == 2){

        if ((tempSample > 400) && (tempSample < 2000)){
            // zero deg.C offset is about 1165
            X = tempSample; //signed 16bit goes from -32k to +32k
            Y = ((int16_t)(15*XY));
            XY = (int16_t)(X + Y)/16;
            MSPtemp = ((XY-1175)*64)/210;
            //MSPtemp = ((XY-1290)*64)/280;
        }

        if ((ADC12CTL1 & ADC12BUSY) == 0){ //NOT SURE ABOUT THIS BEING IN A IF()...
            ADC12CTL0  &= ~ADC12ENC;
            ADC12CTL3  = ADC12TCMAP | ADC12CSTARTADD_2;//select current
            ADC12CTL0  = ADC12ON | ADC12ENC | ADC12SHT1_8 | ADC12SHT0_8;
            ADC12CTL0 |= ADC12SC;
        }
    }
```

```
    if (nextADC == 3){
      if ((curr1Sample > 0) && (curr1Sample < 4000)){
        Xc = ((curr1Sample-30));
        Yc = (15*curr1);
        curr1 = (uint16_t)(Xc + Yc)/16;
        CURR3V = curr1*16/100;//adjusted to match this R2 board. Should theoretically be *16/131
      }
      nextADC = 1;
      HVPS_MCPvoltSample();

    }
    if (ESAvoltSampleFlag == 1){
      ESAvoltSampleFlag = 0;
//#TODO-fix ratio
      ESAvolt = ((ESAvoltSample * 15 )/39)*5;//voltage divider of 56.5M/20k
    }
    if (MCPvoltSampleFlag == 1){
      MCPvoltSampleFlag = 0;
      X3 = MCPvoltSample;
      Y3 = (15*XY3);
      XY3 = (X3 + Y3)/16;
      MCPvolt = ((( (uint16_t)(XY3 * 29)) / 16)); // for a 3000:1 divider, based on 2.5V reference.
(29/16)=1.8125
      HVPS_MCPcurrSample();

    }
    if (MCPcurrSampleFlag == 1){
      MCPcurrSampleFlag = 0;
      X2  = MCPcurrSample;
      Y2  = (15*XY2);
      XY2 = (X2 + Y2)/16;
      MCPcurr = ((XY2)/41);//measured in uA

    }
    //put additional sensors here


//USCI for I2C
//I2C Start talking to DAC.
    if (MainTimerFlag & BIT5){
      MainTimerFlag &= ~BIT5;
      HVPS_DACsetChA();
    }

    if (MainTimerFlag & BIT6){
      MainTimerFlag &= ~BIT6;
      //?????????????
    }

//USCI for UART
    //other stuff to deal with interrupts

    //clock fault?
    __no_operation();
    CSCTL0_H = CSKEY >> 8; //Unlocks clock system
    CSCTL5 &= ~LFXTOFFG;
```

```
      SFRIFG1 &= ~OFIFG;
   }
}
//END OF MAIN LOOP

//*****************************************************************************************
*****************************************************
//*****************************************************************************************
*****************************************************
//********************************************* INTERRUPTS
***********************************************************************
//*****************************************************************************************
*****************************************************
//*****************************************************************************************
*****************************************************

#pragma vector = TIMER0_B0_VECTOR
__interrupt void TB0CCR0_ISR(void)
{
   MainTimerFlag |= BIT0;
   LPM0_EXIT;
}

#pragma vector = TIMER0_B1_VECTOR
__interrupt void TB0CCR1_7_ISR(void)
{

   __no_operation();
   a = TB0IV;//TB0IV gets cleared when you read it, so you have to save a copy
   if (a == 0x02){
      MainTimerFlag |= BIT1;
   }
   else if (a == 0x04){
      MainTimerFlag |= BIT2;
   }
   else if (a == 0x06){
      MainTimerFlag |= BIT3;
   }
   else if (a == 0x08){
      MainTimerFlag |= BIT4;
   }
   else if (a == 0x0A){
      MainTimerFlag |= BIT5;
   }
   else if (a == 0x0C){
      MainTimerFlag |= BIT6;
   }
   LPM0_EXIT;
}

#pragma vector = ADC12_VECTOR
__interrupt void ADC12_ISR(void)
{
   if (ADC12IFGR0 & BIT1) //internal temperature sensor
   {
      tempSample = ADC12MEM1;
```

