

A Detector for Counting Single Photons at 795 nm

Nicholas D. Farrow

A thesis submitted to the faculty of the University of Colorado
in partial fulfillment of the requirements for the degree of
Bachelor of Arts

Department of Physics

Committee:
Heather Lewandowski (advisor)
John Cumalat
Nikolaus Correll

Defense Date:
April 4th 2011

Contents

- 1 Introduction**
 - 1.1 Introduction to photon detection
 - 1.2 Application of single photon detection: Cold Molecules
- 2 Project Proposal: A Single Photon Detector for 795 nm photons**
 - 2.1 Considerations in detector design
- 3 Detector technology: Photomultiplier Tubes and Avalanche Photodiodes**
 - 3.1 Photomultiplier Tubes
 - 3.2 Avalanche Photodiodes
- 4 Quenching Circuit**
 - 4.1 Comparison of Passive and Active Quenching
 - 4.2 Description of the Circuit
 - 4.3 Photon Detection and Monostable Behavior of the Circuit
 - 4.4 Simulating an Avalanche Photodiode
 - 4.5 Circuit detection efficiency
- 5 Thermoelectric Cooling**
 - 5.1 Introduction to Thermoelectric Coolers
 - 5.2 Designing a Heat Sink for the Thermoelectric Coolers
 - 5.3 Accurately controlling the temperature
 - 5.4 Calibrating the Temperature
 - 5.5 Tuning the Temperature Controller
- 6 Physical Construction of the Module**
 - 6.1 Fabrication of the Counting Module
- 7 Characterizing the Behavior of the APD**
 - 7.1 A Source of Light
 - 7.2 Measuring the Breakdown Voltage of the APD
 - 7.3 APD Sensitivity, Relation to Overvoltage
 - 7.4 APD Overvoltage Relation to Temperature
 - 7.5 APD Dark Count Rate Relation to Temperature
 - 7.6 APD Dark Count Rate Relation to Overvoltage

- 7.6 APD Dark Count Rate Relation to Overvoltage

- 8 Counting the Pulses**
 - 8.1 Connecting the Counting Module to the Computer
 - 8.2 Analyzing the Screenshots with LabVIEW

- 9 Final Data and Conclusion**
 - 9.1 The Experimental Procedure
 - 9.2 Dark Counting Rate of the Detector, Effect of Temperature
 - 9.3 Sensitivity of the Detector, Effect of Overvoltage
 - 9.4 Sensitivity of the Detector, Afterpulse probability
 - 9.5 Future Improvements

Chapter 1

Introduction

1.1 Introduction to photon detection

Single photon detection is becoming a widely used technology in many fields of science. This is especially true in physics as experiments progressively approach quantum scales. There are numerous experiments that require detectors with single photon detection capability. Lifetime fluorescence measurements and single molecule detection rely heavily on single photon detection technology. Single photon detection is also used in medical imaging in devices such as PET and CT scanners. Single photon detection has become common in the fields of astrophysics, particle physics and condensed matter physics. An emerging field of physics, quantum cryptography, could not exist without the ability to detect single photons.

Single photon detection is an inherently difficult process. There are a range of photon detection technologies available, but each have their own limitations. Most single photon detectors employ either photomultiplier tubes or avalanche photodiodes as their primary detector. Photon detectors may be used for a variety of applications, including counting photons, measuring rates of photon emission, and measuring time correlations of photon emission.

1.2 Application of single photon detection: Cold Molecules

Our group is interested in experimenting with, and understanding the quantum behavior of cold polar molecules. The group has recently used cold molecules to study the interactions of atomic rubidium and molecular NH_3 (ammonia). Next, we would like to focus our attention to another exciting polar molecule, the free radical NH . Understanding the dynamics of NH is important to many fields of science. NH can be detected in interstellar gas, and may aid astrophysicists in understanding the process of star formation. NH is also a reaction intermediate in the combustion of organic compounds found in fossil fuels.

Observing and measuring ultracold chemical interactions (collisions) at the quantum level allows scientists the opportunity to investigate these molecules at the quantum scale. By applying cold molecule methods to NH, we hope to elucidate some of the more interesting behaviors and properties of this interesting, and important molecule.

Because our lab group is interested in studying cold molecules, a photon detector would be a valuable resource for many applications in our lab. Although complete detectors are commercially available, there are relatively few to choose from. They are also rather expensive, with prices starting at a few thousand dollars. Of those available, dark rates around a few hundred Hz are the norm. To achieve these lower dark rates, the photodetectors utilize small active-area detectors. Dark rate may be further lowered by cooling the detector. Commercially available detectors that employ cooling typically cool their detector to around 30 °C below ambient temperature. Cold molecule trap densities are low enough that the directional photon flux out of them may be a small fraction of the dark rate of these detectors. To optimize the detection of these photons in cold molecule experiments, it is desirable to have a large active area, and a detector with dark rates on the order of, or below that of the signal. No detectors are available that meet both of these needs. It was therefore decided to construct our own single photon detection module that could be built to custom suit the needs of the lab group.

The detector is intended to be used with cold atom/cold molecule experiments. In these experiments, a small cloud of molecules is held in an electrostatic trap. The photon detector will be aimed at the cloud of trapped molecules and will be used to count photons emitted from the cloud. One of the proposed experiments will be to study “resonant quenching” in collisions between NH (in the singlet delta state denoted $^1\Delta$) and ground state ($5S_{1/2}$) rubidium. This collision is near resonant; the NH can transfer its energy to a Rb atom, itself decaying to the NH ground state, meanwhile pumping the Rb up to an excited state (almost exclusively the excited $5P_{1/2}$ state). This excited Rb state is “electric dipole allowed” to decay back to the ground state of Rb, emitting a photon. Counting these photons will allow the lab group to have a better understanding of what is happening in the cloud of molecules.

It would be ideal to design the detector to be sensitive to only a small range of photon wavelengths centered on 795 nm. If detection is limited to only 795 nm photons, it can be inferred that photon counting correlates to Rb decay in the trap. Measuring the photon detection rate would then give a measurement of the interaction rate of the atoms and molecules in the trap.

Chapter 2

Project Proposal: Design and Build a Single Photon Detector for 795 nm photons

2.1 Considerations in detector design

Building a new scientific instrument requires great care in planning and precision in construction. The detector should be sensitive to photons within a narrow window of wavelengths, centered at 795 nm photons. An optical band pass filter may be used in conjunction with the detector to eliminate some sources of noise from extraneous light sources. Another important source of noise is the dark count rate of the detector. Dark count rate can be lowered by cooling the diode.

The detector must as sensitive as possible. Sensitivity in optical detectors is expressed as quantum detection efficiency, or simply quantum efficiency (QE). Quantum efficiency is a value representing the photon detection probability - expressed as a percent - that a photon that hits the detector will be detected. There are many factors contributing to the sensitivity of the detector. This will be explained in more detail in another chapter.

There is limited space available on the end of the vacuum chamber (where the molecule clouds are trapped). Many types of equipment are used here, including photomultiplier tubes, TOFMS, REMPI, and MCP¹ detectors to name a few. A primary design goal was to keep the detector compact. Ideally, a small of a case as reasonably possible should be used to keep the unit compact. A compact unit would allow room for other equipment to be used simultaneously in future experiments. Detector space should be conserved, even if that space is currently not occupied.

¹ TOFMS = time of flight mass spectrometer
REMPI = resonantly enhanced multiphoton ionization,
MCP = microchannel plate
TOFMS + REMPI + MCP = ion detector

Finally, the detector needs to be easy to use, and relatively inexpensive to construct. Most commercially available detectors are self contained units that are just plugged in, some even with USB outputs. It is desirable to emulate some of the features of these commercially available detectors. The detector is designed as a self contained module, but does require some accessories for power and active diode cooling. The temperature goal for the project is to cool the device to much lower than commercially available detectors. Precise temperature control should be easy to operate. Cryogenic (liquid nitrogen) cooling is sometimes employed for this purpose, but would significantly increase the cost and complications associated with operating the device. Ultimately, the temperature goal was met by cooling the device using only electronic components, and components that were readily available in the lab.

Goals of the detector

There are four primary design criteria that the detector should meet.

1. The detector should be sensitive enough to measure a10 Hz signal.
2. The detector should be specific in its wavelength detection ability.
3. The detector should maintain a compact size.
4. The detector should be easy to use and to integrate with existing lab equipment.

Chapter 3

Detector technology: Photomultiplier Tubes and Avalanche Photodiodes

3.1 Photomultiplier Tubes

To begin the search for the best detector technology, we first examined our possibilities with photomultiplier tubes (PMTs). PMTs consist of a series of electrodes called dynodes. Incoming photons strike the first dynode, called the photocathode, ejecting an electron via the photoelectric effect. That electron is accelerated along an electric field gradient until it strikes a secondary dynode causing a shower of electrons to be released. These electrons then cascade down the tube impacting on more dynodes and increasing the electron count as they go. Hamamatsu is probably the market leader when it comes to photon counting modules. They offer a range of photon counting modules spanning the range from UV to IR, including both PMT and diode based modules. However most available products are analog devices; despite being marketed as photon counters, their output format is in the form of amperes per watt of incident light.

Of the few photon counting modules offered that output a digital signal, there is one that's peak sensitivity is near 800 nm, (Hamamatsu H7422P-50) and its dark rate is 125 – 375 Hz with a QE of 12% [7]. The specified project goals hope to exceed this performance. This ruled out using a PMT as the primary detector component.

