

AN AVALANCHE PHOTODIODE-BASED GAMMA CAMERA
PROTOTYPE

DESIGN AND CONSTRUCTION OF AN AVALANCHE
PHOTODIODE-BASED GAMMA CAMERA PROTOTYPE

BY JAMIE L.S. SMITH, B.ENG
Medical Physics and Applied Radiation Sciences

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree
Master of Sciences
McMaster University

©Jamie Smith. November 2007.

MASTER OF SCIENCE (2007)
Medical Physics and Applied Radiation Sciences
McMaster University
Hamilton, Ontario

TITLE: Design and Construction of an Avalanche Photodiode-Based
Gamma Camera Prototype
AUTHOR: Jamie L.S. Smith, B.Eng
SUPERVISOR: Dr. Troy H. Farncombe
NUMBER OF PAGES: xii, 106

Abstract

Objective: The design of a prototype avalanche photodiode gamma camera and an adaptable and expandable acquisition board. To test the performance of the acquisition board a photomultiplier base camera is constructed and utilized in various configurations. The APD-based gamma camera is tested on the same acquisition board to evaluate its performance.

Results: Spectra have been collected and analysed utilizing four various detector/system configurations. i) a PET module, utilizing a PMT coupled to a pixilated BGO crystal with a high energy Na-22 source. ii) a single PMT extracted from a PET module akin to the one mentioned previously. and coupled to NaI(Tl) with a Co-57 source. iii) the same PMT coupled to a CsI(Tl) crystal with a variety of sources. including Tc-99m, Ga-67 and Co-57. iv) an APD coupled to a CsI(Tl) crystal with a variety of sources including Tc-99m and Co-57. Images have been collected from multiple camera configurations, all by a 2 by 2 array. APD's with CsI(Tl), PMT's on NaI(Tl) and a PET module with PMT's on BGO.

Conclusion: These requirements have been met with the current design and the system has been tested in a *controlled* environment. collecting spectrum while operating a single channel and for some configurations, images were collected operating a multi-channel array. A number of steps can be taken to improve upon this prototype. The detector itself may benefit from a pixilated crystal and current APD technology. The acquisition board should be transferred to a multiple printed circuit board design allowing it to be implemented in a black-box manner. Additionally, the current USB bridge limits the transfer rate of data between the DSC and the PC. as a result, alternatives should be investigated.

Acknowledgements

This work would not have been possible without the encouragement of all my family and friends. I have been very fortunate to have so many friends and family to lean on for their support throughout this experience - far too many than I can list here, but I hope they know who they are. Also, I would specifically like to thank Kenrick Chin for his advice and guidance throughout the project.

This work was made possible by generous financial support from the
Canadian Arthritis Network.

Contents

1	Background	1
1.1	Introduction	1
1.2	The Atom	1
1.2.1	Introduction	1
1.2.2	Atomic Structure	2
1.2.3	Nucleus	2
1.2.4	Classifications	2
1.2.5	Mass and Energy	3
1.2.6	Nuclear Stability	4
1.3	Radioactive Decay	4
1.3.1	Introduction	4
1.3.2	Mechanisms of Decay	5
1.3.3	Decay Modes	5
	a) Alpha (α)	6
	b) Beta- (β^-)	6
	c) Beta+ (β^+)	6
	d) Electron Capture (EC)	6
	e) Isomeric Transitions (IT)	7
1.3.4	Decay Rates and Half-Life	7
1.4	Photon Interaction with Matter	9
1.4.1	Photoelectric Effect	10
1.4.2	Compton Scatter	10
1.4.3	Pair Production	11
1.4.4	Attenuation of Gamma Rays	12
	Atomic Number	12
	Density	13
	Thickness	13
	Gamma Ray Energy	13
	Mass Attenuation Coefficient	13
1.5	Conclusions	14

2	Radiation Detection	15
2.1	Radiation Detection	15
2.2	Current Gamma Cameras	15
2.2.1	Collimator	16
2.2.2	Scintillation Crystal	19
	Physical Geometry	22
2.2.3	Photosensors	23
	Photomultiplier Tubes	24
	Avalanche Photodiode	27
2.2.4	Amplification and Shaping	32
2.2.5	Multi-channel Analyzer	33
	Analog to Digital Conversion	34
2.2.6	Acquisition Modes	38
2.2.7	Reconstruction & Correction	38
2.3	Applications	40
2.3.1	Current Imaging Applications	40
	Dual Modality	40
2.3.2	Project Application	42
	Rheumatoid Arthritis	43
2.3.3	Future Applications	44
3	Radiation Detection based on APD's: Results	45
3.1	Scintillation Crystal: CsI	45
3.1.1	Specifications	45
3.2	Selection of APD	46
3.2.1	APD Dark Current (I_d) vs. Bias	49
3.2.2	APD Gain vs. Bias	51
3.3	Preamplifier: CR-100	53
3.3.1	Noise Sources	54
3.4	Pulse shaper: CR-200	55
4	Multi-Channel Analyzer	57
4.1	Peak Detection	57
4.2	Event Trigger	59
4.3	Digital Signal Controller (DSC/DSP)	61
4.3.1	SCI	62
4.3.2	ADC	62
4.4	Master CPU	64
4.5	System Operation	65
4.6	System Resolution	67
4.7	Conclusion	68

5	Single Channel Experiments with CsI and an APD	70
5.1	Single Crystal - Single APD	70
5.1.1	Isotope Energy Spectrum	74
5.2	Double Sized Crystal - Single APD	76
5.3	Effects of Magnetic field on APD Spectrum	78
6	APD Gamma Camera	80
6.1	Images	80
6.1.1	PET Module	81
6.1.2	PMT on NaI(Tl)	82
6.1.3	APD on CsI(Tl)	83
7	Conclusions	87
	Appendices	90
A	APD Pulse Size Sample Calculation	91
B	Peak Detection Circuits	93
C	Digital Signal Controller Code	95
D	Digital Signal Controller Photo	102

List of Figures

1.1	Radioactive Decay	9
1.2	Photoelectric Effect	10
1.3	Compton Scatter	11
1.4	Interaction Dominance	12
1.5	Gamma Absorption	14
2.1	Pinhole Collimators	17
2.2	Pinhole Image Magnification	17
2.3	Parallel hole collimation	18
2.4	Scattered Event	19
2.5	Compton Smear	20
2.6	Crystal Properties	22
2.7	Photomultiplier Tube	24
2.8	Force on Electron in a Magnetic Field	26
2.9	P-N Junction	28
2.10	APD Design	29
2.11	APD Signal, noise and gain	31
2.12	Preamplifier Pulses	32
2.13	Gaussian Shaper Pulse Profile	33
2.14	Analog MCA Design	34
2.15	Non-linearity in Analog to Digital Conversion	37
2.16	Anger Logic Diagram	39
3.1	CsI Emission Spectrum	46
3.2	S8644-1010 Response Spectrum	47
3.3	APD Active Area	47
3.4	APD Photo and Dark Currents	50
3.5	APD Photo and Dark Currents	52
3.6	Optimal APD Gain Bias Voltage	53
3.7	Preamplifier Coupling Circuit	54
3.8	Optimal APD Gain Bias Voltage	55

4.1	Effects of limiting slew rates	58
4.2	Peak Detector: Droop rate	58
4.3	Typical Op amp Summing Circuit	60
4.4	Op amp Summing Circuit with voltage follower buffers	60
4.5	Analog to Digital Converter: Functional Block Diagram [12]	64
4.6	Analog to Digital Converter: Timing Diagram [12]	65
4.7	Analog to Digital Converter: Oscilloscope Traces	67
4.8	Spectrum collected from 50Hz Reference Pulser	68
4.9	Functional Diagram of gamma camera acquisition system	69
5.1	APD A Spectrum	71
5.2	APD B Spectrum	72
5.3	APD A Spectrum	73
5.4	1x2 Crystal Experiment	75
5.5	a) Gallium-67 Spectrum with an APD detector, b) Ga-67, Co-57 and Tc-99m spectrum collected with a PMT detector	76
5.6	1x2 Crystal Experimental Setup	77
5.7	1x2 Crystal Experiment Spectrum	77
5.8	APD spectrum in Magnetic Field	79
6.1	Image from PET module - 5 minute scan	81
6.2	Image from four PMT coupled to NaI crystal.	82
6.3	APD Image Grid	84
6.4	APD images	86
B.1	PKD01 Circuit Diagram with Trigger threshold	93
B.2	Hex CMOS Inverter used to convert 3.3V logic to 5V logic to operate the PKD01	94
B.3	Summing Circuit - output connects to Trigger Threshold Input	94
D.1	Digital Signal Controller Photo	102

Organization of this Thesis

Chapter 1

This chapter will introduce the fundamentals of radioactivity. Understanding the nature of radiation is crucial to the study of emission tomography - radiation and its interaction with matter is the basis upon which everything in emission tomography is built. Everything in emission tomography begins with the presence of radiation and our ability to detect its presence. This chapter first presents the sources and characteristics of radiation and reviews interactions of gamma rays with matter that are most commonly encountered in emission tomography. With an understanding of the nature of radiation, discussions about detection and measurement of radiation will follow.

Chapter 2

Chapter 2 aims to present a general overview of how radiation is detected in a typical gamma camera. The individual components of a gamma camera are detailed here - discussions present the material in a manner designed to highlight the issues that were encountered in the design of the proposed gamma camera. In addition, this chapter serves to provide a background against which we can later compare and contrast our gamma camera design to typical clinical cameras currently in operation. Finally, this chapter will close by discussing the project objectives and the details of its ultimate application, and potential future applications.

Chapter 3

In this chapter we present the details of the new solid-state camera detector design. In general, the overall approach to detecting radiation has not

changed however, new components allow for new capabilities. The specifications and test resolution of these new components are presented here in addition to, where applicable, any new constraints or obstacles introduced as a result of the updated components are also discussed.

Chapter 4

This chapter details the design and operation of the data acquisition system developed specifically for this gamma camera. An indepth review of the individual components, their operating specifications and tested capabilities is provided. The component integration and system operation is described at length.

Chapter 5

An imaging system, at its heart, is an array of spectrum collecting detectors. In this chapter the spectra collected from various detector arrangements are presented and discussed. Spectra collected from PET detector modules, single PMT's mounted to CsI(Tl), NaI(Tl) and APD's mounted to CsI(Tl) are all discussed. The light response and bias effects on APD performance is evaluated in detail. The effects of in crystal light spreading is investigated and the magnetic robustness of APDs is evaluated.

Chapter 6

Chapter 6 moves from a single detector arrangement to an array of detectors. Images collected, again from various detector arrangements are presented and discussed.

Chapter 7

Chapter 7 presents a review of the achievements of the project and a discussion of the future work. The areas of potential improvement are identified and specific improvements are discussed.

Chapter 1

Background: Nuclear Imaging

1.1 Introduction

Nuclear imaging a very beneficial diagnostic tool that allows the collection of images that offer not anatomical information, but functional information. This special ability relies on two unique properties of radiation: First, radiation (depending on the energy) can pass through relatively large amounts of matter (ie. exit the body) and secondly, it originates from unstable atoms which can be attached to different molecules. What this means is that a batch of some chemical (sugar for example) can be marked with a radioactive atom on each molecule and injected into a patient. The molecule will then travel to locations in the body specific to its biological role, and all the while emitting radiation that can be detected from outside the body. There are many types of radioactive decay, but ultimately nuclear imaging deals with detecting gamma rays (in current clinical environments).

1.2 The Atom

1.2.1 Introduction

For thousands of years the atom has been considered the basic building block of the physical world. In the last hundred years or so great scientific advances have been made about our understanding of atomic structure and mechanics. From the plum pudding model proposed by J.J. Thompson in 1903 to the Bohr model in 1913 and to current quantum mechanics - our understanding has grown in leaps and bounds, and with it, so has our repertoire of phenomena that we can exploit to our advantage.