```
        nextADC = 2;
    }
    else if (ADC12IFGR0 & BIT2) //current amp
    {
        curr1Sample = ADC12MEM2;
        nextADC = 3;
    }
    else if (ADC12IFGR0 & BIT3) // MCP voltage
    {
        MCPvoltSample = ADC12MEM3;
        MCPvoltSampleFlag = 1;
    }
    else if (ADC12IFGR0 & BIT4) //ESA volt
    {
        ESAvoltSample = ADC12MEM4;
        ESAvoltSampleFlag = 1;
    }
    else if (ADC12IFGR0 & BIT5) //MCP curr
    {
        MCPcurrSample = ADC12MEM5;
        MCPcurrSampleFlag = 1;
    }
    LPM0_EXIT;
}

//i2c TX INTERRUPT - big fat state machine inside interrupt
#pragma vector = USCI_B0_VECTOR
__interrupt void USCI_B0_ISR(void)
{
    __no_operation();
    if (UCB0IFG & UCTXIFG0)
    {
        if(DACch == 1){
            if (i2cstep == 0){
                UCB0TXBUF = 0x31; // send DAC the command to update ch.A (this is the converter's control pin)
                i2cstep = 1;
            }
            else if (i2cstep == 1){
                UCB0TXBUF = DACval_1H; // (70/255)*2.5V = about 0.7V (just divide by 100, its close enough)
                i2cstep = 2;
            }
            else if (i2cstep == 2){
                UCB0TXBUF = DACval_1L;
                i2cstep = 3;
            }
            else if (i2cstep == 3){
                UCB0CTLW0 |= UCTXSTP;
                i2cstep = 10;
            }
        }
        else if(DACch == 2){
            if (i2cstep == 10){
                UCB0TXBUF = 0x32; // send DAC the command to update ch.B (this goes to the OP747 opamp for
controlling the converter's voltage regulator)
                i2cstep = 11;
            }
```

```
            else if (i2cstep == 11){
                UCB0TXBUF = DACval_2H;
                i2cstep = 12;
            }
            else if (i2cstep == 12){
                UCB0TXBUF = DACval_2L;
                i2cstep = 13;
            }
            else if (i2cstep == 13){
                UCB0CTLW0 |= UCTXSTP;
                i2cstep = 20;
            }
        }

        else if(DACch == 3){
            if (i2cstep == 20){
                UCB0TXBUF = 0x34; // send DAC the command to update ch.C for ESA comparator
                i2cstep = 21;
            }
            else if (i2cstep == 21){
                UCB0TXBUF = DACval_3H;
                i2cstep = 22;
            }
            else if (i2cstep == 22){
                UCB0TXBUF = DACval_3L;
                i2cstep = 23;
            }
            else if (i2cstep == 23){
                UCB0CTLW0 |= UCTXSTP;
                i2cstep = 99;
            }
        }
    }
    UCB0IFG = 0;
    LPM0_EXIT;
}

//*******************************************************************************************
*****************
//*******************************************************************************************
*****************
//********************************************    FUNCTIONS
//*******************************************************************************************
*****************
//*******************************************************************************************
*****************

void HVPS_setup(void) {
    //Watchdogg Setup
        WDTCTL = WDTPW | WDTHOLD;              // Stop watchdog timer
        //dont forget to timeout any while loops
        REFCTL0 |= REFVSEL_0 + REFON;          // Enable internal 1.2V reference, used for temp sensor (also
maybe used for checking 2.5V ref?)

//Ports Setup
```

```
	PM5CTL0 &= ~LOCKLPM5;			// Disable the GPIO power-on default high-impedance mode to
activate previously configured port settings
    HVPS_portsZero(); //makes sure they all start zeroed so I can safely use |= operators

	P1DIR |= BIT3; //ESA relays pin
	P2DIR |= BIT4; //Doubler pin
	P4DIR |= BIT1; //This is the watchdog kick pin- toggles whenever the watchdog gets kicked.
    P3DIR |= BIT4; //troubleshooting, no flight purpose. Also this pin can directly output SMCLK(may be handy?)

	P4DIR |= BIT7; //P4.7 goes to enable pin for the converter regulator. Converter can only output high voltage if
this goes high.
	//P4OUT |= BIT7; //turns on converter regulator

	PJDIR  |= BIT4 | BIT5; //Pins for the 32kHz crystal
	PJSEL0 |= BIT4 | BIT5; //Pins for the 32kHz crystal

	//I2C Using P1.6 SDA, P1.7 SCL
	P1SEL1 |=   BIT6 | BIT7   ;//This sets the UCB0 to use these pins.  Turns out setting P1DIR is unnecessary.