3.2 Avalanche Photodiodes

Next, the market was examined for avalanche photodiode (APD) based modules. APDs are silicon based photodetectors which are constructed similar to regular diodes in that they are formed from the junction of p-type and n-type semiconductor materials. A reverse-bias voltage is applied to the APD such that the diode is biased above its breakdown potential. An incoming photon is detected by the photoelectric effect, as it destabilizes one of the electrons in the semiconductor material. The free electron is accelerated by the electric field, liberating other electrons in the process, initiating an avalanche of current which can then be measured. APDs have typical QEs around 80%,

much better than PMTs covering the same spectral range. This makes APDs much more appealing since our signal is expected to be on the order of the noise rate. Again, a search was conducted of commercially available APD counting modules, but none were identified with outstanding characteristics or ones that were optimized for our spectral range. The best option available is a line of single photon modules by PerkinElmer (SPCM-AQR-1X). They all offer the same detection efficiency, 55% for 800 nm, and their dark count ranges from 500 Hz to 25 Hz depending on how much you have to spend. Up to date prices are not listed, but are assumed to be upwards of \$5000. [8]




Figure 3.1: A comparison of the quantum efficiencies of a typical avalanche photodiode with a typical photomultiplier tube. Note that detection efficiencies of photomultiplier tubes are sometimes given in the form of a spectral sensitivity. (Source 3.1a: [6], 3.1b: [10])

An APD is a special type of silicon photodiode, but is closely related to regular photodiodes. Photodiodes are made from three principle semiconductor layers sandwiched together. The three layers are a p-type (abundance of holes), an n-type (abundance of electrons), and a lightly doped depleted region² in between. Together these constitute a PIN diode. Avalanche photodiodes are different from regular diodes in that they typically use a much thinner depleted region, and the applied bias voltage is usually much greater than that used with regular diodes. To be used as photon detectors, a *reverse*-bias voltage must be applied across the three layers. When incoming photons impact on electrons in the depleted region, the electrons are freed from the substrate and travel along the electric field gradient within the substrate. As the electron moves with the potential, it gains energy from the electric field imparted by the bias voltage. In APDs, the higher bias voltage in the depleted region causes the freed electron to gain sufficient energy that it knocks other electrons free. This causes the electrons to avalanche. The avalanche can then be sensed as a current through the diode.

² The depleted region is also known as the intrinsic region.




Figure 3.2: Cross section of a typical avalanche photodiode showing the layered semiconductor design. (Source: [11])

Avalanche photodiodes may be used in two separate modes of operation. In typical APD use, the APD is treated just like a regular photodiode: light enters the diode causing current avalanches and the current out of the APD is measured. In this mode, the current gain is specified in the form of A/W. Noise is specified in the form of a dark current. A dark current is a steady current through the diode that is present, even when no light is admitted to the diode. In this configuration, the APD is acting as an analog device. Photon counting rates can be approximated by integrating the measured output current.



Figure 3.3: An image of the APD used in the photon counting module. The PerkinElmer C30902 Silicon APD is available in two package types. This module uses the one on the left. The APD on the right is used with fiber optics.

Alternatively, the APD can be used in a more digital format, in what is called ‘Geiger’ mode operation. In this configuration the APD is biased *above* its breakdown voltage. APDs used in this mode are referred to in the literature as a single photon avalanche diodes (SPADs). There is no difference between an APD and an SPAD except for the bias voltage setting³. ‘Geiger mode’ is a reference to Geiger counters used to detect radiation.

³ There may be slight differences in manufacturing as some APDs may be manufactured with the intention of using them in Geiger mode, but in reality they are interchangeable.

In Geiger mode, incoming photons cause current avalanches in the APD much greater in magnitude than in analog mode. The avalanche pulse, 10^8 electrons, effectively causes the diode to break down. An APD conducting in breakdown will continue to conduct unless the current is shut off. Stopping the current flow through an SPAD after breakdown is called 'quenching' the diode. Quenching is necessary; if left in conduction, the sustained current will eventually damage the APD. Furthermore, when an APD is in breakdown, incoming photons will have no further effect. Quenching the APD will reset the diode so that it is ready to count another photon.

Chapter 4

Quenching Circuit

4.1 Comparison of Passive and Active Quenching

There are two methods of quenching the SPAD, passive and active. In passive quenching, a load resistor is placed in series with the diode. The load resistor acts as a current limiting resistor so that a sustained avalanche current is not allowed. A typical load resistor value ($100\text{ M}\Omega$) is chosen so that current is limited to be below the latch-current of the diode. The latch-current is the current through the diode necessary to sustain breakdown. A series current through both the diode and the resistor, causes the voltage drop to be shared between the two components such that the voltage across the APD falls below its breakdown threshold. When the voltage across the APD drops below breakdown, the avalanche current stops, and the diode is quenched. Voltage across the diode then rises as the voltage across the load resistor falls back to zero. Once the voltage across the diode rises back above breakdown voltage, the APD regains its ability to detect another photon. The time to complete this passive quenching cycle depends on the RC time constant of the load resistor and the inherent capacitance of the APD. This also limits the maximum counting resolution that may be obtained with the passive quenching circuit to usually less than 1 MHz.

The other method of quenching the APD uses an active feedback circuit. This type of circuit is called an active quenching circuit (AQC). Active quenching circuits are employed when a counting resolution greater than that achievable by passive quenching is desired. Our detector uses an active quenching circuit.

The photon detection and avalanche quenching circuit must perform two functions in conjunction with the APD. First, the circuit must act as a discriminator, able to detect the onset of an avalanche pulse. Second, the circuit must be able to provide a secondary current pulse to counter the primary avalanche pulse. The quenching circuit is connected to the APD through only one of its terminals, so the quenching circuit needs to be able to perform both of these operations at the *same point* of the circuit. This type of active quenching circuit is said to be in the coincident terminal configuration (as opposed to the opposite terminal configuration)

The benefit of using a coincident terminal configuration is that the APD has a free terminal that is not directly connected to the quenching circuit. The free terminal is used to apply the bias voltage. Since the bias voltage is not applied by the circuit itself, the circuit can be used generically with *any* APD.

The quenching circuit itself acts as a monostable trigger. A simple way to document detection of a photon event is to count the pulses generated by the quenching of the avalanches, rather than trying to count the avalanches directly. In this way, the circuit is behaving as an amplifier. Every avalanche triggers a quench, and every quench triggers a pulse out of the module. To minimize noise on the circuit, the electronics are housed inside the module case, but the quench pulse needs to be measurable outside of the case. The most practical method to export the signal outside the case is through a BNC connector, followed by a coaxial cable to our pulse counter. Coaxial cables have inherent capacitances, which could interfere with the comparator's ability to successfully quench the APD. To avoid this problem, another pulse amplifier is used between the output of the comparator and the BNC port. Specifically, a transistor gate driver (Zetex ZXGD3004E6) is used that is able to respond to the 10 ns width of the comparator pulse⁴. Experimentally, the transistor driver performed better than expected. It had the effect of sharpening the comparator pulse and nearly doubling its amplitude to 2V.⁵

4.2 Description of the Circuit

The circuit used in this project was adapted with some modifications from a publication [2]. The circuit contains both digital and analog components on the same board. The digital and analog components are combined in an interesting arrangement. The primary sensing component on the board is the AD8611, prominently displayed in the middle of the circuit diagram. The AD8611 is a high speed voltage comparator with *two* outputs: Q and Q-NOT. The comparison inputs are on the analog part of the board, while the logic outputs are on the digital side of the board. On the far left of the circuit diagram are the rest of the digital components (three NOT gates) used in the active quenching and sensing circuit. There are a handful of resistors and diodes and a capacitor in between that constitute the analog part of the circuit. Most of the other parts on the board (the components on the bottom and left side of the schematic) are the voltage regulators and the line driver. These are not particularly relevant to the operation of the quenching circuit. The circuit was designed so that a single +15V line could be used to power all of the components, with the exception of the APD.⁶

⁴ We originally tried to use a line-driver MIC4420, but its response time was too slow to register a pulse

⁵ Going into the project, I wanted TTL logic levels to be employed as the output format, but the timescales involved in generating and measuring these pulses made TTL impractical.