This section will discuss the various properties of the atom and more specifically how those properties can be utilized in the world of nuclear imaging.

1.2.2 Atomic Structure

Modern atomic theory models the atom as two distinct components, the *nucleus* and the *electron cloud*. It can be thought of as a small core of mass composed of neutrons and protons with nearly massless electrons orbiting around it. The electron is considered to have negative charge, the protons have an opposite charge (positive) and the neutrons have no electric charge.

1.2.3 Nucleus

The nucleus, composed of neutrons and protons, makes up the vast majority of the mass of the atom. A proton and neutron have roughly equal mass, but an electron contributes only $1/2000^{th}$ the mass of a proton to the total mass of an atom. The electrons orbit the nucleus in stable energy states only changing states if they gain or lose energy.

1.2.4 Classifications

Due to the sizable array of atoms known to man, numerous attempts of organizing elements has been taken. The current periodic table remains the most robust model to date. Any element on the periodic table can be described with the following notation,



Where A is the mass number, Z is the number of protons or atomic number and X is the elemental short form (i.e. H for hydrogen or Na for sodium). In fact it is actually redundant to specify both the atomic number and the element as they provide identical information; however, it is often easier to recall an element by name than by number of protons.

Atoms with similar subatomic arrangements share similar properties and because of this, various groupings have been defined to aid in classification of atoms. Isotopes are groupings that all contain the same number of protons meaning they have the same atomic number, Z (ex. 1H and 2H both have 1 proton). A table of common isotopes used in nuclear

imaging is provided in table 1.1 on page 7. If nuclides share the same number of neutrons, they are referred to as isotones (ex. ^{12}B and ^{13}C both have 7 neutrons). If the sum of neutrons and protons for two groupings is the same then they are considered to have the same mass (ie. the same mass number A) because electrons are relatively insignificant, as such these groupings are referred to as isobars. If two groupings share the same number of protons and neutrons but are in different energy states they are nuclear isomers (ex. ^{99}Tc and ^{99m}Tc are nuclear isomers).

1.2.5 Mass and Energy

The Atomic Mass Unit (amu) was introduced because the metric standard, the kilogram, was far too large to be practical for dealing with subatomic particles. One amu is equivalent to 1.66×10^{-24} grams. As such, the mass of a neutron is 1.00783 amu and a proton 1.00866 amu. An electron 'weighs in' at only 0.00055 amu. Like the amu, the electron volt or eV was introduced because the joule was also too bulky. One eV is equivalent to 1.602×10^{-19} Joules. It is derived from the amount of kinetic energy an electron gains by passing through an evacuated electric field of one volt potential.

When the various constituents come together to form an atom, their combined mass is actually less than the sum of the individual parts – this is the result of the parts, together, reaching a lower energy state. The difference in mass is realized in energy lost by the atom when it is formed as described in the equation below.

$$\Delta m = Zm_p + (A - Z)m_n + Zm_e - m_{Z,A}$$

Where Z is the number of protons, m_p denotes the mass of a proton, $(A - Z)$ is the number of neutrons, m_n is the mass of a neutron, m_e is the mass of an electron and $m_{Z,A}$ is the mass of the formed atom. The energy associated with this mass difference is referred to as the binding energy for the atom.

For example, given a ^4He nucleus has a total mass of 4.003 amu but sum of two neutrons, two protons is actually 4.033. When the atom is formed it reaches a lower energy state and the excess energy, the binding energy, can be thought of as the energy needed to hold the nucleus together.

1.2.6 Nuclear Stability

The binding energy of the nucleus needs to exceed the electrostatic repulsion of neighbouring protons. The like charges will try to repel and if the binding energy of the nucleus is not sufficient, they will succeed. As a consequence, the number of neutrons must increase quicker than protons to maintain an dominant binding energy. The result is that typically, a stable atom will have more neutrons than protons, though at low proton counts, atoms may find stability with equal numbers of protons and neutrons. Having a greater number of neutrons than protons tends to contribute to stability, but other factors play a role as well, such as spin coupling of the proton and neutrons.

1.3 Radioactive Decay

1.3.1 Introduction

Radioactive decay is the processes by which an unstable atom releases some energy in order to move towards a more stable energy state. The energy released is often classified as one or more of three types, alpha (α), beta (β) and gamma (γ). The gamma emission is the only one of the three that does not have mass. A particular grouping of protons and neutrons is referred to as a nuclide; if that grouping is unstable and emits radiation it is referred to as a radionuclide.

A stable state for an atom is referred to as the ground state. In a stable atom, the associated protons, neutrons and electrons all occupy their ground energy state. When one of these particles is found in an elevated energy state, the atom is considered to be in an *excited state*. At some point the elevated energy of the atomic particle will be released and the particle will return to its stable energy state.

The time spent in this elevated state can be very short or very long – some times have been measured to be as short as 10^{-23} seconds and some as long as 10^{19} years. This process is referred to as radioactive decay and is a probabilistic event and so these times are better understood as averages over an extended period of time. A consequence of the probabilistic nature of radioactivity in a counting experiment designed to count the decays of an amount of radioactive material over a set period of time, the uncertainty in the counts is equal to the square root of the counts. This indicates that

more accurate measurements of activity require either longer measurements or higher count rates. The associated probability distribution is best described as a Poisson distribution, which is used to describe an event that happens rarely but has very many opportunities to happen. However, the associated derivation and mathematical application is beyond the scope of this thesis and will not be discussed here.

The original atom is referred to as the *parent* and the newly created atom is the *daughter*. The time spent for an atom, on average, in the excited state is inversely proportional to how radioactive it is. That is, an isotope that takes years, on average, to decay is considered to be only slightly radioactive, while an isotope that decays relatively quickly is referred to as highly radioactive or sometimes as *hot*.

1.3.2 Mechanisms of Decay

There are four conditions that must be accounted for in order for a decay to occur spontaneously. These are

- Net Charge
- Baryon Number (quarks less antiquarks)
- Lepton Number (leptons less antileptons)
- Mass / Energy

If, over all four of these conditions, the energy defect is negative the reaction will require extra energy input to cause a decay which can come in the form of external radiation, for example.

1.3.3 Decay Modes

The most common decay modes are Alpha (α), Beta-minus (β^-), Beta-plus (β^+), Electron Capture (*EC*) and Isomeric Transition (*IT*). Alpha, Beta-minus and Beta-plus decay are named after the particles they emit in the decay process. An alpha particle is two neutrons and two protons from the nucleus of an atom. A beta-minus particle is an electron and beta-plus particle is a positron.

a) Alpha (α)

In alpha decay the energy of the alpha particle is proportional to the difference in the mass of the unstable atom and the two resulting parts (the alpha particle and the resulting daughter product). As shown in equation 1.1, a radioactive particle X decays into Y and produced an alpha particle and releases energy.

**b) Beta- (β^-)**

In beta-minus decay a neutron becomes a proton. To balance the creation of the positive charge of the proton an electron is also created and thus no net change in charge occurs. see equation 1.2. The energy of the beta-minus particle is proportional to the difference between the mass of the parent and the mass of the daughter product, shared between the e^- (electron) and $\bar{\nu}_e$ (antineutrino). The mass of the electron is sufficiently small enough to be ignored for most calculation requirements.

**c) Beta+ (β^+)**

Beta-plus decay is the process through which a proton becomes a neutron. shown in equation 1.3. Again, to balance charge, a positive charge is created in the form of a positron. The energy of the beta-plus particle is proportional to the mass of the parent less the mass of the daughter and the mass of two electrons. Also, in both beta decays, it should be noted that technically an antineutrino or neutrino ($\bar{\nu}_e$ and ν_e respectively) are also created to conserve lepton number.

**d) Electron Capture (EC)**

When an atom contains too many protons in the nucleus and not enough energy to emit a positron (as in beta-plus decay) then electron capture is the sole mode of decay (equation 1.4). The result is the emission of a neutrino and an x-ray often from the K-shell of the daughter atom. The energy, which is shared between the neutrino and the x-ray is proportional to the difference in mass between the parent and daughter products.



e) Isomeric Transitions (*IT*)

There are two steps by which an isomeric transition occurs. gamma-emission and internal conversion. First internal conversion: the excited nucleus interacts (typically) with the K-shell electron and ejects it (provided the energy exceeds the binding energy) as in equation 1.5. As such, this process will appear like beta-minus decay, but because the particles are defined by their production and not their composition, it is not beta-minus decay. Secondly, the ejected electron leaves a hole in the shell, which, when filled, results in the emission of an x-ray. As shown in equation 1.6.



Table 1.1: Common Radiotracers [1]

Element	Imaging Energy (keV)	Half-life
Carbon-11	511	20.3 m
Nitrogen-13	511	10.0 m
Oxygen-15	511	2.0 m
Fluorine-18	511	110.0 m
Gallium-67	93, 185, 300	3.3 d
Rubidium-82	511	1.2 m
Technetium-99m	141	6.0 hr
Indium-111	171, 245	2.8 d
Iodine-123	159	13.3 hr
Iodine-131	364	8.0 d
Thallium-201	72	72.9 hr

1.3.4 Decay Rates and Half-Life

The time that it takes for an amount of radioactive isotope to decay to half of its original activity is termed the half-life ($t_{1/2}$). Connected to this idea is the mean lifetime (τ) which is described as the average amount of time a particle can expect to exist before undergoing decay. Another

term which is used often when discussing decay rates is the decay constant (λ) and it is simply the inverse of the mean lifetime. The activity (A) describes the number of decays a substance of N particles undergoes per second and is derived here. For simplicity assume that once an atom decays it does not decay again. The change in the number of radioactive atoms is equal to the number of atoms present (N) times some decay constant (λ) and times the change in time (t) as in equation 1.7.

$$-\Delta N = \lambda N \Delta t \quad (1.7)$$

In general the equation can be written as.

$$dN = -\lambda N dt \quad (1.8)$$

Rearranging equation 1.8 gives.

$$\frac{dN}{N} = -\lambda dt \quad (1.9)$$

Integrating this yields equation 1.10.

$$[\ln N] = [-\lambda t] + C \rightarrow N(t) = C'e^{-\lambda t} \quad (1.10)$$

where C is a constant of integration. Solving equation 1.10 for $N(t)$ yields

$$N(t) = N_0 e^{-\lambda t} \quad (1.11)$$

Since activity is a measure of the number of decays in a unit of time, we can define the activity at a given time with respect to how many atoms change to daughter products.

$$A(t) = -\frac{dN}{dt} = \lambda N(t) \quad (1.12)$$

Bringing together equation 1.12 and 1.11 and substituting λN_0 for A_0 then the activity at time t is given by.

$$A(t) = A_0 e^{-\lambda t} \quad (1.13)$$

The activity is dependant on the number of radioactive particles in the sample, but for any given isotope the half-life is considered constant and is described by equation 1.14 by solving equation 1.11 for the case of $t = t_{1/2}$ and $N(t_{1/2}) = \frac{1}{2}N_0$.

$$t_{1/2} = \frac{\ln(2)}{\lambda} = \tau \ln(2) \quad (1.14)$$

Half-life is measured in seconds, minutes, hours, days or even years and the units for activity are often measured in either becquerels (Bq), the number of disintegrations per second, or historically in curies (Ci), 3.7×10^{10} disintegrations per second.




Figure 1.1: Radioactive Decay of ^{99m}Tc

As shown in Figure 1.1 once the activity of an amount of radioactive material and its half life are known, the activity can be predicted for any time after or before that point. In the case above, notice that the activity drops from 100% to 50% in about 6 hours, corresponding to the half life of $^{99m}Tc \approx 6$ hours.

1.4 Photon Interaction with Matter

As a photon travels through matter there is a certain probability that it will interact with the atoms of that material. This probability is dependant on the energy of the photon and the properties of the material it is traveling through. It is this interaction with matter that allows photons to be detected and their energy measured which is desirable for event localization in nuclear imaging. There are three outcomes for a photon traveling through matter; it will be absorbed, scattered or transmitted. If it is absorbed or scattered it is considered to be *attenuated*. The probability of photon attenuation increases with material thickness, material density and material Z number; it decreases with increasing energy of the photon.