//ADC12 Setup -note previously was ADC10
	//select which pins do what
	ADC12CTL0 = ADC12SHT0_2 | ADC12SHT1_2 | ADC12ON; //Sets number of ADC12CLK cycles to
sample for. 8 means 256 cycles, 5 means 64 (see user guide)

	//Dont do ADC12ENC until you're done changing settings

	ADC12CTL1 = ADC12SSEL_0 | ADC12SHP;//selects MODOSC (about 5MHz)
	ADC12CTL2 = ADC12RES_2;//selects 12bit mode
	ADC12CTL3 = ADC12TCMAP | ADC12CSTARTADD_1;//sets temp sensor to be Channel 30 (0x1E)

	//This ADC12 has 32 independent conversion and control buffers. But do I need to use more than just one?
	//For all of them: ADC12VRSEL_4 sets positive reference as external, negative reference as AVSS
	ADC12MCTL1 = ADC12VRSEL_4 | ADC12INCH_30;//Internal Thermometer
	ADC12MCTL2 = ADC12VRSEL_4 | ADC12INCH_8; //Current amp. output (5V/A?)
	ADC12MCTL3 = ADC12VRSEL_4 | ADC12INCH_0; //MCP voltage divided
	ADC12MCTL4 = ADC12VRSEL_4 | ADC12INCH_12; //ESA voltage buffered from the OP747
	ADC12MCTL5 = ADC12VRSEL_4 | ADC12INCH_2; //MCP return current-across 24.9kohm.
	//ADC12IER0 = ADC12IE0; //enables interrupts for memory location 0
	ADC12IER0 = 0xFFFF;//enables interrupts for all the ADC channels (this is risky, fix later!)
	ADC12IER1 = 0xFFFF;

//Clocks Setup
	FRCTL0 = FRCTLPW | NWAITS_1; // FRAM wait states because clock might be 2fast4u (>8MHz)

	CSCTL0_H = CSKEY >> 8; //Unlocks clock system
	//CSCTL1   = DCORSEL | DCOFSEL_4; //16MHz?
	CSCTL1 = DCOFSEL0 | DCOFSEL1;     //set DCO to 4MHz
	CSCTL2 = SELA__LFXTCLK | SELS__DCOCLK | SELM__DCOCLK;
	CSCTL3 = DIVA__1 | DIVS__1 | DIVM__1;     // Set all dividers
	CSCTL4 &= ~LFXTOFF;
	SFRIFG1 &= ~OFIFG;//Clear oscillator fault
	//CSCTL0_H = 0; //lock clock system? (recommended by TI's code)


//UCB0 for I2C setup   //Using P1.6 SDA, P1.7 SCL
```

```
UCB0CTLW0 |= UCSWRST;
    UCB0CTLW0 |= UCMST | UCMODE_3 | UCSYNC | UCSSEL__SMCLK;
    UCB0BRW = 40;
    //UCB0CTLW1 = UCASTP_2;  //automatic STOP (???)
    //UCB0TBCNT = 0x01; //bytes of data to send (??)
    UCB0I2CSA = 0x0F; //slave address      // Set slave address-DAC should be 0x0F since A1 and A0 pins are
tied high. First 5 bits of DAC address are always 00011 (remember i2c addresses are only 7 bits.
    UCB0CTLW0 &= ~UCSWRST;              // Clear SW reset now that we're done changing settings

    UCB0IE |= UCTXIE0;                  // this enables the TX interrupt on UCB0

//Timers Setup
  //TimerB for ESA relays
    //Unlike on first revision, relays only switch through software. No direct connection between timer and relay
    //Timer B has 7 capture/compares, so I can do everything together
    TB0CCR0   = 40000;//TA0CCR0 is the period of a whole switching cycle (around 100Hz normally) When this
gets triggered, all relays turn off.
    TB0CCR1   = 36500;//doubler on
    TB0CCR2   = 1000;//doubler off (Dont make this too small, or it will get stuck) (alternatively just turn off the
doubler with the ESA)
    TB0CCR3   = 36500;//ESA relay on
    TB0CCR4   = 4000;//Begin ADC sampling sequence
    TB0CCR5   = 10000;//Begin talking to DAC
    TB0CCR6   = 65530;//adjusting the clock

    TB0CTL    = MC__UP | TBSSEL_2  | TBCLR;  //    //Selects UP MODE and tells it to use SMCLK
(DCO),clears timer. Don't need TBIE, that enables an overall interrupt for the counter ticking over, but can use other
flags instead

    TB0CCTL0 |= CCIE;               //CCIE means Capture Compare Interrupt Enable
    TB0CCTL1 |= CCIE;               //
    TB0CCTL2 |= CCIE;               //
    TB0CCTL3 |= CCIE;               //
    TB0CCTL4 |= CCIE;
    TB0CCTL5 |= CCIE;
    //TB0CCTL6 |= CCIE;

 //Timer A0 is used for compensating timerB over long-term
    TA0CTL    =  MC__UP | TASSEL_1  | TACLR;
    TA0CCR0   = 0xFFFF;