⁶ The bias potential across the APD is provided by a high-voltage Op Amp which is configured to be a simple power supply

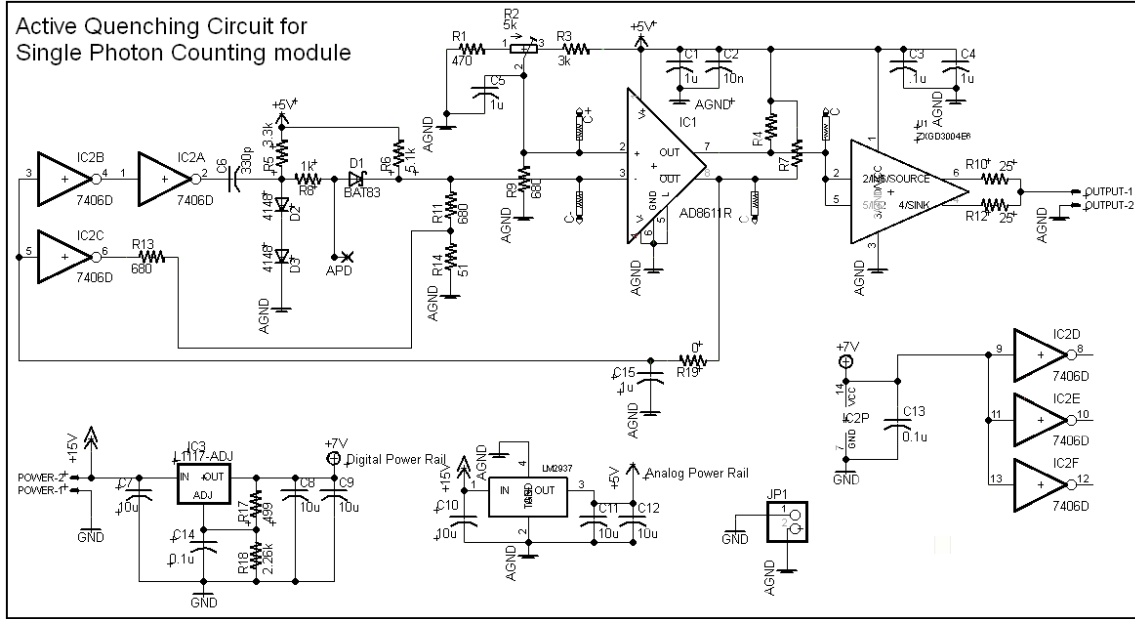


Figure 4.1: Circuit diagram of the electronics used in the active quenching circuit

The circuit detects an avalanche from the APD which triggers the circuit's monostable behavior. The monostable loop spans both the digital and analog components. The circuit spends the majority of its time in the resting state waiting for a photon. In the resting state, a steady current flows from the 5V line, through R5, R8, D1, R11, R14, to ground. This series of components forms a (somewhat complicated) voltage divider that establishes a resting potential of about 0.75V on the inverting input of the comparator. The potential on the non-inverting input of the comparator is established by another voltage divider. This voltage divider is formed by the combination of R3, R2, and R9. R2 is the trimpot that is adjusted to set the threshold of the comparator. The potential on the non-inverting amplifier is adjusted so that it is about 30 mV below the inverting input. The comparator has two outputs, Q and Q-NOT. In the resting case the comparator's output Q, is normally low, while Q-NOT is high (+5V TTL). Q-NOT is fed back into the circuit, while Q is used as the output off of the board. Q-NOT is connected directly to the digital logic chip containing six NOT-gates. There are two NOT-gates (2B & 2A) between Q-NOT and the capacitor C6. This means that the logical value of gate 2A matches Q-NOT and so normally keeps C6 charged (+7V). Capacitor C6 plays a critical role in the quenching of the APD. There is an optional low-pass filter between Q-NOT and the NOT gates, but it is currently permanently shorted.⁷

⁷ This low-pass filter may be used if it is desirable to slow down the loop. Slowing the feedback loop increases the hold-off time, which will increase the chances of a successful quench. This would be implemented if a high afterpulse rate becomes a problem

4.3 Photon Detection and Monostable Behavior of the Circuit

A large reverse-bias voltage is applied to the free terminal of the APD that keeps the APD biased a few volts over its breakdown potential. A *negative* voltage is applied to the APD to obtain the desired bias voltage (typically -185V at room temperature). When a photon is detected by the APD it initiates an avalanche pulse causing APD breakdown. When the APD conducts, it appears to the circuit as a short to a negative potential. The small current that would normally be flowing from R8 through diode D1 is now pulled down through the APD⁸. Since D1 is no longer conducting, the potential at the inverting terminal of the comparator drops. This drop in potential is sensed by the comparator and triggers the comparator to change its output state. Q-NOT now goes to logical low (0V TTL). The low logic level reaches NOT-gates 2B and 2C simultaneously. The logic gates each have a 4 ns propagation delay. The output of 2C goes high (+7V) which propagates through resistors R13 and R11 back to the comparator. This will shut off the comparator pulse. Meanwhile, the output of the third NOT gate, 2A, is forcefully pulled low by 2B going high. When gate 2A is pulled low, a -7V charge is pushed off of capacitor C6 which becomes the quenching pulse for the APD. This pulse quickly lowers the potential across the APD below its breakdown threshold terminating the avalanche.

The circuit will return to its resting state when the comparator registers the signal from logic gate 2C. The propagation delay through the comparator is about 15 ns, causing its output pulse to last for approximately 20 ns. After this time, Q-NOT will return to a logical high state. Again, this reaches the two parallel logic gates simultaneously. The output of 2C simply goes low allowing current to resume through diode D1. The output of gate 2A will go high (+7) to match Q-NOT. The capacitor C6 recharges, and any latent circuit ringing should be dumped through diodes D2 and D3. At the conclusion of this, the circuit has been reset, and is ready to detect another pulse from the APD. Each half of the cycle takes approximately 20 ns to complete. Therefore the reset time of the circuit is 40 ns.

The comparator output Q is used only as the pulse signal off of the board. Before the pulse leaves the board, it passes through a gate-driver. The gate-driver is remarkably fast (1 ns), and is able to source and sink tremendous current (8 A). This component is much more responsive than the comparator; it is mostly behaving as a voltage follower for the comparator. The gate-driver's high current capability should allow it to easily drive the 20 ns pulse through any length coaxial cable that would be needed, without having to worry about the pulse being lost due to stray capacitance in the line.

⁸ The diode D1 is specially chosen. It is a Schottky barrier diode. This diode has important advantages relevant to pulse detection. Barrier diodes have lower forward resistance and lower noise generation than typical diodes.

Parts list:

Part #	Description
AD8611ARZ	two-output high speed voltage comparator
74AC04MTR	14-pin Hex Inverter (SOIC)
LM2937IMP	voltage regulator, dropout-type, 5V
LM1117IMP-ADJ	adjustable voltage regulator (set for 7V)
ZXGD3004E6TA	IC gate driver/MOSFET
BAT83S-TR	Schottky barrier diode
1N4148 SMD	surface mount diode
resistors	surface mount resistors (assorted)
capacitors	surface mount capacitors (assorted)

4.4 Simulating an Avalanche Photodiode

Avalanche photodiodes are relatively expensive components. At a few hundred dollars apiece, they are not a circuit component that can be easily replaced. In order to characterize the behavior of the circuit, we wanted to be able to simulate the photodiode. Initially, a simple function generator was used to provide quick negative pulses to the circuit. The output of the function generator was connected directly to the APD terminal of the board. The function generator was able to supply pulses of $-2V$, 10 ns in duration to the quenching circuit.

This stimulated the comparator to trigger a pulse, but introduced tremendous noise on both of the comparator terminals. The non-inverting input to the comparator must be completely steady in time. Otherwise the comparator threshold will fluctuate, resulting in erratic behavior. It was also noted that the resting potential of the function generator was completely dominating the resting potential of the inverting input. This was not letting the circuit maintain its own passive resting potentials throughout the board.

Next, it was decided to use traditional components to simulate the internal capacitance and resistance of an APD. In its resting state, a Geiger mode APD has trapped charge carriers in stable positions in the depleted region. An incoming photon destabilizes the charge carriers causing the avalanche pulse. In this way, an APD avalanche is similar to a discharging capacitor. To simulate this behavior, the simulated APD pulse from the function generator is applied through the simulated APD circuit, so

that the -2V pulse is transmitted through the capacitor. The resistor isolates the function generator from the circuit, allowing the circuit to maintain its own resting potentials.

This APD simulation gave a slight improvement to the pulse detection ability of the quenching circuit. However, It did not give insight into the quenching circuit's ability to quench the APD. Quick pulses from the function generator are a reliable means of triggering the monostable behavior of the circuit, but these pulses are *programmed* to be short. In the final application of the circuit, the quenching feedback mechanism must be used to terminate the triggering pulse. In order to completely simulate an APD, a more advance circuit is required.

To make use of the feedback mechanism of the quench circuit, the output of the logical NOT gate would be used to terminate the pulse from the function generator. This time, a field effect transistor (FET) was employed as our simulated APD. When the APD breaks down, it appears to the circuit as a short to ground (or lower potential). In this respect, the APD is behaving similarly to a transistor. The simulation circuit used an AND gate coupled to a FET. The function generator is now used to apply a TTL pulse to one terminal of the AND gate. The other terminal of the AND gate is connected to the output of the logic chip on the quenching circuit. In its resting state, the NOT gate of the circuit is outputting a logical 1. When the function generator is pulsed, the AND gate will trigger the FET to conduct. The quench circuit senses this, and a quench is triggered. The quench pulse makes its way to the NOT gate, whereupon the NOT outputs a logical 0. The logical 0 is fed back to the AND gate, terminating the signal to the FET, and so closing the short to ground.

This circuit simulation was very successful. The quenching circuit could reliably terminate the triggering pulse, but circuit sensitivity was still an issue. The successful implementation of a simulated APD gave us the confidence to proceed forward with the project construction. Simulating the APD also allowed us to identify weaknesses in the quenching circuit design. Issues that could be resolved were excess noise on some of the lines, and the inadequacy of our line driver's ability to propagate the signal off of the board⁹. The board was redesigned with more noise reducing capacitors and a more responsive line driver. Also, separate ground plates were introduced to separate the digital components from the analog components. A final test of the quenching ability of the circuit would need to wait until the real APD was used.