Three primary interactions are considered when investigating photon interactions with matter. *Photoelectric Effect*, *Compton Scatter*, and *Pair Production*. A brief overview of these is presented in the following subsections:

1.4.1 Photoelectric Effect

This process occurs when an incident photon is absorbed by a bound electron and is therefore not possible with free electrons. The bound electron is ejected with its newly gained energy less its original binding energy, leaving a hole for an outer electron to drop into, as in figure 1.2. The excess energy can be absorbed by an outer shell electron. When this second interaction occurs the ejected electron is referred to as an *Auger electron*. The probability of this photoelectric interaction occurring depends strongly on the atomic number (Z) of the material the photon is traveling through. Typically this ability to interact with the original gamma ray is referred to as the stopping power.




Figure 1.2: Photoelectric Effect Diagram

1.4.2 Compton Scatter

When a photon scatters off an electron in a classical collision type manner the result is a sharing of energy between the electron and the photon and often a change in direction as depicted in Fig. 1.3. Being similar to a classical collision the distribution can be derived as in Eq. 1.15 - however, because the photons entering and exiting are different it is not considered classical scattering.

$$hv' = \frac{hv}{1 + (hv/m_e c^2)(1 - \cos(\theta))} \quad (1.15)$$

where hv' is the scattered photon energy, hv is the incident photon energy, m_0 is the electron rest-mass, c is the speed of light, and θ the angle of scatter.

In general, the greater the energy of the incident photon the more likely it will be scattered forward (ie smaller θ , but also the more energy loss (as a percentage of its original energy) it will suffer if scattered off its path. The actual angle of scatter can be accurately derived from the




Figure 1.3: Compton Scatter Diagram

Klien-Nishina formula which provides a differential cross-section with respect to scattering angle. Derivation of the Klien-Nishina formula is not within the scope of this paper.

1.4.3 Pair Production

Pair production occurs when the energy of the incident photon is greater than twice the rest mass energy of the electron (2×511 keV). If the photon passes close enough to the nucleus of an atom it may disappear and reappear as a positron and electron. The positron soon undergoes annihilation with an electron and emits two gamma rays each of 511 keV. Pair production, as shown in figure 1.4, requires either an atom with significantly high atomic number to occur at lower energies. As a result, Pair production is not typically considered a dominant effect in nuclear medicine imaging.

The dominance of these three effects varies as a function of energy. As shown in Figure 1.4, a typical clinical tracer (ex. < 1 MeV), see table 1.1 would most likely undergo a single, or not multiple, Compton scatters before being reduced to an energy that is increasingly susceptible to photoelectric processes.

Photoelectric effect is the most desirable interaction in gamma ray detection (in a scintillation crystal) because it produces only electrons - no extra lower energy gamma rays left to interact again. In section A of Figure 1.4 the curve suggests that increased number of protons (Z) in the absorber would increase the likelihood that photons undergo primarily photoelectric effect. For a given energy photon, lower Z valued absorbers would be




Figure 1.4: Interaction Dominance: A) photoelectric. B) Compton. C) Pair Production

increasingly likely to allow Compton scatter (perhaps multiple) before the gamma ray energy is reduced to photoelectric levels. Also possible, is that the Compton scattered photon exits the absorber before depositing any further energy.

1.4.4 Attenuation of Gamma Rays

The macroscopic manifestation of gamma rays interacting with matter is termed attenuation. As radiation passes through a material its intensity is reduced as it deposits energy in the material. An appropriate analogy might be a flashlight covered by the palm of a person's hand. The factors that determine how much light can be seen on the backside of the hand include the makeup of the material in the hand (i.e. bone or muscle or skin), the density of that material, the thickness of the hand and of course the brightness of the flashlight to begin with. For simplicity, the following properties will be investigated under the premise that the attenuating material is made of one element, like iron for example, and not a composite of various materials.

Atomic Number

Typically the attenuation of gamma rays increases with increasing atomic number. The simple explanation is that increased atomic number means more neutrons, protons and electrons per atom with which the

gamma ray might interact. For instance, lead, with 82 protons, is relatively adept at attenuating gamma rays where as aluminium ($Z = 13$) given an equal atomic density would be less effective. The actual relationship is $\Delta I \propto Z^3$, indicating that doubling the atomic number of a material increases its attenuation by a factor of eight.

Density

Conceptually along the same lines as Atomic number, the closer atoms are packed together – the mass in a given volume – the more likely a gamma ray will be attenuated.

Thickness

Naturally, the more material a gamma ray has to pass through the more it will be attenuated on average. This is theoretically interchangeable with density - in reality the factor is simply the number of atoms the gamma ray must pass.

Gamma Ray Energy

Lastly, the more energy the gamma ray has, the less likely it is that it will interact with matter. The probability of a gamma ray interacting with matter in a particular way as described in Section 1.4 is dependant on energy. If each of these interaction responses is superimposed, as shown in Figure 1.5, the resulting absorption coefficient generally decreases with increasing energy.

Mass Attenuation Coefficient

A useful measure of a material's ability to attenuate gamma rays is the mass attenuation coefficient, which is a modified version of the *linear attenuation coefficient*. The linear attenuation coefficient is defined as,

$$\mu_{(lmcu)} = \frac{\ln(2)}{x_{1/2}}$$

Where $x_{1/2}$ is the depth which the gamma rays must penetrate to have the intensity reduced to half of the original ('the half value layer'). The mass attenuation coefficient builds on this to normalize by the density of the material. Thus making it more adaptable to applying to specific scenarios.

$$MassAttenuationCoefficient = \frac{\mu}{\rho} = \frac{\mu_{(lmcu)}}{\rho}$$

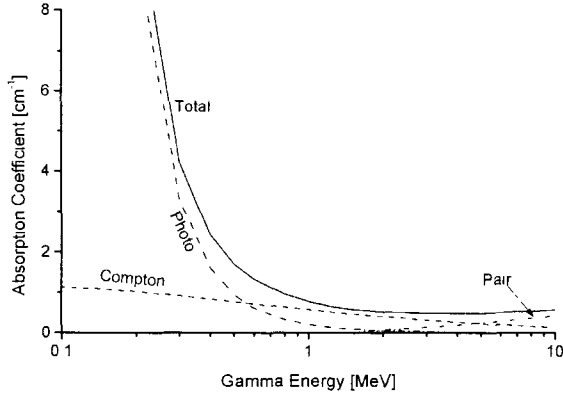


Figure 1.5: Gamma Absorption

The result is a measurement with units of cm^2g^{-1} or area per gram. With comprehensive and widely available databases of Mass Attenuation Coefficients for a large range of gamma ray energies, the actual attenuation of a source of radiation can be estimated rather easily and accurately simply by knowing the density of the material. The general expectation is that the intensity will obey the following relationship.

$$I_r = I_0 e^{-\mu x \rho}$$

where, I_0 is the unattenuated intensity, x is the thickness or depth of the material, ρ is the density of the material, μ_ρ is the Mass attenuation coefficient (not the Linear attenuation coefficient) and I_r is the intensity expected at depth x .

1.5 Conclusions

In the subsequent chapter we will see how the concepts of photon interaction and attenuation come together to result in useful radiation detection. The section presents the general process of gamma camera operation, as well it will discuss the objectives of this project and the background required to understand the related issues.

Chapter 2

Radiation Detector Design Considerations

2.1 Radiation Detection

In general, detectors are classified in one of two ways; by the information they provide or the process they use to detect radiation. For instance, classification by information would include detector types such as *counters* which provide a total number of events in a detector over a specified time for a specified energy range, *spectrometers* which present number of interactions over a number of specified energy ranges for a specified time and *cameras* which provide this same energy information in combination with spatial information for each event. Segregation by detection method would result in groupings such as *Gas chambers*, *Semiconductors* and *Scintillators*. The intersection of the *scintillation* and the *camera* classifications provides a functionality that has lent itself well to medical diagnostic applications. These devices are often referred to as *gamma cameras*.

2.2 Current Gamma Cameras

Most gamma cameras found in any medical imaging facility will operate in virtually the same manner. Perhaps the most logical way to describe the camera is to follow a single photon as it reaches the camera, becomes an electronic pulse and is recorded. The photon originates from a radionuclide that has been located inside the patients body (Ex. injected). As a simple overview the process can be summarized as follows: the first component the photon comes into contact with is the collimator, followed

by the scintillation crystal; from there the photon becomes many lower energy (optical) photons. These low energy photons are spread across a number of photomultiplier tubes. The light that enters a photomultiplier (PM) tube is converted into electrons that cascade down a set of plates gathering more electrons as they go. At the end of the PM tube, the pulse of current is passed to a preamplifier which converts the current pulse into a voltage pulse which it then amplifies further before sending the pulse out to a pulse shaper. The pulse shaper converts the pulse from a long tailed step function into a pseudo-Gaussian pulse where the height of the output pulse is proportional to the energy of the original input pulse. This pulse height can then be recorded and used in combination with the other signals in the array to approximate the original photon energy and the position of the event. These steps will now be described in more detail.

2.2.1 Collimator

In order to have an image appear on an imaging plane, photons must be organized or sorted into locations on that plane that are indicative of their origin; this is called collimation. There are many approaches for achieving this collimation, pinhole and parallel hole collimation are two distinct methods and variations of each do exist. Traditionally, a collimator is made of lead containing an arrangement of holes depending on the design. In a parallel hole collimator, the typical hole shape is hexagonal and the lead walls that separate one hole from the next referred to as septa. In a pinhole camera a single circular hole is common, but variations exist that accommodate multiple pinholes. The size of the hole affects the resolution of the camera as well as the sensitivity – how much radiation it lets through.

In ideal pinhole collimation all photons must pass through a common point (ex. A small hole in a piece of lead.) in order to reach the imaging plane. As shown in Figure 2.1(b) for any given point on the image plane, the pinhole reduces the number of possible locations of gamma ray's source to a single straight line of potential sites of origin. If the lead sheet and pinhole didn't exist then for any given point on the image plane there would be infinite lines of possible origin – that is, any point in front of the plane could potentially emit a gamma ray that would hit that point, as depicted in Figure 2.1(a)

The pinhole collimator inherently offers another advantage – magnification or minification. The location of the pinhole in relation to the




Figure 2.1: (a) Without pinhole collimation no meaningful image is formed. (b) With pinhole Collimation points on the imaging plane are (ideally) mapped to a unique vector.

subject and the imaging plane determines how the image is scaled with respect to the subject size. The magnitude of magnification is equal to the ratio of y/x as shown in figure 2.2. Naturally, when the pinhole is positioned equidistant from both the subject and the imaging plane, there is no magnification.




Figure 2.2: Pinhole Collimation used to magnify an image

The other common collimator is the parallel hole collimator. While the pinhole collimator should be separated from the image plane, the parallel hole collimator should be seated up against the image plane (the

front of the gamma camera). Figure 2.3 depicts how the parallel hole collimator only allows gamma rays that are traveling perpendicular to the image plane to be accepted while rays that stray from these perpendicular paths are removed. The end result is the same as with the pinhole collimator in that each point in the image plane has only one line of response from where the gamma ray could have originated however no magnification is possible in a straight parallel hole collimator. Magnification can be achieved, however, with collimators that do not have their septum perpendicular to the imaging plane – also known as a converging collimator. A collimator with diverging geometry would, as expected, provide minification.




Figure 2.3: Parallel hole Collimation

In a parallel hole collimator, often some gamma rays will penetrate through a septum of the collimator passing into the next hole. When this occurs the gamma rays can scatter in the septum and be redirected toward the detector, and if detected, the event will be measured with a lower energy than a gamma ray that passes straight through the collimator untouched. This effect can also occur in pinhole collimators near the edge of the pinhole.

Energy Acceptance Window

Although collimators are essential, they cannot remove all scattered events.