  __delay_cycles(10000);
  //Now that it's had some time to stabilize, re-clear the oscillator fault. (not sure if necessary)
  CSCTL0_H = CSKEY >> 8; //Unlocks clock system
  CSCTL5 &= ~LFXTOFFG;
  SFRIFG1 &= ~OFIFG;
}

void HVPS_MCPvoltSample(void){
  ADC12CTL0  &= ~ADC12ENC;
  ADC12CTL3  = ADC12TCMAP | ADC12CSTARTADD_3;//select addr3
  ADC12CTL0  = ADC12ON | ADC12ENC | ADC12SHT1_10 | ADC12SHT0_10; // define number of clock
cycles to sample and hold
  ADC12CTL0 |= ADC12SC;
}
```

```
void HVPS_ESAvoltSample(void){
   ADC12CTL0  &= ~ADC12ENC;
   ADC12CTL3  = ADC12TCMAP | ADC12CSTARTADD_4;//select addr4
   ADC12CTL0  = ADC12ON | ADC12ENC | ADC12SHT1_10 | ADC12SHT0_10; // define number of clock
cycles to sample and hold
   ADC12CTL0 |= ADC12SC;
}

void HVPS_MCPcurrSample(void){
   ADC12CTL0  &= ~ADC12ENC;
   ADC12CTL3  = ADC12TCMAP | ADC12CSTARTADD_5;
   ADC12CTL0  = ADC12ON | ADC12ENC | ADC12SHT1_10 | ADC12SHT0_10; // define number of clock
cycles to sample and hold
   ADC12CTL0 |= ADC12SC;
}

void HVPS_DACsetChA(void){
   __no_operation(); //this just provides a place for the debugger to stop.
   if (DAC_Enable == 1)
   {
      if((UCB0STATW & UCBBUSY) == 0){ //make sure its not still doing a prior transaction
         DACch = 1;
         i2cstep = 0;
         UCB0CTLW0 |= UCTXSTT + UCTR; //transmit mode, start condition
      }
   }
}

void HVPS_DACsetChB(void){
   __no_operation();
   if (DAC_Enable == 1)
   {
      if((UCB0STATW & UCBBUSY) == 0){ //make sure its not still doing a prior transaction
         DACch = 2;
         i2cstep = 10;
         UCB0CTLW0 |= UCTXSTT + UCTR; //transmit mode, start condition
      }
   }
}

void HVPS_DACsetChC(void){
   __no_operation();
   if (DAC_Enable == 1)
   {
      if((UCB0STATW & UCBBUSY) == 0){ //make sure its not still doing a prior transaction
         DACch = 3;
         i2cstep = 20;
         UCB0CTLW0 |= UCTXSTT + UCTR; //transmit mode, start condition
      }
   }
}


void HVPS_portsZero(void){ //makes sure they all start zeroed so I can safely use |= operators
   P1DIR = 0x00;
   P2DIR = 0x00;
```

```
   P3DIR = 0x00;
   P4DIR = 0x00;
   PJDIR = 0x00;

   P1OUT = 0x00;
   P2OUT = 0x00;
   P3OUT = 0x00;
   P4OUT = 0x00;
   PJOUT = 0x00; //Deal with this later: unused pins of PJ should be set as outputs to avoid any issues with floating
pins triggering weird JTAG things.

   P1REN = 0x00;
   P2REN = 0x00;
   P3REN = 0x00;
   P4REN = 0x00;
   PJREN = 0x00;

   P1SEL0 = 0x00;
   P1SEL1 = 0x00;
   P2SEL0 = 0x00;
   P2SEL1 = 0x00;
   P3SEL0 = 0x00;
   P3SEL1 = 0x00;
   P4SEL0 = 0x00;
   P4SEL1 = 0x00;
   PJSEL0 = 0x00;
   PJSEL1 = 0x00;
}

void HVPS_kick(void) //restarts the watchdog counter at zero so it doesn't reset.
{
   //WDTCTL = WDTPW | WDTHOLD;
   WDTCTL = WDTPW | WDTCNTCL | WDTSSEL_1;// | WDTIS0; // (Watchdog Timer Count Clear)-this resets
the counter. Don't worry, this bit doesn't have to be manually cleared.
   //WDTIS0 sets it to divide by 512, which works out to about 16ms. THAT'S PRETTY QUICK-MIGHT HAVE
TO CHANGE IT IF THE SWEEP RATE IS SLOWER THAN 100Hz!!!
   P4OUT ^= BIT1;
}


/*******************************************************************************************
*******************************************************
 *                                TO DO
 *Write USCI code for UART
 *Make sure unused PJ pins are not floating.
 *
*/
```