4.5 Circuit detection efficiency

The quenching circuit is designed to operate in conjunction with the APD such that a negative high-voltage is applied to the free terminal of the APD. This negative voltage is supplied by a bipolar op-amp capable of supplying +/- 1000 V at +/- 40 mA (KEPCO Bipolar Operational Amplifier BOP 1000M). For the purpose of this experiment, the voltage range was limited from approximately -200 V to 0 V, and current was limited to +/- 2 mA. (A word of caution to anyone using this supply; the voltage limits are ignored when the BOP is switched off. The output voltage rises appreciably

⁹ A MIC 4420 was first used as the line driver

into the positive when powered off, presumably due to discharging of internal capacitors or inductors)

The sensitivity of the circuit to pulses may be adjusted by varying the trimpot which sets the comparator level. If the comparator threshold is set too low, the comparator will miss pulses from the APD (the APD will self quench via passive methods). If the comparator threshold is set too low, the circuit will self-oscillate. As reported elsewhere in the literature [3], an ideal threshold is approximately 30 mV to reliably detect a pulse.

Complications arising from the circuit detection efficiency of actual photons were responsible for a considerable delay in the project timeline. When the actual APD was inserted into the circuit, it was thought that the APD should be protected from passing too much current. Initially, the circuit was using a current limiting resistor in series with the APD. This configuration is used in some types of passive-active hybrid quenching circuits [3]. From tests of the circuit, it was found that the resistance value of this series resistor considerably changes the magnitude of pulses transmitted to the comparator. With a 25 k Ω resistor in series, pulses to the comparator were around 1 V in magnitude. Adjusting the discriminator level gave mixed results. The circuit appeared to be either ultra-sensitive to pulses or non-responsive. Compounding this problem, the series resistor greatly increased hold-off time required for a successful APD quench, which the active quench circuit was not designed for. This resulted in pulses being generated in the APD before the APD returned to its resting bias voltage. These little pulses, that never fully quench, prevent the APD from ever accumulating enough charge to cause a single significant avalanche pulse. Effectively, the APD is paralyzed, and cannot trigger the comparator [3]. Adjusting the discriminator level alone, could not fix the underlying problem.

The sensitivity issue was resolved when, paradoxically, applying light to the diode stopped the pulses. Probing the APD here, confirmed the APD was in fact paralyzed. The error here delayed an accurate measurement of the dark rate by about a month. Removing the current limiting resistor and re-adjusting the discriminator level allowed the detector to reliably emit pulses when exposed to light.

Chapter 5

Thermoelectric Cooling

5.1 Introduction to Thermoelectric Coolers

Thermoelectric coolers (TECs) achieve their heat pumping ability by using the Peltier thermoelectric effect. The TEC consists of two ceramic plates sandwiching an array of p-type and n-type semiconductor materials. The semiconductor materials are doped in such a way that charge carriers transport heat with them as they conduct through the array. The p-n materials used in the TEC have a remarkable property; they are chosen so that heat is pumped along the same gradient in both materials, while the current through them is opposite [9].




Figure 5.1: A cutaway view of a thermoelectric cooler. (Source: [12])

TECs are commercially available in a wide range of sizes and heat pumping capacities. In general, the larger the surface area of the TEC, the more heat it can pump. Heat pumping capacity scales linearly with the surface area of the TEC. TECs have a

maximum heat pumping efficiency inherent to the device. At low temperature differences (a few degrees), the heat pumped by the TEC is approximately linear with the current through the TEC. The TEC creates its own waste heat in the process of pumping heat. The heat generated by the TEC is easily computed by using the formula $P = IV$, where I is the current through the TEC and V is the voltage applied to the TEC. At low power levels, the TEC is primarily pumping external heat from the hot side to the cold side. As the pumping power is increased, the TEC must pump more of its own waste heat, so its pumping efficiency begins to decrease. At I_{max} (the maximum current through the TEC) the TEC is pumping 100% of its own power so its efficiency is 0 and no further cooling occurs. To achieve practical heat pumping levels, the current should be set to between 1/3 to 2/3 of I_{max} . The maximum amount of heat is pumped within this range.

Under normal operating conditions, each TEC will maintain a certain temperature difference across its ceramic plates. TECs may be stacked to achieve greater temperature differences. In a stack of TECs the total temperature difference is a sum of the temperature differences from each individual TEC within the stack. In a stacked configuration however, each TEC must also be able to pump the waste heat generated from the TECs above it. It therefore becomes impractical to make large stacks of TECs. To mitigate this, the TEC stacks generally follow a tiered configuration. The TECs are stacked with smaller ones on top and larger ones on the bottom. This is done so that the larger TECs are able to handle the waste heat from the smaller TECs on the top. TEC stacks are also commercially available and we use one in this project. (Tellurex M2-40-1503-3)

Again, there are options for powering TECs: passive and active. The simplest approach is passive. In passive cooling, a steady DC current is used and the TEC is allowed to reach thermal equilibrium across its ceramic plates. In this case the heat pumped is constant in time as long as the temperature of the ceramic plates does not change. If the temperature of the cold side is able to fluctuate, an active cooling approach may be desired. In this case, a temperature controller is used to monitor the temperature at the cold side, and the current through the TEC is actively adjusted to change the cold side temperature at the set point.

5.2 Designing a Heat Sink for the Thermoelectric Coolers

TECs push heat from one side to the other. Heat is removed from the cold side and deposited on the hot side. There must be some sort of heat dissipating mechanism on the hot side, otherwise the TEC would overheat. Commercially available single photon modules that are TEC cooled typically have a large heat sink fixed to one side, along with an equally large fan to move air across the fins.




Figure 5.1: Hamamatsu H7422P photon detector module, showing its compact size and heat sink fins with cooling fan.

Taking a cue from these modules, we initially experimented with fin-style heat sink cooling as well. A DC current was passed through the TEC to cool a block of aluminum. The temperatures dropped for awhile, but soon unexpectedly began to rise. Increasing the current to the TEC only worsened the problem. The case temperature was not measured but was very hot to the touch. Fans blowing air on the fins did not help. The temperature difference that we were aiming for required more heat to be pumped than could be dissipated by the fins. It was decided that heat sink cooling would not be a feasible option for this project.




Figure 5.2: Heat sink configuration that was considered and tested. The diagram shows the 3-tier TEC stack attached to a large aluminum block centered in the case. A fin-style heat sink was applied to the bottom of the case.

Next, water cooling the block was investigated. Water cooling uses a running stream of water passed through a metal block which acts as the heat sink. Water flows through channels drilled through the block of metal, absorbing the heat and carrying it out of the block. Flow rate through the tubes becomes an important consideration. Ideally as much water is moved through the tubes as possible, as quickly as possible. The volume of water that flows through a tube increases with the square of the tube radius. Therefore a few large through-holes are better than a bunch of small-diameter through holes. On the other hand, a large tube lowers the water pressure in the tube, and the water flows more slowly. If the water flow is too slow, the water will begin to saturate with heat before it passes completely through the heat sink. Therefore a careful balance is needed. The tubes are most efficient when placed as close to the heat load as possible. In our initial design, the tubes were placed directly behind the face of the heat sink. This was found experimentally to be an effective means of cooling the diode. Using water as a heat sink also had the unexpected benefit of cooling the entire case of our module to the water temperature. This is beneficial because it automatically sets our baseline temperature to around 60 °F, as opposed to room temperature. The water comes from

outside the building – underground – and so is automatically at underground temperature. This also means that the water temperature, and hence TEC hot-side temperature will fluctuate with the seasons. In the summer, the water temperature was measured to be relatively steady with day to day levels between 16.6 °C and 17.0 °C (62 °F). In the winter, the water temperature was once measured at 11.9 °C (54 °F).



Figure 5.3: Water cooling the module. On the left, the fin-style heat sink was removed and holes were bored through the aluminum block for water cooling. The water passed directly through the aluminum block and the case. The water kept the block cool, but the desired temperature (-70 °C) could not be achieved with a single TEC stack. On the right is the heat sink configuration that is used in the final construction. In this arrangement, a two stage cooling approach is used. Heat is pumped from the 3-tier stack into the large aluminum block centered in the case. Another group of TECs is used to pump heat from this block into a water cooling block attached to the bottom of the case.

5.3 Accurately controlling the temperature

In JILA, temperature controllers are readily available as they are commonly used with laser diode cooling¹⁰. In practice, diode lasers must be kept at a constant temperature somewhere around 30 °C. The actual temperature will vary from diode to diode, but the important part is that they must be kept at a very *constant* temperature. It should be noted that at 30 °C, the temperature controller will be doing a roughly equivalent amounts of heating and cooling to keep the diode temperature constant.

Temperature controllers operate via a negative-feedback circuit. A thermistor is employed to monitor the temperature at the cold side. A thermistor is a resistor whose resistance is proportional to its temperature. The datasheet of the thermistor gives the thermistor resistance as a function of its temperature, and this will be an exponential function. The thermistor that we are using has a precise resistance of 2 kΩ at 25 °C. (NTC Thermistors DC95F202W).

The thermistor was epoxied in placed near the diode using thermally conductive epoxy. The temperature controller measures the resistance of the thermistor, and tries to keep it at a set value using a negative feedback circuit. The temperature controller is ‘programmed’ by specifying a set resistance it will attempt to match. A comparator in the temperature controller compares the resistance of a set resistor to the resistance of a thermistor using a standard comparator configuration.