As seen in Figure 2.4, a gamma ray initially traveling away from the detector can be scattered into the field of view. The events that are scattered can not be differentiated based on their direction; however, they often can be distinguished by a reduced energy. The process is actually executed after the photomultiplier tubes, once the energy information has been collected.




Figure 2.4: Scattered Photons may still reach the detector

By setting a window of acceptance around the known gamma energy (i.e. $140 \text{ keV} \pm 10\%$ for ^{99m}Tc) events that have undergone significant Compton scatter and thus significantly out of place can be identified by a reduced energy and subsequently ignored. Additionally, by having an upper bound of acceptance on the window it is possible to ignore piled up events – that is, when two photons are detected within irresolvable span of time.

2.2.2 Scintillation Crystal

Just as a photon traveling through the body will undergo interactions with matter, so will the photons interact with a scintillation crystal. Converse to in-body scatter, interactions in the crystal are essential to reducing a gamma ray into measurable (lower energy) photons. Two primary interactions are considered when investigating matter interactions in a crystal, Photoelectric Effect and Compton Scatter as described in Section 1.4

The role of a scintillation crystal is to convert high energy gamma rays into low energy optical photons. In order to reduce incoming gamma-rays into low energy photons multiple steps must occur within the scintillation crystal. First the gamma-ray must interact with the crystal to produce free electrons. The crystal is doped to create what are termed *luminescent centers* which are essentially holes; atoms in the crystal structure which have a valence electron deficit. With the electron released by photon interactions in the crystal, it travels in the conduction band and it soon finds a *luminescent center*, fills the hole in the atom's valence and

emits a photon. The energy of the newly emitted photon is that of the electron's energy less the binding energy the electron now experiences in the valence. In NaI, a common scintillation crystal, one 3eV photon is produced for every 25eV (*approximately*) of radiation.

The ideal situation is where the entire photon is stopped via photoelectric interaction because the entire incoming photon is converted into electrons for luminescent centers. However, Compton scatter does occur in higher energy gamma rays that are used in nuclear imaging. The result is that the Gaussian photo peak corresponding to photoelectric interactions is trailed by what is referred to as the *Compton smear*. The result is a spectrum near the photo peak similar to Figure 2.5.




Figure 2.5: Typical Spectrum exhibiting a photo peak and Compton smear

Thallium doped Sodium Iodide (NaI(Tl), discovered in the 1940's, was the first inorganic scintillation material discovered and is still the most widely used scintillator today. There are now many competing crystal composites but each leaves something to be desired. Because no crystal offers the perfect performance for every application, it's important to be able to select a crystal that best suits the needs of the design. There are several figures of merit on which a scintillation crystal is evaluated. Each of these figures is influenced by multiple underlying effects, but the general performance of a scintillator crystal can be roughly gleaned from evaluating a few sweeping parameters. It should be noted that near-endless work can be done at a depth far beyond this generalization in order to optimize the performance of the crystal in a given system. Nonetheless, these guiding parameters are stopping power, emission wavelength, light output and scintillation time.

1. *Stopping power* is directly linked to the density of the crystal, and it is an indication of the ability of a crystal to stop incoming gamma

rays. For example, everything else constant, a lower density crystal would, on average, require a greater thickness to stop a photon of a given energy than that of a higher density crystal. Density is actually a combination of a number of underlying factors, particularly the **crystal pack structure type and the effective Z value of the atoms that make up the crystal unit lattice.**

2. *Emission wavelength* refers to the wavelength that the crystal will emit, and is typically referred to as the wavelength of maximum emission. The maximum is quoted because emissions are not monochromatic, but cover a fairly wide spectrum. **This figure is important in matching the crystal effectively to a photomultiplier and should ideally align with the peak response of the photomultiplier.** Worth mentioning is that when designing a crystal it is favourable to have the emission wavelengths not overlapping the absorption wavelengths to avoid self-absorption of emission light.
3. *Light output or light yield* determines the number of photons produced per keV of incident gamma ray energy. As expected, a crystal with low light output will require much more sensitive measurement and will be more susceptible to noise. As such, the output of a crystal can greatly improve the performance of a system.
1. *Scintillation time* is a measure of how long it takes the crystal to luminesce following an interaction. This time is actually a combination of the rise time and the decay time. The rise time is **strongly influenced by the hole-mobility of the crystal. It is a measure of the time it takes to reach maximum light intensity, and is typically measured as the time it takes to rise from 10% intensity to 90% intensity [1].** The decay time is defined as the time for the pulse to fall from its maximum to an intensity that is $1/e$ of the maximum. **This figure of merit is important in determining the count rate limit of the crystal. Put simply, shorter pulses can be spaced closer together than long pulses.**

Figure 2.6 summarizes some of the common scintillation crystals on the market today.

The numbers in Figure 2.6 suggest why Sodium Iodide (NaI) has to date been the most commonly used scintillation crystal. Despite its lower density, it offers an emission peak which matches well with most photomultiplier tube spectral response peaks and it produces a significantly

Material	Density (g/cm ³)	Z _{eff}	Emission Peak (nm)	Decay Time (μs)	Light Yield (photons/MeV)
NaI	3.67	51	415	0.23	37,700
Bi ₄ Ge ₃ O ₁₂	7.13	76	505	0.30	8,200
CsI(Tl)	4.51	54	540	0.68	64,800
BaF ₂	4.89	53	310,220	0.62, 0.0008	9,950
GSO	6.71	59	430	0.06	10,000
LSO	7.4	65	420	0.04	30,000
LaBr ₃	5.08	47	360	0.0016	80,000
LaCl	3.64	60	350	0.03	50,000

Figure 2.6: Common Scintillator Crystal Properties [1–4]

large number of photons in a very short decay time. However, a NaI crystal would have a poor emission peak for an avalanche photodiode (APD) detector, as discussed later. APD spectral response is typically at shorter wavelengths, thus CsI crystals find useful application in solid-state detector systems. The CsI(Tl) produces a very desirable amount of light, but its drawback is the relatively slow decay time for the response - ranging from $0.68\mu\text{s}$ to as high as $1.0\mu\text{s}$ thus potentially limiting the cameras achievable count rate. The high density of the $\text{Bi}_4\text{Ge}_3\text{O}_{12}$ (BGO) crystal means that it is more effective at stopping higher energy photons than most crystals at comparable thickness. This advantage means BGO can often be found in PET cameras, as their operation depends on the collection of 511keV gamma rays. Not included in this table, cost is another aspect which often has an affect on final design outcomes.

Physical Geometry

Beyond the inherent properties of a crystal there are macroscopic modifications that can be made to improve the performance of a scintillator. The simplest gamma camera arrangement is a single solid crystal mounted to an array of photomultipliers. However, when an interaction occurs the light spreads out inside the crystal. If the light is spread too thin then the light might be lost in the background noise of some photomultipliers thus distorting both the total energy and the position information for that event. For instance, in the extreme case, the light is spread across all

photomultipliers such that no detector experiences anything above normal background and the event goes undetected. Contrariwise, a certain amount of light spreading is desirable. In order to extract location information the light must be shared by a group of photomultipliers.

The thickness of a crystal is not only selected to optimize light spreading, but also to ensure that incoming photons can be effectively stopped. For instance, if a crystal was too thin then perhaps an incident photon, after a single Compton scatter, could exit the crystal, not depositing its entire energy in the crystal. Light spreading can be adjusted by controlling the crystal thickness, however, only beyond the thickness required to stop the incident photons.

In working with these limitations certain techniques have been developed to allow increased control over light distribution. By cutting the crystal into what is essentially pixels and separating the pixels with reflective material the light can be guided down to the individual photomultipliers. However, if all the light is directed solely to a single photomultiplier then the spatial resolution is limited by the size of the photomultiplier face. Therefore, a balance exists between guiding the light and spreading the light and as such, a pixilated crystal will typically not be cut through the thickness of the crystal but only to a certain depth and then the remaining crystal is left as a solid piece.

Additionally, when an event occurs near an edge of a crystal there is potential for some light to escape. To counter this crystals are typically covered in a reflector on all (but at the photomultiplier interface). Often the reflector is a metal-oxide powder or Teflon which reflects the light back into the crystal.

2.2.3 Photosensors

With the exception of detector materials like CZT (cadmium zinc telluride), most detection systems require a photosensor to collect the crystal's emission light and convert it to a pulse of electrons. Two major devices exist today, *Photomultiplier tubes* (PMT's) and *Avalanche photodiodes* (APD's). APD's can, in a black-box manner, be thought of as the solid state version of the PMT.

There are some significant performance differences between APD's and PMT's. Photomultiplier tubes operate at gains on the order of one

million, whereas avalanche photodiodes operate at gains in the region of about one hundred. A second and equally significant difference is in the construction materials. PMT's do not have a negative magnetic susceptibility and as such will distort local magnetic fields and their own operations are negatively affected by magnetic fields. The APD's, being solid state, are capable of operating in magnetic fields and are able to be constructed with materials that will not negatively affect a magnetic field. The majority of clinical cameras in operation today are operating using photomultiplier tubes however it is reasonable to believe that within a number of years these will be replaced with the quickly improving solid-state system.

Photomultiplier Tubes

The photomultiplier tube has been the workhorse of nuclear imaging since they began to replace film systems with electronic systems. The PMT was first invented well over 50 years ago and although many different variations are available today, they all operate on the same principal described here. *Note: There exist a large number of variations in geometry and implementation of components so for simplicity the photomultiplier described here will be of typical design.* As shown in Figure 2.7, nearly the entire device is housed in a glass tube which is evacuated and sealed.




Figure 2.7: Typical Photomultiplier Tube design (Source: Wikimedia Commons)

The entrance window is made from a glass that is transparent to light in the spectrum of scintillator emissions. Once through the entrance window, the optical photon produced in the scintillator will interact with the photocathode via the photoelectric effect and generate photoelectrons which will transmit through the photocathode and accelerate towards the first dynode. When the photoelectron crashes into the first dynode it

liberates a number of electrons in the metal dynode. The dynode stages are all electrically charged with increasingly higher positive voltages (up to thousands of volts) so that the liberated electrons are accelerated towards the next dynode. Each dynode stage is positioned so as to best direct the ejected electrons towards the next dynode. As the electrons cascade down the PMT they each eject more electrons at every dynode they crash into. The result is that at the end of the PMT there is a larger multiplication of photo current than would be possible directly from a photocathode alone. When the multiplied electrons eventually hit the anode, they create a sharp drop in the anode voltage. Typically the anode signal is AC coupled, via a capacitor, to remove the DC bias on the anode, leaving only the sharp drop in voltage which returns, relatively slowly, to baseline.

In an arrangement that finds the anode at the front and the cathode at the end of the tube, the result is a sharp rise in signal instead of a sharp drop. Typical gains for a photomultiplier tube are on the order of a million times, however this gain can be reasonably adjusted simply by adjusting the bias voltage applied across the dynodes. Additional gain matching between individual PMT's is often done by adjusting a potentiometer (adjustable resistor) in the dynode resistor network.

Quantum efficiency (QE) is a figure of merit used to indicate the sensitivity of a photosensor. It is an accurate measurement of a photosensor's ability to create electron-hole pairs from photons hitting the active sensor. The values are wavelength specific and often measured over a range of wavelengths when characterizing a photosensors efficiency.

In a photomultiplier tube gamma camera, quantum efficiency is influenced primarily by the photocathode properties, namely its reflectivity, photoelectric cross-section, and its thickness, but also the transmission properties of the entrance window is critical [1]. Typical quantum efficiencies for a photocathode range from 20 to 30%, but recent developments produce expensive PMT's with as high as 50% quantum efficiency [1].