¹⁰ Diode lasers are also mounted on TECs

The temperature controller is loaded with an array of set resistors connected to a knob. The tuning knob is used as a coarse adjustment to set the set-point resistance. The temperature controllers used in JILA are designed to operate within a fairly narrow window centered around 30 °C. To modify a temperature controller to operate over a wider range of temperatures, it was necessary to change out all of the set resistors. The replacement resistors were chosen to span a range of resistances corresponding to set points between 0 and -80 °C.

The temperature controllers also have a fine-adjustment knob. This knob does not contribute to the set-resistance. Rather it acts as proportional gain control for the output of the resistor comparison. The details of this gain control are not particularly insightful, but what is important is the scaling factor of its output. Normally constructed JILA temperature controllers are able exercise gain control to within +/- 10% of the output. To achieve the wide tuning range desired between set points, it was necessary to change out this part of the temperature controller circuit as well. Two resistors and the 10-turn trimpot were replaced so that the gain range was 50% to 125%.

5.4 Calibrating the Temperature

Inside the module, the APD is mounted in an aluminum block that also contains the temperature sensing thermistor. The thermistor was placed as close to the APD as possible. It is hoped that the two share an identical temperature. The thermistor used in the module is a 2 k Ω thermistor. The datasheet asserts that this thermistor is suitable for use over the range -80 °C to 150 °C, and guarantees its accuracy to +/- 2% over the range 0 °C to 80 °C. However, outside this range it does not specify its accuracy. Since we will be using the thermistor to operate our temperature controller to temperatures down to -70 °C, it was necessary to determine the behavior of the thermistor below of its specified tolerance range. A thermocouple temperature sensor (Fluke 54II) with K-type probe was used for the temperature measurements.

The probe was inserted into the front end of the diode mount (aluminum block) without the diode present. The probe was placed close to the thermistor by inserting it into one of the grooves that are used to pass wires to the diode. The module was set-up for water cooling. The current through the TEC was manually varied to achieve different temperatures to the diode mount. Styrofoam insulation was extensively used around the diode mount to help the module achieve its coldest temperatures. Temperatures were allowed to stabilize at steady state heat flow before recording their value. (It takes 30 – 45 minutes for temperatures to asymptotically approach their stable value when dynamic heating is not employed; this part of the project took a long time to complete). When the temperature did not change after a few minutes, I recoded both the temperature and the thermistor resistance. Approximately 60 thermistor values were recorded over a temperature range spanning 20 °C to -68 °C. The data was fed into a graphing program (Origin) and a best fit line through the data was computed using the functional form of an

exponential decay,

$$R = R_0 + A e^{-\left(\frac{T-T_0}{B}\right)}$$

where T is the independent variable. A and B are constants to be determined. R_0 and T_0 are the resistance and temperature at known values. ($R_0 = 2 \text{ k}\Omega$, $T_0 = 25 \text{ }^\circ\text{C}$) In practice, these were also determined by the best-fit line.

In hindsight, perhaps this is not the best functional form to use for the best-fit line. A more appropriate choice would be to use the Steinhart-Hart equation [13]. This equation, with three free parameters, is used to model the resistance of a semiconductor at different temperatures. A simplified variant of the Steinhart-Hart equation, dropping the least-significant parameter and combining the remaining two, is called the B-parameter equation:

$$R = R_0 e^{-B\left(\frac{1}{T} - \frac{1}{T_0}\right)}$$

where T is the independent variable. B is the constant to be determined. R_0 and T_0 are the resistance and temperature at known values.

A new best-fit line was later re-computed using this functional form. Both equations pass through most of the temperature points of interest on the plot.

Thermocouple vs. Thermistor




Figure 5.4: A plot of two best-fit lines through the data points. The blue line, which does not go through most of the data points, but does pass through the data point at approximately $20 \text{ }^\circ\text{C}$, is a plot of the equation provided by the datasheet. Note that it is inaccurate at temperatures below $-10 \text{ }^\circ\text{C}$. The B-parameter equation is also in better

agreement with the higher temperatures. The exponential decay diverges at higher temperatures, but provides the best fit at the lower temperatures.

5.5 Tuning the Temperature Controller

The temperature controllers designed by the JILA electronics shop use a PID feedback scheme and come with a Ziegler-Nichols tuning instruction sheet. Ziegler-Nichols is one of the available tuning algorithms used to optimize the feedback loop of PID circuits. Ziegler-Nichols tuning is the algorithm used with the diode lasers, and is reliably employed for this purpose. It was natural to assume that Ziegler-Nichols tuning would work for the APD cooling.

The Ziegler-Nichols tuning and calibration algorithm begins by deactivating the integral and derivative part of PID. Then the natural period of the circuit is measured. The circuit will oscillate (approximately sinusoidally) between two current output values that are centered around the set point. From discussions with Terry Brown¹¹, I learned that this free-running oscillator arises from what is happening on a macroscopic scale at the level of the TEC/heat reservoir/thermistor level¹².

It is worthwhile to explain the cyclical behavior here. To begin the cycle, the TEC is conducting current so that it is cooling. It takes a finite amount of time from when the TEC begins to conduct, to when the thermistor will measure a difference in temperature. A rise in thermistor resistance will trigger the TEC to stop cooling. While the TEC is off, the temperature of the diode mount will approach equilibrium, whereupon the thermistor continues to become colder. This will trigger the temperature controller turn on the TEC, but this time for heating. The pattern repeats itself for a heating phase to complete the cycle. The cooling/heating oscillation repeats indefinitely keeping the temperature around the set point¹³. The important part here is the time required for heat to propagate to the thermistor. The time required for heat conduction constitutes one-quarter of the period of the cycle. Therefore the natural period of the cooling/heating cycle scales proportionally with the size of the system. In the case of diode lasers operating around room temperature, the time constant is around a few seconds. When the Ziegler-Nichols algorithm was applied to this module, it was found that the time constant was around 55 to 60 seconds.

Ziegler-Nichols uses an aggressive feedback mechanism that relies on having a system with a short time constant. When tuned correctly using Ziegler-Nichols, the heating and cooling is always $\frac{1}{4}$ cycle out of phase with the natural period of the system. This allows the temperature to rapidly adapt to a new set-point, while the aggressive cycling keeps the system at the set point.

¹¹ JILA electronics shop

¹² The heat reservoir is the diode mount

¹³ The goal of Ziegler-Nichols tuning is to tighten this oscillation frequency to be as close as possible to the set point.

In the case of the photon module, the natural heating/cooling period of the relatively large system had the effect of actually destabilizing the temperature of the module. Ziegler-Nichols tuning of TECs makes use of both their cooling and heating capabilities. With the module, the desired temperature range (-70 °C) is so far below ambient, that it is much more difficult to cool the diode than it is to heat the diode. In fact, any amount of heat applied will be too much heat. If any amount of heat is applied, it must diffuse through the diode mount before it is measured, and by that time the total amount of applied heat becomes a significant load that needs to be pumped back out during the cooling phase.

At this point, it was decided that Ziegler-Nichols tuning would not work for the design; it is too aggressive. Using the Ziegler-Nichols tuning algorithm without the heating, was attempted, but again the temperature cycled while the TEC cooling was turned on and off. The time constant of the diode mount is just too large to use a PID tuning algorithm. A benefit of having a large time constant is that this corresponds to a large thermal reservoir. A large reservoir is beneficial if passive cooling is employed since small thermal fluctuations from the diode are easily damped out while the temperature difference across the TEC is maintained at a steady state level. The large thermal reservoir acts to keep the temperature gradient constant. From this point on, it was decided to use the temperature controller in P (proportional) mode only. Experimentally, this was found to be a very stable mode of operation, especially at the lower temperatures (anything below -30 °C).

After completing the modifications to the temperature controller, and calibrating the thermistor, it was possible to compute a plot that shows all of achievable set points, along with their computed temperatures. The behavior of the modified the modified temperature controller was tested and it was in close agreement to the calculated temperatures over a range of values.




Figure 5.5: This graph is used as a means of converting a desired temperature setting to a resistor setting on the temperature controller. The coarse adjust knob is chosen to the desired set-point resistance. Next, the 10-turn fine adjust knob is set to the desired temperature. The lines shown here were calculated using the best-fit line computed during the thermistor calibration. A fine-adjust setting of 5-turns equates to a gain of 1, corresponding to the exact value represented on the calibration curve. The accuracy of this plot was determined by tracing the fine-adjust knob at a coarse setting of 130 kΩ. It was verified that the measured temperature matched the prediction to within 1 degree.

Chapter 6

Physical Construction of the Module

6.1 Fabrication of the Counting Module

The module is made from individual components that were acquired from many different sources. Many of the components are readily available from online retailers. Some of the components were machined from a block of aluminum. All electrical components were purchased and soldered together in the lab. All of the metal components were precision machined in the JILA staff shop. The gasket lid case was purchased, however its inside and outside surfaces were milled smooth to improve temperature conduction. All non-circuit related parts were designed to keep the heat conduction through the module as efficient as possible. All of the faces of the metal components that interface with the TECs were polished to a mirror finish. This is important to keep the components in good thermal contact. Furthermore, all of the heat conduction interfaces were coated with a thin layer of thermal compound (zinc oxide).