To get a perspective on how the signal propagates, consider a single 140 keV photon, which produces on average in the vicinity of 6000 light photons. Not all of these photons will exit the crystal and find the photocathode. With the sum of losses in the crystal and in transport to the

photocathode, perhaps half[†] of the photons will hit the photocathode. From the remaining 3000 photons, the photocathode might produce 600 photoelectrons. However, these electrons are then multiplied across the dynode stages producing a pulse of 600 million electrons or a charge of 9.612×10^{-11} coulombs (or ≈ 96 pC). This amount of charge is quite small and requires very sensitive amplifiers to boost the signal.

Photomultiplier tubes are not designed to operate in magnetic fields. The photomultiplier tube requires electrons to travel freely through a vacuum. When a moving electron is exposed to a magnetic field, as in Figure 2.8, it will experience a force on it which is perpendicular to both its velocity and the magnetic field vector. The consequence of this force in a photomultiplier tube is disastrous; the electrons no longer reach their carefully positioned targets and instead are pulled off course, in whatever direction the magnetic field dictates. As with any amplification system, the signal is most sensitive prior to amplification. In the case of PMT's this is between the photocathode and the first dynode.

Despite all of the losses and inefficiencies that propagate along the way, the PMT is still able to offer a very desirable signal to noise ratio - but when in a magnetic field, the losses are so great that the device is rendered useless.

No Field



Magnetic field into page




Figure 2.8: Left: In zero magnetic field the electron experiences no extra force. Right: Force on electron in magnetic field.

[†] Assuming the ideal situation of complete deposition of energy at (the sufficiently large) crystal-pmt interface without any intra-crystal reflection, half of 4π will be incident on the PMT - the other half will be directed away from the detector.

Avalanche Photodiode

As stated earlier in this section, on the surface, avalanche photodiodes operate much like a photomultiplier tube. That is, they convert a low energy photon into electrons, those electrons are multiplied and the result is an electrical pulse with a height proportional to the energy of the incident light that can then be manipulated and measured with analog and/or digital circuitry. In actuality the internal mechanisms for multiplications are slightly different between APD's and PMT's.

A typical avalanche photodiode cross-section is presented in figure 2.10. An avalanche photodiode is a semiconductor device. A semiconductor is primarily made up of two types of materials (N and P) both of which are commonly composed mostly of a Silicon (*Si*) crystal structure. In a *Si* crystal structure there are four neighbours that a *Si* atom shares valence electrons with - creating strong bonds and leaving no valence electrons unbound. A pure *Si* crystal structure is considered to be an effective electrical insulator. However, the electrical properties of this crystal structure can be modified by replacing a *Si* atom with a different element. This process of replacing of the atoms in a crystal structure is referred to as *doping*.

In the case of an N-type material, a (*Si*) atom would be replaced with an element which has extra valence electrons. If that *Si* atom was replaced with a Phosphorous (*P*) atom, while leaving the neighbours as *Si* atoms, then since phosphorous has five valence electrons, it would use four electrons to bond to the *Si* neighbours and would have a single electron (*donor*) left unbound (still loosely bound to the atom though).

In P-type material, a (*Si*) atom is replaced with an element which has less than four valence electrons such as Boron (*B*). In this case, there is one space (*hole*) still left empty which can be filled by a free electron.

Both types of doping, (N & P) result in increased electric conductivity of the material. In fact only a very small amount of doping will result in very drastic changes in the material properties. Both *donors* and *holes* can propagate through a material, however, because of the nature of the *holes*, they tend to propagate slower than *donors*. In the N-type material the donor electrons can easily be excited into the conduction band and in the P-type material electrons in the valence band can be excited into the band gap, mobilizing holes in the valence band.

Although it is interesting to increase the conductivity of a material by adding small amounts of another element, more interesting is what happens when a P-type and N-type layer are placed side-by-side. The junction that forms due to neighbouring, contrasting band gaps is termed a P-N junction. An area at the junction forms, termed the depletion region, where the electrons from the N-type and the holes from the P-type join. This depletion region forms naturally as the holes are attracted to the electrons locally. However, beyond that region in either direction, the original band gaps are maintained (see figure 2.9).

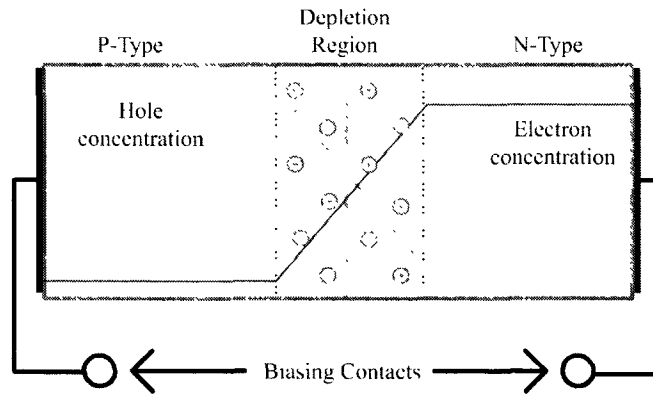


Figure 2.9: P-N Junction

Consider a P-N junction semiconductor exposed to an external voltage (applied across the biasing contacts as in figure 2.9), for instance from a battery, such that a positive voltage was experienced on the P-side and negative voltage on the N-side. In this situation the donor electrons would move away from the contact point towards the junction and likewise the holes would propagate away from its contact towards the junction as well. When the electrons and holes meet at the junction (depletion region) they will combine, thus allowing more holes and electrons to move into place. The result is that a current loop is formed and the battery is able to drain its potential.

To contrast to this, if the junction was insulated such that the electrons could not reach the holes and thus could not combine, then eventually enough electrons would pile up such that an electron density is reached that equals the potential of the battery (the electrons repel each

other - increasing electron density results in the electrons pushing back against the battery potential) and a steady state would be reached after a short flow of current. The same would happen with the holes on the other side of the insulator.

Now imagine that the bias voltage is reversed (*reverse bias*). The electrons are pulled away from the junction toward the contact and the holes do the same thing. The result here is that a transient current again, but ultimately it widens the depletion region which greatly increases the potential barrier thus acting like the contrasting example above. With the widened depletion region there is no current flow, however due to thermal effects a small current will 'leak' through. Another point of interest is, if the reverse bias across a P-N junction increases beyond a certain threshold, the leakage current can suddenly become so significant that it begins to behave like a short circuit. This point at which this sharp increase in current occurs is termed the breakdown voltage.

In an APD, as depicted in figure 2.10, the AR-layer is an anti-reflective layer that improves the absorption of light into the APD. Below the AR-layer is a semi-conductor device with P and N layers just as described above, as well as a depletion region. Also labeled in the figure is the anode and cathode across which the bias voltage is applied, but also it is these connections through which the signal is measured, super-positioned on the DC bias voltage (this is why a reliable and noise-free bias voltage supply is essential for the operation of an APD system).




Figure 2.10: Typical Avalanche Photodiode design

Typical APD's require a lower bias voltage in comparison to

PMT's, usually on the order of a few hundred volts, but because of its small cross-section this works out to an electric field of several volts per micrometer. In an APD the bias voltage is applied in the reverse direction, so ideally no current is initially flowing. However, when a photon enters the window and deposits energy into the device, it can raise the energy of a single electron (in the simplest case) which is then pushed into the conduction band, allowing it to move through the device under the influence of the bias voltage. As this electron accelerates through the multiplication region (assuming its energy is high enough) it will create more electron-hole pairs, via a process called 'impact ionization', essentially knocking electrons out of their bound state as it runs into them [1]. The gain of an APD then is determined by how many electron-hole pairs can be created from a single electron elevated into the conduction band.

When an APD is in the dark, it would *ideally* not allow any current to be generated. However, just as in a simple P-N junction, when the APD is reverse biased it is subject to leakage, which is generally referred to as *dark current* (I_D) when no light is incident on the window. Statistical fluctuations in this dark current adds to the noise and is partially affected by the gain (M), such that increasing gain also increases dark current noise. Total dark current is made up of surface leakage current and bulk current, only the bulk current is multiplied by the gain factor as shown in Eq.2.1, because it flows through the multiplication region where as surface leakage does not [5].

$$I_D = I_{DS} + I_{DB} \cdot M \quad (2.1)$$

where, I_D is the total dark current, I_{DS} is the surface leakage current, I_{DB} is the bulk leakage current and M is the gain. Additional noise appears in the statistical nature of the multiplication process: this is referred to as the *excess noise factor* (F) as described by equation 2.2 [6].

$$F = Mk + (2 - \frac{1}{M})(1 - k)... \quad (2.2)$$

where k is derived from the ratio of the probability for electrons to generate further electron-hole pairs to the probability for holes to do the same [1]. An APD in the dark presents a total spectral noise current describes as in Eq.2.3 [5].

$$i_n = \sqrt{2 \cdot q \cdot (I_{DS} + I_{DB} \cdot M^2 \cdot F) \cdot B} \quad (2.3)$$

where i_n is the total spectral noise, q is electron charge, and B is the bandwidth of the system.

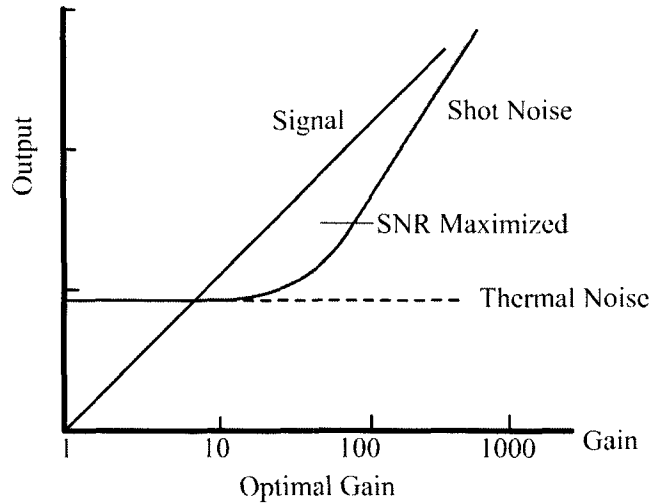


Figure 2.11: APD - Signal and Noise vs. Gain

As a result of this noise and the nature of the operation of the APD, in an ideal situation an optimal gain can be determined.

The thermal noise represents a baseline noise that must be overcome to achieve a positive S/N ratio, beyond this point exists a 'sweet spot' in which the signal gain increases much faster than the shot noise (noise due to random fluctuations in current, resulting from the fact that electrons have discrete units of charge). However, beyond this 'sweet spot', at a certain gain the shot noise will begin to increase at a rate greater than the gain can overcome. This is depicted in figure 2.11, where the separation between the output signal and the shot-noise is largest around the optimal gain. Of course, because SNR is a ratio and not just a separation, reducing thermal noise can significantly improve the signal to noise ratio.

Critical to an APD's operation are its gain characteristic, its shot noise and its thermal noise. However, in an imaging application setting, other properties need to be factored into the assessment. Of primary importance is the active area of the APD. In this realm, APDs are no different from PMTs in that the ideal design will offer 100% active area. In reality, there is an overhead amount of 'dead' surface space required in the fabrication of the APD such that 100% active area is (for now at least)

unattainable. However, by increasing the overall size of the APD, while maintaining the minimum overhead ‘dead’ area, it is possible to improve the active area percentage. There is always, of course, a tradeoff; the tradeoff for increasing the APD’s active area is the increase in noise. However, in an imaging application, the counterpoint is that more light can be collected in a single APD thus improving the signal quality. Which factor outweighs the other is dependant on the details of the system.

2.2.4 Amplification and Shaping

Preamplifier

Even in PMT systems some amplification is required, but it is especially important in APD systems where the signal is significantly smaller. The amplifier is ideally placed as near to the optical sensing device as possible. Unfortunately, the signal is most prone to noise and interference at this point.

Without a reliable way to remove the noise, amplifying the signal will also amplify the noise. By amplifying the signal before noise is introduced, the amplitude of the signal becomes much larger than the amplitude of the noise, so if noise is introduced at a later stage it will be much smaller with respect to the signal. A preamplifier will first convert the current pulse generated by the optical sensor device into a voltage pulse (shown in figure 2.12(a)) which will then be amplified in the second stage of the preamplifier device.




Figure 2.12: a) Preamplifier Pulse Profile, b) Pulse Pile-up

Gaussian Pulse Shaper

The signal that is output from the preamplifier has a similar profile to the pulse that entered the preamplifier, in figure 2.12(a). Although the signal is larger now, it is still difficult to reliably sample the pulse height without limiting the maximum detection count rate. Notice the pulse starts with a very sharp rise and a tail that slowly returns to ground (baseline). This slow 'baseline restoration' would limit the count rate significantly if we needed to wait for this to reset before we could measure again - or events would pile up, as in figure 2.12(b). In reality the height of the second pulse in figure 2.12(b) should be measured from the point it starts to rise to the peak, not from the baseline to the peak, as it has to be measured (technically). So events that pile up can introduce another issue with measuring this sort of signal.

A pulse shaper solves these problems by creating a Gaussian shaped pulse that has a height that is proportional to the height of the original pulse and returns to ground quickly after. This is depicted in figure 2.13; the rise is not quite as sharp as the original, but it quickly returns to baseline, allowing a second pulse to be accurately measured.




Figure 2.13: Gaussian Shaper Pulse Profile

2.2.5 Multi-channel Analyzer

Initial gamma-camera designs utilized a resistive charge division readout design to extract position information. In these systems the data from the array of photomultipliers was reduced to sums of rows and

columns which were passed through a network of resistors producing four signals, X+, X-, Y+ and Y-. These signals were initially used to control the electron gun on a type of cathode-ray tube, thus displaying the image on a screen. Some of these old systems have been brought into the 'digital age' by intercepting the four signals and digitizing them on a computer to collect images. A typical modified analog system might look something like the schematic presented in Figure 2.14.

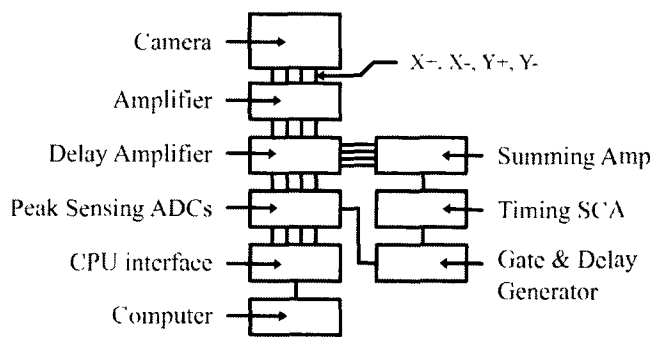


Figure 2.14: Typical Analog MCA Design

Modern systems are moving towards 'parallel readout' design. In this design each photomultiplier signal is digitized almost immediately (pre-amplification and shaping may be required) and a computer is used to calculate the position of the events. Each signal is filtered, shaped and digitized with a multi-channel analog to digital converter. The summed signals are used to trigger the ADC sampling process just as in the analog design.

Analog to Digital Conversion

Aside from variations in where the signal is converted from analog to digital, there is also variability in how it is converted; that is the operation and capability of the ADC itself, which will be discussed here.

1. Resolution and Accuracy

ADC resolution is ultimately a result of dividing the acceptable analog input voltage range (Full Scale Range) by the number of discrete values the ADC is capable of producing. Typically this is expressed as

the number of bits it works with. For instance an 8-bit ADC will divide the full scale range into 2^8 or 256 discrete values. In general the resolution of an ADC is described by the following relationship.

$$R = \frac{V_{FSR}}{2^M - 1}$$

Where R is the resolution in volts, V_{FSR} is the full scale range in volts and M is the number bits the ADC utilizes.

In practice, however, a number of other factors influence the resolution and accuracy of an ADC. The signal to noise ratio and the non-linearity of the ADC both have a strong influence on the resolution of a converter. The SNR is primarily dependent on the signal in question, the condition of the circuitry and exposure to external interferences. While SNR is primarily an external effect, the non-linearity is inherent in the converter. Non-linearity occurs (in successive approximation ADCs) when the converted values are not equally spaced across the full scale range. An ideal ADC would have a linear transfer function: resulting in a unique output for every input (within the resolution of the system).

To understand the effects of non-linearity let us consider first a linear 2-bit ADC operating at a full scale range of 4V. The bit values available range from 0 to 3, that is, there are four possible digital values available to describe the input voltage. In this case the ranges would be as follows.

Table 2.1: Non-linearity Example

Result	(MSB LSB)	Range
0	(00)	0V - 1V
1	(01)	1V - 2V
2	(10)	2V - 3V
3	(11)	3V - 4V

Notice that the least significant bit (LSB) adds a value of 0 volts when it is false and a value of 1 volt when it true, where as the most significant bit adds nothing when it is false and a value of 2 volts when it is true. Non linearity occurs when the threshold

between a bit being on or off is not centered in its respective voltage range. In the case of least significant bit the consequence of the threshold tending towards the upper end is an increased conversion of falses than trues across that specific range. This error results in high frequency oscillation in the transfer function, which directly affects a collected spectrum in the same manner. More significant the bits present lower frequency distortions, but at larger amplitudes. Of course, the further the threshold is from equally dividing the segment the further the transfer function deviates from linear.

As shown in figure 2.15(a), when the least significant bit is not shared equally across its voltage range there is high frequency oscillations and when only the upper bit is distorted (figure 2.15(b)) significantly less (20% offset vs. 80% offset) than in the least significant bit case, there is a much larger error, but the oscillations occur at a lower frequency. As expected, when multiple bits are offset then the resulting transfer function is the superposition of the individual effects (figure 2.15(c)).

2. *Sampling Rate*

There is always a finite amount of time required to convert a signal to a digital value. In application, it's important to also take into account the time required to store the data, if storage is required. The time required to convert a signal depends on the type of ADC and the aptitude of the design and hardware of the ADC.

3. *Analog to Digital Algorithm Designs*

Some common designs for analog to digital converters include Flash ADCs and Successive-Approximation ADCs.

Flash ADCs use a voltage ladder and a network of comparators to quickly output the digital value in parallel nearly instantaneously. Flash ADCs are the quickest by design, limited only by the propagation delay of a comparator. However, the cost, as well as physical size, increases with increased resolution much faster than other designs.

Successive-approximation ADCs use an iterative algorithm in order to arrive at the correct value. They utilize a digital to analog converter (DAC) and a comparator. Starting by enabling the most significant

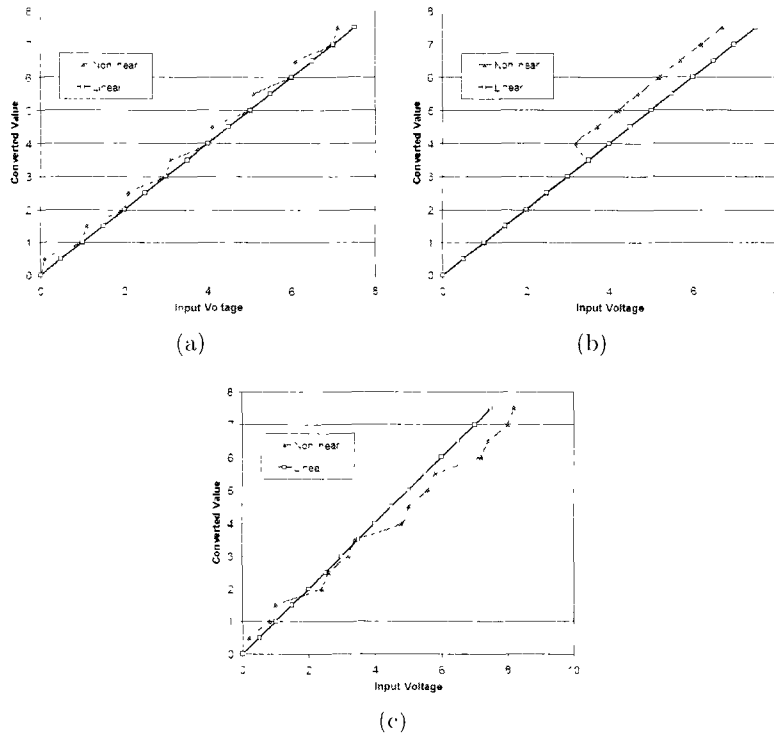


Figure 2.15: Non-linearity in Analog to Digital Conversion (a) Non-linearity in Least Significant Bit (b) Non-linearity in Most Significant Bit (c) Superposition of Non-linearity effects.

bit on the DAC, the algorithm checks if the current DAC output is greater or less than the target voltage. If the current value is greater than the target it will undo the most recently enabled bit and enable the next lowest bit. If it finds the current value is less than the target, it will simply enable the next bit. Once the change is made, the DAC output is again compared to the input. In this manner the ADC moves through all the bits; once all bits are checked, the value in the DAC register (sometimes referred to as the successive approximation register or SAR) can be read. Successive approximation ADCs are more affordable and take up less physical size than the Flash ADC. However, because each bit requires a clock cycle, increasing resolution results in increasing the time per sample.

2.2.6 Acquisition Modes

What happens once the data is digitized is completely dependent on application. The data may be collected by a microcomputer which passes the information on to a PC for more intensive processing. The microcomputer may or may not do some of the processing itself. The advantage of having the microcomputer do some processing is that the amount of data to transfer to the PC can be reduced, thus reducing transfer times. However, serial transfer speeds, with USB have reached a point that it is feasible to simply dump all collected data right to a PC (termed list-mode collection) which then takes care of all the processing.

The means of data collection currently can fall into four categories: Static, Dynamic, Gated and List-Mode. *Static* imaging is a collection of data over an extended period of time, representing a sum of the radiotracer distribution over that time. This modality is effective for studies with slow radiotracer metabolism or when the metabolism itself is not the focus, instead the distribution of radiotracer is of primary interest. Expectedly, the extended collection time lends well to low count rate applications - which can be useful for limiting patient dose in a clinical setting. A *dynamic* study is somewhat analogous to a movie; it aims to depict how the radiotracer distribution changes with time and is useful for assessing organ function. *Gated* imaging is similar to a dynamic study, but instead synchronizes the data to a patient generated rhythm such as respiratory or cardiac rhythms. This approach is useful in cardiac imaging, resulting in a series of composite images relating to the phases of the cardiac cycle. The gating process is achieved by collecting rhythm data in concert with the emission data (ex. Electrocardiogram data). Finally, *list-mode* is a catch all tactic that aims to essentially record all data in the least processed form achievable. The goal is, by collecting the data in the rawest form possible, the data can be examined using a number of techniques including, potentially, each of the above described means. The limitations for list-mode collection are first and foremost computing capacity as collecting data before it has been reduced to manageable sizes can be quite cumbersome.

2.2.7 Reconstruction & Correction

In order to generate an image from the collected data, a process called *Anger logic* is commonly employed. The general idea of Anger logic

is straight forward, but has a tendency to produce distorted images. In order to form more representative images improved algorithms have been developed, but often they rely on Anger logic as a starting point. Anger logic takes individual signals from an array of detectors and weights each signal according to its known location in reference to a central point. As an example, consider a four detector system arranged in a 2×2 square (figure 2.16), where the light from the crystal is shared amongst the four detectors.