To minimize heat conduction to the cold side, single strand wire-wrap wires were used as electrical conductors from between the APD and the electronics board. Large grain Styrofoam was cut and placed (rather packed) all around the APD mounting block. Styrofoam has nearly the same thermal conductivity as air, but it is a better insulator than air because it prevents convection air currents that would otherwise develop inside the housing. The aluminum center block is surrounded by a thin sheet of Styrofoam on the three sides that would otherwise be exposed to the wall of the housing. The fourth side (base) is in direct thermal contact with TECs which drain heat to the water-cooling block. With this configuration, the aluminum housing is completely separated into two compartments so air from the cold side cannot flow back to the warmer side which contains the electronics, and vice-versa.




Figure 6.1 An exploded view diagram showing the APD mounting configuration. A brief description of the parts:

1. The aluminum block which sits in the middle of the case of the module. This block acts as an large intermediate thermal reservoir. It rests on top of four high performance TECs (TE Technology HP-127-1.0-0.8) which pull heat out of this block and pump it into the water cooling heat-sink through the bottom of the case. In normal operation, this block may be cooled down to -30°C .
2. The 3-tier TEC stack is mounted onto the face of the aluminum block. (Tellurex M2-40-1503-3)
3. A smooth-finish aluminum plate is mounted to the top of the TEC stack. This plate is approximately 1 mm thick, and forms the foundation of the APD mount. Thin grooves are cut into this face to permit the passing of wires to the APD.
4. The main part of the APD mount. This component is about 1 cm thick. The large hole here is designed so that the APD mount may accommodate large diameter APDs. This block is held to the foundation by 4 screws.
5. The adapter is used to mount small size APDs in the larger sized hole.
6. The C30902 avalanche photodiode used in the project.
7. An antireflective coated window is used to seal the aperture opening. The window has been glued into part 8.
8. A block of aluminum to hold the AR coated window. Gluing the window to this block, instead of directly to the case allows the easy removal of the window for maintenance. (The window needs to be removed to gain access to the screws in the APD mount)
9. The thermistor has been permanently connected to the APD mount using thermally conductive epoxy.
10. A large diameter APD may also be used in the module to increase the light collecting ability of the module. We plan to try this size of APD in future experiments.

Not shown: Styrofoam is used around the mount to prevent convection air currents. Parts # 2-4 are held to the block (#1) by an utlem plate.

Chapter 7

Characterizing the Behavior of the APD

7.1 A Source of Light

A consistent attenuated light source was constructed. The light source consists of a generic red LED mounted to a 1 inch disk of aluminum. A BNC connector was used so that the LED could be connected to a 15 V lab power supply. A 2 kohm resistor was used in series with the LED to limit the current. The LED had a dim glow when lit by 15 V. A capacitor was placed in parallel with the LED to act as a low-pass filter for the LED current source. This was done so that the LED would not fluctuate in brightness in the case of noise on the power line. A 1 inch optics tube was obtained to house the light source. A sheet of Teflon was obtained from the shop, and was punched into 1 inch disks. The Teflon disks were stacked inside the 1 inch tube to the desired attenuation thickness. Lastly the LED mount was used to cap the optics tube with the Teflon disks sealed inside. Upon powering the source in a darkly lit room, I could not see light penetrating through the Teflon sheets. The light source intensity was not measured, but by construction, is assumed to be constant light in light output.




Figure 7.1: Construction of the uniform light source

7.2 Measuring the Breakdown Voltage of the APD

To measure the breakdown voltage as function of temperature, the following procedure was used. The temperature controller was set to maximum cooling until a steady temperature developed. The constant light source is connected to the device and is powered so that the LED illuminates the diode. The output from the device is monitored on the oscilloscope for pulses. The supplied bias voltage is set so that pulses are seen, but are not continuous.

In operation, the temperature controller passes a current through the thermistor to measure the diode temperature. The temperature controller converts the resistance into an error (proportional to how much the thermistor and set resistance disagree). The resistance of the thermistor measured with a multimeter while it is in the feedback loop of the temperature controller. Therefore, it is impossible to manually measure the thermistor resistance while the temperature controller is operational. However, the temperature controller outputs its error (deviation from set-point) through a BNC port which *can* be monitored on the oscilloscope. Typically the error will overshoot its steady state level at most twice while the temperature stabilizes¹⁴. When the temperature was deemed stable (and hence the diode was as cold as possible), the thermistor was quickly removed from the temperature controller and connected to a digital multi-meter.

Once the thermistor is disconnected from the temperature controller, the diode immediately begins to heat due to passive thermal conduction, primarily through the TEC stack. As the diode begins to heat, its breakdown voltage will increase. The heating of the diode may be seen on the oscilloscope as pulses become more sparse. When pulses from the detector cease, the temperature of the thermistor, as measured with the multimeter is recorded. The bias voltage supplied to the APD is then increased by a few volts whereupon pulses reappear on the oscilloscope. The APD continues its passive heating and this process is repeated taking as many data points as possible. Breakdown voltages were recorded for temperatures between -50 °C and 5 °C.¹⁵

¹⁴ If the thermistor is not connected to the temperature controller, the temperature controller will not cool the diode. When the thermistor is disconnected from the temperature controller, the temperature controller will measure an infinite resistance, which corresponds to the thermistor being too cold. In such a case with a bipolar temperature controller, the temperature controller would immediately supply maximal-positive current to heat the thermistor, but in this experiment, no positive current is supplied to the temperature controller, so this is disallowed.

¹⁵ The APD heated much more quickly at the colder temperatures, so data below -30 °C was difficult to obtain in this manner. If I were to do this again, I would not use the temperature controller; rather I would manually set the current through the TECs using a power supply, and measure the thermistor temperature more precisely as the bias voltage was adjusted so that pulses just begin to appear on the oscilloscope.

Breakdown Voltage vs. Thermistor Resistance

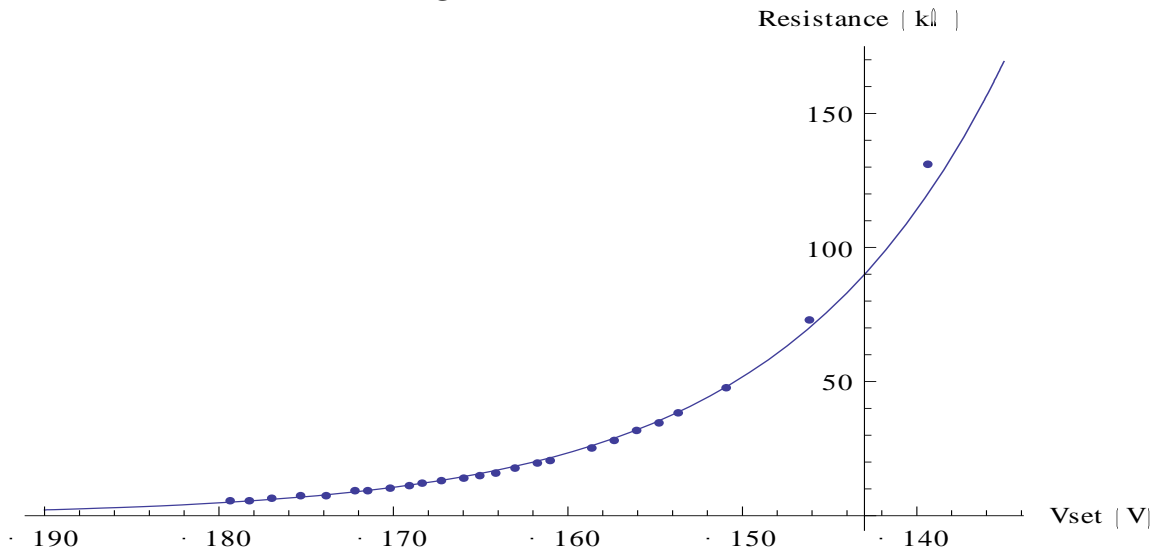


Figure 7.2: A plot of the best-fit line through the data points. When the thermistor resistance is converted into a temperature, this function becomes a straight line. This is shown in Figure 6.4.

7.3 APD Sensitivity, Relation to Overvoltage

In Geiger mode, the photon detection probability steadily rises with the applied bias voltage above breakdown. The voltage above the breakdown potential does not have a consistent terminology in the literature. One convention that may be used is to call this the overvoltage. At precisely breakdown voltage (overvoltage = 0V), the detection probability is 0%, and should rise to a detection probability of 50% with an overvoltage of around 16 V. The relationship is entirely non-linear, and there is not a given function that can describe the relationship. There is some ambiguity in the datasheet as to the performance of the diode in the range of overvoltage from 0 V to 6 V. Presumably, each diode may behave differently. The datasheet gives the approximate behavior only for 22 °C. For the purposes of characterizing the behavior of the detector, this relationship needs to be found experimentally. Literature publications indicate that the sensitivity as a function of overvoltage does not change with temperature [1]. To clarify, the breakdown voltage *does* depend on temperature, and consequently the applied bias voltage (the overvoltage) *will* be a function of temperature.




Figure 7.3: Geiger mode photon detection probability vs. overvoltage at 22 °C for the C30902 Silicon APD. (Source: [6])

7.4 APD Overvoltage Relation to Temperature

The breakdown voltage of the diode decreases linearly with temperature. This linear relationship results from the thermal expansion (contraction) of depleted region of the diode. As the diode cools, the depleted region contracts, linearly increasing the magnitude of the electric field within depleted region. Exploiting this linear relationship is important to characterizing the diode over a range of temperatures. The temperature coefficient of the diode was measured to be +0.744 V/°C.