Figure 2.16: 2×2 detector array

For simplicity consider only a single event, depositing energy at a single site. Each detector in the array will report how much energy it collected as a function of its distance from the event. A convenient approach is to allow the origin to be in the center of the image - corresponding to the center of the four detectors (where they all meet). The signal from detector 1 contributes to the $-X$ and $+Y$ direction, the signal from detector 2 $+X$ and $+Y$, detector 3 contributes to $-X$ and $-Y$ and finally detector 4 to $+X$ and $-Y$. To estimate where the event occurred in the X direction we take the difference between the $+X$ signals and the $-X$, likewise to determine the Y component. These values are then scaled by the total signal gathered by all four detectors to normalize the values. For the 2×2 case, the following equation is used.

$$X = \frac{(E_2 + E_4) - (E_1 + E_3)}{\sum_{i=1}^4 E_i}$$

where X is the coordinate of the estimated position of the event, E_i is the energy collected from detector i . Likewise for the Y coordinate,

$$Y = \frac{(E_1 + E_2) - (E_3 + E_4)}{\sum_{i=1}^4 E_i}$$

However, this mapping of energy into position results in a number of distortions. The distortions are caused by three primary factors: the *non-linear response* of position calculations, the *non-uniformity* across the field of view and variations in *detector sensitivity*. The non-linear response in the Anger logic is inherent in the Anger logic and can be corrected by using a parallel hole phantom of known grid dimensions to record an image that can be mapped to the known grid dimensions. Every image taken hereafter will undergo that mapping for correction. Secondly, to correct for the photo peak response across the field of view, a spectrum at each location is generated from the data and the photo peaks are aligned and the acceptance windows adjusted accordingly. The third factor is addressed by implementing a flood source over the entire field of view to generate a correction map which will remove the residual non-uniformities due to different photosensor efficiencies. Once these three major correction maps are generated they can be applied to any images gathered on the system to achieve reasonable distortion correction. In reality, the correction maps should be updated often to ensure any internal or environmental changes do not go unaccounted for.

2.3 Applications

2.3.1 Current Imaging Applications

Once it is possible to image radiation an entire world of applications opens up. By either utilizing an inert radiotracer or designing one that metabolizes in a desirable way, the secrets that the imaging modality can reveal are only as limited as our creativity. The gamma camera (or SPECT) has been used in imaging the thyroid, lungs, heart, liver, kidneys, brain and more. What is imaged is not just dependent on where the camera is positioned; the real fortitude derives from the radiotracer. The workhorse of radiotracers that appears in most imaging studies is Technetium-99m (^{99m}Tc), with its relatively low gamma ray energy (140 keV) and reasonably short half-life (6.02 hrs). Technetium-99m, is readily attached to many specially designed tracers used to target specific areas defined by form or function.

Dual Modality

More recently, combinations of imaging modalities have become increasingly popular for their ability to merge complimentary

benefits [13–17]. Dual modality has been primarily comprised of the addition of a CT (computed tomography) to complement a SPECT or PET scan. The CT provides a density map and the utility of this is two-fold: first, in the reconstruction process, it allows for improved attenuation correction algorithms, but secondly, the CT image can be layered with the SPECT image (for instance) providing an anatomical framework for viewing the radiotracer functional image.

New developments in specific hardware (namely solid state photomultipliers) have made it feasible to attempt dual modality imaging that replaces the CT density image with MRI. Due to its limited ability to provide soft tissue contrast, a CT image has been limited in its utility. However, by overcoming the traditional gamma camera's sensitivity to magnetic fields, it has become possible to take advantage of the strengths that MRI imaging offers and combine this with the functional imaging that gamma-ray imaging offers.

A series of solid state breakthroughs have made it possible to improve the robustness of gamma cameras (and thus SPECT cameras). With the replacement of photomultiplier tubes with solid state devices (i.e. avalanche photodiodes or CZT detectors) it is anticipated that gamma cameras will become significantly more reliable in magnetic environments.

Avalanche photodiodes have greatly improved in the last few decades, positioning themselves to potentially take over the market strongly held by photomultiplier tubes for over half a century. APD's began to make their way into the field of medical imaging in the mid-eighties and early nineties, collecting spectra in a wide energy range on a variety of scintillator types (including NaI(Tl) and CsI(Tl) and down to energies as low as 5.9 keV x-rays [21]. Since then APD's have improved from their once small active area to much larger sizes with ever decreasing dead space, they have improved their noise characteristics and moved on to modified designs that significantly improve their gain capabilities [22–26].

Naturally, this APD work advanced the interest from radiation detection to applications of imaging. The vast majority of work to date with APD's in an imaging atmosphere have been applied to positron emission tomograph (PET) for small animal imaging [27–31]. As APD's evolved further, modifications became available just as in the progression of PMT's, position sensitive APD's were soon on the market and being tested [26,32]. Due to the APD's ability to operate in magnetic fields with

which PMTs could never cope. APDs began to be tested and employed in multi-modality imaging, fusing PET with MRI [36,37]. Although experiments with PET/MRI imaging had been conducted earlier using fibre-optics to separate the PMTs from the magnetic fields of the MRI [33–35], avalanche photodiodes offer a more compact design, removing the awkward and expensive fibre-optics.

Typically, the application of APD's in medical imaging has been clearly focused on PET/MRI imaging systems [27–31]. In PET, two $511keV$ gamma rays are detected. This high energy gamma ray results in stronger signals produced from the scintillator crystal. In SPECT imaging energies are often significantly weaker (often 100 - 300 keV). This leads to weaker signals which are more susceptible to signal degradation from noise and interference. The favourable signals from a PET system may be why many projects have focused on PET/MRI.

The objective of this project was to develop a gamma camera prototype that operates with an acquisition system designed for general purpose applications. Specifically, the purpose is to design an APD-based camera to operate in the vicinity of an MRI camera. It must be able to adapt to a range of signal pulse heights and operate at sufficiently rapid count rates while still remaining cost effective.

2.3.2 Project Application

The required application for dual-modality (MRI-SPECT) imaging is to image in detail the structural details of the metacarpophalangeal joint (knuckle) obtained with MRI together with the functional information derived from the radiotracer distribution. Two radiotracers have been proposed for this application, with two complementary utilities. A ^{99m}Tc labeled MDP molecule will preferentially collect at areas of increased bone turnover. Using ^{111}In labeled white blood cells, the inflammation in the metacarpophalangeal joint can be highlighted. The ability to observe the distribution of these two radiotracers while overlaid with the detailed soft tissue contrast obtained with MRI should make possible further investigation of the origins of rheumatoid arthritis.

Rheumatoid Arthritis

The fusing of SPECT images with MRI images permits one the ability to investigate the locations of increased uptake of radiotracers while providing high-resolution structural information of soft-tissues. One immediately recognizable application of this type of facility is the study of arthritis and other rheumatic conditions. Arthritis and its associated ailments have a dramatic effect on our population. Almost 4 million, approximately 1 in 6, Canadians over the age of 15 are affected by arthritis and other rheumatic conditions [38]. Of those 4 million Canadians, two thirds are female and well over half are under 65 years of age [38]. These numbers are increasing at a significant rate – it is estimated that by 2026 the number will grow by 50%, from 4 million to 6 million [38]. Those who suffer arthritis have been found to experience more pain, activity restrictions and were more likely to require aid in daily activities and also report worse self-rated health, more disrupted sleep and depression than those with other chronic conditions [38].

The most recent studies estimate the economic burden of arthritis to be over 4.4 billion dollars [38]. Of this amount, 20% are direct costs while the remaining 80% is classified as indirect costs [38]. Indirect costs refer to lost production and mortality costs while the direct costs include treatment, care and rehabilitation expenses. An American study in 2003 found the total cost incurred by both arthritis and other rheumatic conditions totaled 128 billion dollars, well above the 86 billion dollar expense in 1997 [39].

The last decade has been marked with rapid advances in research and new highly effective drugs for treating Rheumatoid Arthritis [40]. It is no longer viewed as a chronic disease resistant to medical therapy as it once was, but as an acute or sub-acute disease that is manageable with early intervention.

Early disease modifying anti-rheumatic drugs (DMARD) have shown promise in reducing the signs and symptoms at all stages of RA. Studies with DMARDs have indicated that early detection is valuable in providing effective treatment [40]. Studies have shown that even a delay of a few months in the application of DMARDs can negatively affect the long term effectiveness [40].

2.3.3 Future Applications

Recent developments in other fields have brought to light the potential application of a magnetically robust SPECT camera in cancer treatment. The technology is still in its infancy, however, the potential exists to use magnetic nanoparticles to deliver radiation doses to tumors with controlled precision. Consider a magnetic carbon nanoparticle which contains two radioactive substances. The first is designed to treat the tumor (perhaps an alpha emitter), the second is selected for its ability to be imaged. The particles could be injected into the patient and manipulated into place near the tumor using magnetic fields. To ensure accurate placement of the particles it would be ideal to be able to detect their distribution in the body; this may be best achieved with a gamma-camera that could operate reliably in a strong magnetic environment.

Chapter 3

Radiation Detection based on Avalanche Photodiodes: Results

3.1 Scintillation Crystal: CsI

3.1.1 Specifications

The crystal selected for this prototype is a CsI(Tl) crystal. It was primarily selected because the emission wavelength ($\sim 540nm$) matches well with the optimal absorption wavelength of the APD's that were desired in this design. However, it does possess other desirable properties. It has a reasonably high effective Z value (54) and a density of $4.51 g/cm^3$. The decay time of CsI(Tl) however is much slower than would be ideal; at 1000ns, this kind of decay time limits the maximum count rate possible by a detector system. In table 2.6 CsI(Tl) is presented in comparison to some other common detector crystals.

The spectral output of CsI(Tl) peaks just under $550nm$ as shown in figure 3.1, which if compared to the APD response for the S8664-1010 model in figure 3.2 (peaking near $600nm$) is a very acceptable match.

In order to minimize light loss between the scintillation crystal and the avalanche photodiode through reflection, coupling grease is sandwiched between the two. It is important that the grease is free of air bubbles because the bubbles would cause reflection and scatter.

The scintillation crystal available was a single solid block of




Figure 3.1: Emission Spectrum from CsI(Tl) and CsI(Na) [8]

CsI(Tl). 6mm thick. The block was too large for a 2×2 testing array, and was cut down to fit. However, due to less than ideal cutting equipment, the cut edges most likely contributed to light loss. It is wrapped in Teflon tape to reflect scintillation light back towards the crystal. Fortunately, CsI(Tl) offers almost twice as much light output than NaI(Tl).

3.2 Selection of APD

The avalanche photodiode selected for this prototype is a Hamamatsu S8664-1010 (see Figure 3.3). It offers an effective active area of 10 mm by 10 mm, but more importantly the ratio of inactive area (*dead space*) to total area is greatly decreased from smaller models. The S8664-1010 has a dead space percentage of just less than 50% ($49.7\% \pm 6.3\%$); that is to say, there is more active detector than dead space. In a similarly designed model of 5 mm by 5 mm active area, the same dead space ratio increases to 73.8 ± 9.2 [7]. Also available are metal can models, but these too have less desirable active area ratios. Increased dead space leads to lost light from the scintillator. The more light collected by an APD from a scintillation event the better the signal that can be expected. Due to APD's low gains (at the time of purchase) every photon of light that can be collected is important, thus selecting an APD with the lowest dead space will act to somewhat counter the negative effects of low gain.

The spectral response for the S8664-1010 peaks just below 600nm with a nominal full width half max of 500 ± 20 nm as shown in figure 3.2. To maximize energy collection, it is desirable to select a scintillator with an emission wavelength that best matches the spectral response peak of the




Figure 3.2: S8644-1010 Response Spectrum [7]

photo multiplier. Upon review of crystal properties in figure 2.6 it can be seen that CsI(Tl) with an emission peak at 540nm is best positioned to operate with these APD's.




Figure 3.3: Active area of S8664-1010 [7]

The breakdown voltage for this APD is quoted at 400 V [7], but in our experience, it is more often above 400V than it is below, but also never exceeding 500 V . Typical dark currents (I_D) are claimed to be between 10 and 100 nA [7]. In our tests the dark currents were found to be larger than those claimed by Hamamatsu. The source of this discrepancy is believed to be caused by our inferior testing conditions. However, these conditions perhaps more accurately represent the operating conditions of the ultimate gamma camera design. Test results for APD photo current, dark current,

gain and signal to noise ratios are presented in Section 5

Table 3.1: APD Specifications [7]

Type	Spectral Response Range (nm)	Peak Sensitivity (nm)	Quantum Efficiency * (%)	Typ Break-down Voltage (V)	Max Break-down Voltage (V)	Typ Dark Current (nA)	Max Dark Current (nA)	Terminal Capacitance (pF)
S8664 10x10	320 to 1000	600	70	400	500	10	100	270
S8664 5x5	320 to 1000	600	70	400	500	5	50	80

It is possible to calculate the expected size of a current pulse from a single event if the absorbed energy (E_{abs}), detector efficiency (ϵ) and gain (G) are all known, as in equation 3.1 below. Refer to appendix A for a detailed sample calculation using this system's specifications.

$$Q = E_{abs} \epsilon G \quad (3.1)$$

Care should be taken when choosing an approach for integrating devices. Each connected device will affect the components it is connected to, and each component will have some effect on it too. In most cases, components are designed to minimize these effects, however, occasionally this is not feasible and special attention must be paid to how components will interact.