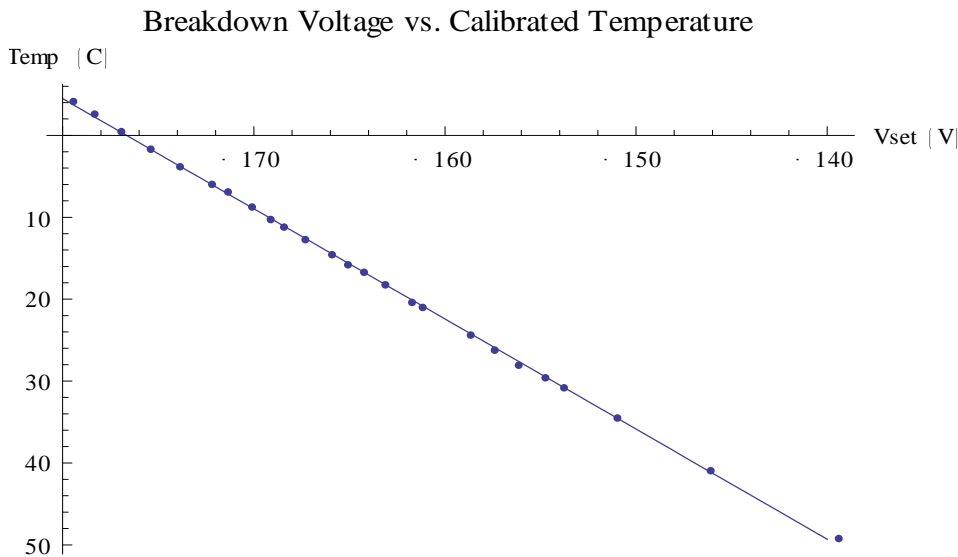


Figure 7.4: A plot showing the APD breakdown voltage linear relationship to temperature.

7.5 APD Dark Count Rate Relation to Temperature

Most (if not all) of the dark counts from the module arise from a thermally generated avalanche in the APD. Fundamentally, temperature is manifest as a vibration (thermal oscillation) of the crystal structure of the APD. The thermal oscillations inside the APD will sometimes destabilize a charge carrier in the depleted region. The charge carriers normally reside at stable regions introduced in the depleted region by doping in the manufacturing process. If a charge carrier is knocked loose by the thermal oscillation, it will initiate an avalanche pulse. Cooling the APD will decrease the dark counts by means of lowering the thermal oscillations. However, there may be a limit to this. It has been proposed that the charge carriers move more slowly the colder that they are. This may have the unintended consequence of trapping charge carriers in unstable position following an APD quench [3]. If the charge carriers are trapped in unstable positions, then they are more likely to be knocked loose. This may significantly increase the afterpulsing rate of the APD. It is unknown if this is a real phenomenon or if we will operate the module at temperatures where this becomes a problem.




Figure 7.5: Typical dark count vs. temperature at 5% photon detection efficiency (830 nm) for C30902 Silicon APD in Geiger mode. Note that a 5% detection efficiency equates to a bias overvoltage setting of 2V. (Source: [6])

7.6 APD Dark Count Rate Relation to Overvoltage

According to the datasheet for the diode, dark counts are more likely to occur at higher overvoltages [6]. In fact, the dark count rate is supposed to follow the same curve as the photon detection rate as seen in Figure 6.3. However, this statement does not fully take into consideration the dark count rate's relation to temperature. If the dark counts follow the same curve as the light counts, then the curve is probably scaled with temperature. We are relying heavily on the fact that dark counts decrease with temperature, while the sensitivity of the APD to light should remain constant. This relationship will be extensively measured before the detector is put to use in actual experiments.

Chapter 8

Counting the Pulses


8.1 Connecting the Counting Module to the Computer

The photon detector module will output a very short (about 10 ns) pulse with a magnitude of slightly over 1 V for each detection event. The module is connected to an oscilloscope so that the pulses can be monitored. The oscilloscope is connected to the computer network in JILA. In principle, any computer in JILA can access the oscilloscope on the network. A LabVIEW computer program was used to connect to the oscilloscope through the network. Each time the LabVIEW program executes, it requests a single screenshot from the oscilloscope. The oscilloscope is left in free-running mode (untriggered) while the APD module is emitting pulses. In principle, when the LabVIEW program requests a screenshot, it will receive a completely random sampling in time of the APD module.

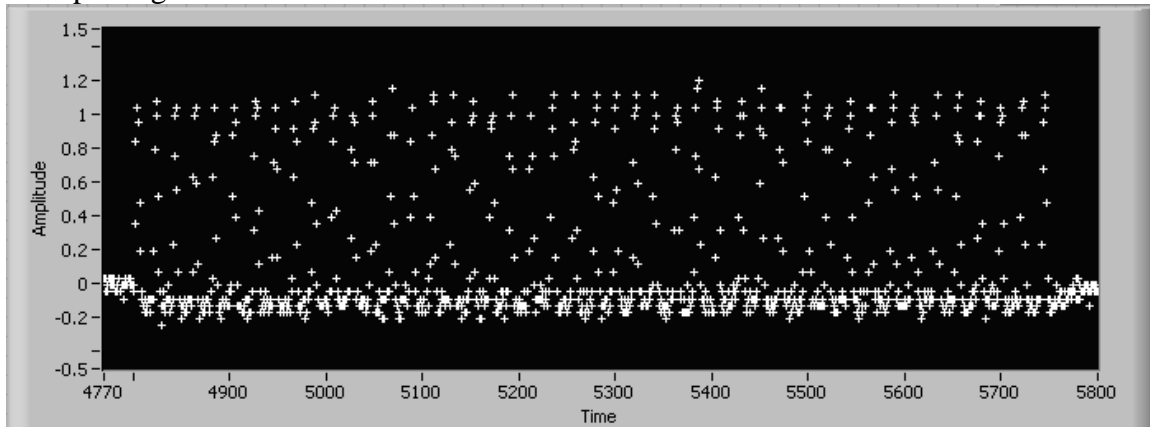
8.2 Analyzing the Screenshots with LabVIEW

A program to count pulses was created for LabVIEW. LabVIEW programs are called VIs (virtual instruments). The photon counting VI was adapted from another VI that was being used to count pulses from a PMT. There were many changes applied to the pre-existing VI to enable more efficient pulse counting from the APD. Features of the supplied VI that were kept included the VI's ability to read a screenshot from an oscilloscope over a network connection, the ability to find and subtract off the DC voltage offset, and the VI's ability to tally pulses as they were identified from the oscilloscope trace. A significant change was made to how the VI identified pulses. Since the electronics inside the detector module react the same way to every detection event, the pulses coming out of the module are extremely uniform in height and width. This allows a simple discriminator threshold to be used to identify the beginning and end of a pulse. This made pulse detection much more efficient and allowed closely spaced (40 ns) pulses to be resolved. Fine resolution control is preferred, since that allows the oscilloscope to send longer time snapshots. Individual pulses can be resolved from screenshots up to 40 microseconds in duration.


Another modification to the original VI was including the ability to read the time-stamp information from the oscilloscope. This allows the VI to associate a time with every pulse event. This is important for several applications of the counting program. Depending on the time-resolution of the screenshots leaving the oscilloscope, the counting circuit can measure time between pulses to 100 picoseconds. This allows the VI to discriminate between pulses originating from a unique detection event or afterpulses initiated from an unsuccessful avalanche quench.



Afterpulsing from a dark count



Zoomed in appearance of the pulse train



Individual pulses can still be resolved from the pulse train

Afterpulses may become a serious problem with APDs, especially at high overvoltages. Afterpulses do not arise from photon generated events nor thermal noise, yet in most applications they are treated like dark counts. As previously discussed afterpulses happen following an incomplete (or unsuccessful) quench of the APD, and arise due to charge carriers trapped in unstable positions in the diode substrate. If an afterpulse is going to happen, it will occur within the first 20 ns after the quench. The quench process of our circuit takes approximately 40 ns to complete. That means that the circuit may afterpulse between 40 to 60 ns following the previous pulse event.

The VI is programmed so that if a pulse begins within 60 ns of the beginning of a previous pulse, it is considered to be an afterpulse, otherwise the pulse is considered a novel pulse. Afterpulses are tallied differently than novel pulses. From tallying both novel pulses and all pulses, the VI provides a calculation of the afterpulse rate. This is significant information since it a high afterpulse probability correlates to too high of an applied bias voltage. Since we are able to reject afterpulses from the counting rate, we can get a much more accurate measure of the actual counting rate and/or dark count rate.

Chapter 9

Final Data and Conclusion

9.1 The Experimental Procedure

The final data acquisition with photon counting module is intended to characterize the counting behavior of the module. An important test was to determine the dark count rate at a variety of temperatures. Theoretically, the sensitivity of the APD to light is only determined by its overvoltage, while the dark count rate is determined by both the overvoltage and the temperature. Thus, the dark count rate should decrease with cooling the APD, while its sensitivity to photons remains constant. This will be examined.

During the final data acquisition, the APD was set to a desired temperature using the temperature controller. Then the bias voltage was set so that pulses could be seen on the oscilloscope. This indicated that the overvoltage was positive, and the APD was avalanching. The count rate was determined in both light and dark conditions. Data was taken at a few different bias voltages at each set temperature. When a few overvoltages at the desired temperature were examined, the temperature was lowered by about 10 degrees and the process was repeated. The probed temperatures spanned the range from positive temperatures ($^{\circ}\text{C}$) down to approximately -70°C , the coldest temperature achievable with one power supply used for cooling.

9.2 Dark Counting Rate of the Detector, Effect of Temperature

Perhaps the most important test was to see if the dark count rate could be lowered by cooling to very low temperatures. It was hoped that the dark count rate could be brought down into the single digit range by cooling the APD to -70°C . It is important to measure the dark count rate at the same sensitivity over a range of temperatures. To do this, the overvoltage was adjusted so that a uniform counting rate of approximately 750 KHz was established while the uniform light source was applied to the APD. The light was then turned off and the dark count rate was recorded.

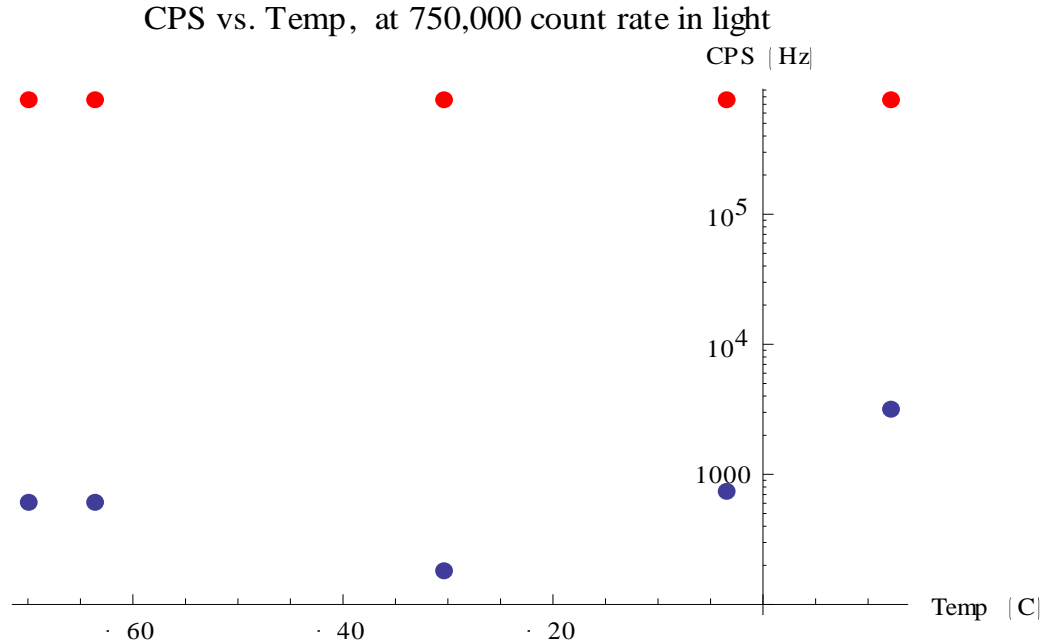



Figure 9.1: The count rate of the detector under both uniform lighting and dark conditions.

Blue dots = data with no light.

Red Dots = data with light applied to the APD.


What was found was that the dark count rate fell with temperature as expected to approximately -30 °C. Then it began to rise with further cooling. This was unexpected. This may result from charge carriers becoming trapped in unstable positions at the very cold temperatures, leading to afterpulses that were not correctly identified. This effect will be carefully examined as the project continues. The lowest dark count rate was 180 counts per second at -30 °C. Unfortunately, there was not enough data recorded to show the behavior at small temperature increments.



Reprint of Figure 6.5 for comparison: Typical dark count vs. temperature at 5% photon detection efficiency (830 nm) for C30902 Silicon APD in Geiger mode. (Source: [6])


9.3 Sensitivity of the Detector, Effect of Overvoltage

A plot of the sensitivity of the APD to overvoltage was made of all the data taken. Ideally, this plot should follow the same form as the sensitivity plot from the datasheet. The acquired data spans all temperatures.



Reprint of Figure 6.3:
Geiger mode photon
detection probability
vs. overvoltage at 22
°C for the C30902
Silicon APD.
(Source: [6])

The counting rate due to overvoltage is plotted for data from the light and data from the dark on two separate plots. The data must be plotted on a linear scale so that it may be compared to the datasheet. Since there are a few orders of magnitude difference between the light plot and the dark plot, they could not be easily plotted together. The datasheet lists the detection probability (%). Our module has not yet been calibrated to a known light source, so the detection probability is still unknown.






Figure 9.2: Data from all temperatures of the counting rate versus overvoltage. There is significant noise in the graphs. This may result from error in the

9.4 Sensitivity of the Detector, Afterpulse probability

The VI is able to discriminate an afterpulse from a novel pulse. The first pulse in a train of pulses is considered to be a novel pulse. All other pulses following the novel pulse in close succession are counted as afterpulses. The afterpulse probability was computed by dividing the total afterpulses by the total of all pulses. The data was acquired from the same dataset that is shown in Section 10.2.




Figure 9.3: Afterpulse probability with respect to temperature at uniform sensitivity. Blue dots = data with no light. Red Dots = data with light applied to the APD.

One result that I find somewhat unusual is that afterpulses are more likely to occur when the counting rate is low. This may result from the initial (novel) pulses being larger in magnitude the longer the APD is at rest. However this was not experimentally verified, and does not fully explain why subsequent afterpulses are also more likely to produce afterpulses of their own (afterpulse probabilities exceeding 50 %) A future improvement would be to histogram how many afterpulses occur after each novel pulse, to see what effect the counting rate has on this relationship.

Another surprising result was that after pulse rates went down with decreasing temperature. This is contradictory to another literature article which explains that at lower temperatures, charge carriers move more slowly, hence they are more likely to be trapped in unstable positions. This should have the effect of raising the afterpulse probability with decreasing temperature.

If the afterpulses are happening later than expected, they will not be identified by the counting program as such. Instead, these delayed afterpulses would be tallied as dark counts. This could simultaneously explain why afterpulses decrease at lower temperatures, while the dark count rate paradoxically increased.

9.5 Future Improvements

Unfortunately, there were many unexpected complications during the project development. It was difficult to keep up with the project timeline. Regrettably, there was not enough time to take all of the data that we would have liked. The project will continue to be refined. It is currently making great progress. The initial data results reveal many areas of interest. The module seems to have the best characteristics around -30 °C. Temperatures around this range will be extensively characterized to see if the dark count rate can be further lowered.


A limitation of the pulse counting program is that the VI must individually request, and examine, 10,000 screenshots from the oscilloscope to count the pulses from 0.4 seconds of time. It takes slightly over an hour for the VI to examine this much data. If the VI is going to be used to examine data from an applied experiment, some other counting interface may need to be arranged. A possible solution to this dilemma would be to incorporate a data buffering mechanism between the oscilloscope and the VI. It would be ideal if 10 contiguous seconds worth of data could be captured, and then streamed in segments to the VI for processing.

The detector will need to be used with a calibrated light source to get a measure of its actual detection probability. Plans for this are currently underway. Most the original design goals have been met. The lowest dark count rate yet identified is 180 counts per second¹⁶. With respect to dark count, our photon detector is already competitive with the commercially available ones, and there is much room for improvement. To me, this qualifies as a success.

¹⁶ Even more recent data that did not make it into the thesis has found temperature (-32 °C) where the dark count was measured to be 80 Hz

Appendix A

A diagram of the electronics used for the APD quenching circuit.



Sources

- [1] Y. Kim, V. Makarov, Y. Jeong, and Y. Kim, Silicon Single-Photon Detector with 5 Hz Dark Counts, Conference on Lasers and Electro-Optics/International Quantum Electronics Conference, OSA Technical Digest (CD) (Optical Society of America, 2009), paper JThE103.
- [2] Mario Stipčević, Active quenching circuit for single-photon detection with Geiger mode avalanche photodiodes, *Applied Optics*, 48, (2009) 1705-1714
- [3] M. Ghioni, S. Cova, F. Zappa, and C. Samori, Compact active quenching circuit for fast photon counting with avalanche photodiodes, *Rev. Sci. Instrum.* 67, (1996) 3440-3448
- [4] I. Rech, I. Labanca, G. Armellini, A. Gulinatti, M. Ghioni, and S. Cova, Operation of silicon single photon avalanche diodes at cryogenic temperature, *Rev. Sci. Instrum.* 78, (2007) 063105
- [5] L. Li and L. M. Davis, Single photon avalanche diode for single molecule detection, *Rev. Sci. Instrum.* 64, (1993) 1524-1529
- [6] http://www.excelitas.com/downloads/DTS_C30902.pdf (APD datasheet)
- [7] http://sales.hamamatsu.com/assets/pdf/parts_H/m-h7421e.pdf
- [8] <http://www.scitecinstruments.pl/detektory/pcm/pdf/SPCMAQR.pdf>
- [9] <http://www.tetech.com/FAQ-Technical-Information.html>
- [10] http://sales.hamamatsu.com/assets/pdf/parts_H/m-h7421e.pdf
- [11] http://en.wikipedia.org/wiki/File:SPAD_Cross-section.gif
- [12] <http://www.kryotherm.ru/imagez/teceng.jpg>
- [13] John S. Steinhart, Stanley R. Hart, Calibration curves for thermistors, *Deep Sea Research and Oceanographic Abstracts*. 15, (1968) 497-503