In the case of the APD, its high voltage bias supply resides on the same contact that provides the signal. The consequence is that the APD provides a small signal superimposed on a large DC voltage. If the large DC voltage was not removed, the subsequent preamplifier would be overloaded and its output would be clipped at best. Instead, the APD is connected via a small capacitor to filter out the DC component while allowing nearly all the variable voltage (ie. signal) to pass through. This capacitor is currently included in the prototype board provided by Cremat, however, when a customized board is developed, the capacitance matching between APD and the board should be considered to reduce signal degradation. With this filter in place the preamplifier only sees the small changes in the signal from the APD and is able to provide significant amplification of this signal while still avoiding any signal clipping.

The APD must also be shielded from external interferences. Due to its relatively small signals, it is important to block as much interference as possible before the preamplifier stage, otherwise the interference will be amplified along with the signal. The APD is sensitive to light, of course, and thus should be shielded from all external sources of light other than that provided by the scintillator crystal. Also, the APD is very capable of picking up any stray radio frequencies (RF), so it is also important to keep the APD in a grounded conductive box that will shield it from this RF interference. However, by placing the APD in a grounded box it is possible to create a ground loop which can allow AC power supply frequencies to interfere with the signal. The prototype avoids these ground loops by placing each of the front-end instrumentation into separate RF shielded containers and connecting them via insulated co-axial cables. The cables provide the connection to ground on a common line for all front-end devices.

The APD is reverse biased using a linear power supply. However, the preamplifiers, shapers and digital electronics are powered by a switching power supply which is in close proximity. Switching power supplies are known to produce RF in the region above 50 kHz [18], which has been detected in the circuitry and measures have been taken to minimize it. Unfortunately, some of this RF can propagate down output lines, resulting in potential interference issues in the front-end instrumentation.

Terminal capacitance is the summation of series capacitance between the APD to the preamplifier. In the case of this design, this capacitance is made up of the capacitance from the APD itself, any capacitance formed by the contacts in the BNC connectors for the co-axial cables, and the capacitance on the preamplifier's circuit board. This capacitance is proportional to the rise-time of the pulse. It should be noted that, if the rise-time is limited by the detector, then reducing this capacitance will not improve the overall rise-time. However, if the rise-time of the detector is not the limiting factor, then attempts to reduce this terminal capacitance can result in improved rise-times, which can result in improved maximum count rates.

3.2.1 APD Dark Current (I_d) vs. Bias

To measure dark current and photo current the APD was enclosed in a metal box having an LED embedded in a hole in the side. The box was covered with felt to absorb any external light. The APD was directly biased

by the linear power supply voltage, with a large resistor ($120M\Omega$) in series between the APD and ground, across which a voltage differential was measured on an oscilloscope and an average value recorded. This resistor is essential to both the protection of the APD and the measurement of the current, however, it is also inherently a source of noise which contributes to the uncertainty in the plots. The voltage drop across the resistor was measured with the LED both 'on' and 'off' to simulate *dark* and *photo* current responses. The current value is deduced from the measured voltage and the known resistor value.




Figure 3.4: APD Photo and Dark currents: APD A has a breakdown voltage of 489 V and APD B has a breakdown voltage of 448.

Figure 3.4 is a comparison of dark current (I_D) and photo current (I_p) between two S8644-1010 APD's (see Table 3.2). The first APD (A) has a specified breakdown voltage (V_b) of 489 V and a dark current (I_d) of 30 nA as measured by Hamamatsu. The second APD (B) has a breakdown voltage measured at 448 V and dark current of 32 nA quoted by Hamamatsu. The increased breakdown voltage of APD (A) is exemplified by the curve shifted to the right along the bias voltage axis.

The gain of the APD is dependent on the bias voltage. The gains can reach as high as 100 when the bias is pushed to the edge of breakdown but are reliably run around 50 times gain. Testing our own APD's with a red (actual emission spectrum unknown) LED showed similar but slightly lower gains.

Table 3.2: APD A and B Specifications

APD ID	Breakdown Voltage	Dark Current [†]
APD A	489 V	30 nA
APD B	448 V	32 nA

[†] Dark Current as specified by manufacturer at $M = 50$

3.2.2 APD Gain vs. Bias

The slight variations in structure that exist between each APD lead to variation in breakdown voltages. The same two APD's as previously used above (APD A and APD B). properties in Table 3.2. were studied to investigate the extent of this effect. Each APD, placed in the apparatus described previously for dark current measurements. was tested from zero bias voltage to voltages beyond their respective breakdown. At lower bias voltages samples were taken every 50 volts. but as the bias voltage neared the breakdown more finely spaced measurements were made, to a minimum of 5 volt separations between measurements. As shown in figure 3.5 and table 3.2. APD B has a breakdown voltage of 448 V. but the APD A breakdown is at a higher voltage of 489 V. This difference results in a horizontal shift in the plot. but the consequence is that at a given bias voltage the photo current for APD A will be much less than APD B - as a result. the gain for APD B will peak at a lower voltage than APD A. The gains measured with these two APD's are shown in figure 3.5 and as it can be seen. APD B with its lower breakdown voltage reaches a higher gain sooner than APD A. In fact. when APD B is at a gain of 50 times. APD A will only operate at a gain of just over 20. at that same bias voltage.

The gain curves appear similar to those by Kataoka and Sato [19,20]. The curves will not look identical. just as these two APD's produce slightly different curves, so will the other APD's used in other experiments. Despite using the same family of APD's as used in this experiment. it appears as though others were able to achieve slightly more steep gain curves. The source of this discrepancy is not known but. it may be due to variations in the characteristics of the single APD used in their experiments or due to the variations in experimental setup.

By dividing the photo-current (signal) by the dark current (noise) we get a good approximation of the signal to noise ratio for the APD. by multiplying that by the photo-current again. we will have an indicator of signal strength. Shown in figure 3.6 is a plot of this figure of merit.




Figure 3.5: Gain Curves for APD A and APD B vs. Reverse bias Voltage.

Consider selecting a bias voltage for APD B near 420 V. this would ensure that not only the clearest signal, but also strongest signal, is achieved for APD B. However, if this same voltage is applied across APD A (optimal bias voltage of 460 V), the result is an APD that is operating far below its optimal capabilities, and moreover, the gain is significantly lower. In an imaging environment, this will result not only in a strongly skewed image, but also possibly lost events. Clearly, this is not the ideal way to operate an array of APD's: a better method is to test each APD for its optimal bias voltage, and to operate each APD at its own optimal bias voltage. The current design does not attempt to do this, but doing so should produce much more desirable results.

The data presented in figure 3.6 is of arbitrary units of gain vs. bias voltage. The gain values are normalized by the values generated at a zero bias reference.

When the APD is struck with photons from a CsI crystal that is exposed to low energy gamma rays, pulses can be seen on the output of the preamplifier. The pulses appear as those in figure 2.12(a) and 2.12(b), and they have been measured to have decay times on the order of 10's of μs up to 100's of μs depending on the pulse height.

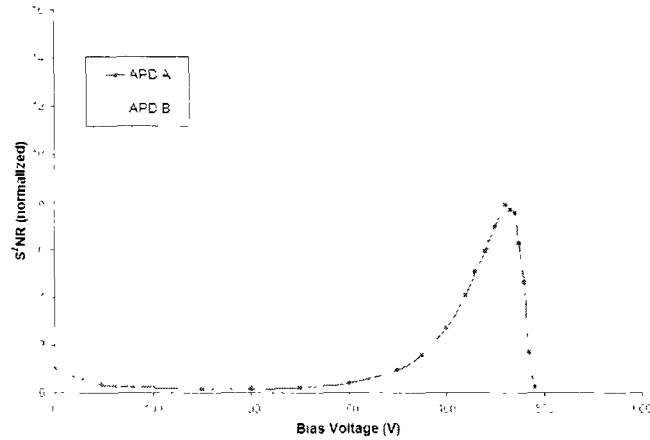


Figure 3.6: Optimal Gain Bias Voltage: APD A has a breakdown voltage of 489 V and APD B has a breakdown voltage of 448 V.

3.3 Preamplifier: CR-100

The preamplifier chosen was selected for both cost and operating specifications. The Cremat CR-110 is a single channel charge sensitive preamplifier that offers gains of 1400 mV/pC allowing the system to detect smaller pulses (at the sacrifice of count rate). The amplifier produces 200 electrons of equivalent noise charge (ENC RMS) when measured with no input signal, but this noise will increase with detector capacitance generated noise, leakage current of the APD, and dielectric losses in the electronics [9].

Given the calculation in appendix A a 140 keV pulse would produce a pulse from 86 mV at a 50 APD gain to a 168 mV pulse at a 100 APD gain.

The preamplifier offers a rise time of a minimum of 7ns, but when the detector capacitance is factored in, can be described by the equation below [9]

$$t_r = 0.4C_d + 7ns$$

Where C_d is the detector capacitance, measured in pico-farads (pf) and t_r is the pulse rise time, measured in ns. The Hamamatsu 8664-1010 has a large terminal capacitance of 270pf because of its large active area, resulting in

an estimated rise time of $115ns$. However, because of other hardware limitations (ie. CsI(Tl) decay time) this rise time is acceptable. A circuit diagram depicting the APD bias interface is shown in figure 3.7.

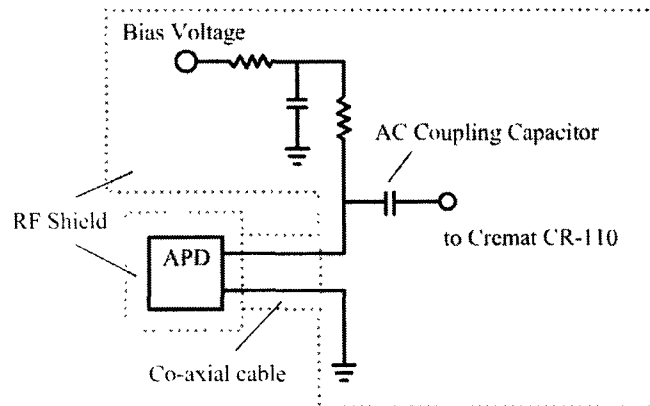


Figure 3.7: Coupling Circuit connecting APD to the Cremat preamplifier

Currently, the Cremat preamplifiers are operating on independent Cremat prototype boards (*CR – 150*) in their own shielded boxes as in Figure 3.7. The board allows the preamplifier to be easily powered with ± 12 volts and provides a large ground plane to improve noise sensitivity. The APD's are housed in a custom made RF shielding box and the two are connected with insulated co-axial cable and BNC connectors.

3.3.1 Noise Sources

According to the preamplifier manufacturer, the primary noise factors are:

1. Thermal noise from the capacitance at the JFET input.
2. Thermal noise in the preamplifiers feedback circuit (and biasing resistors).
3. Detector leakage (as discussed in Chapter 2).
4. 1/Frequency noise produced at the contact point between the detector and the preamplifier.
5. Lossy dielectric material near the preamplifier input.

The resultant output signal from the preamplifier on a long time scale appears as a step function that rises sharply and very slowly decays back towards the baseline as depicted in figure 3.8.




Figure 3.8: Preamplifier Pulse as seen by the Pulse Shaper

3.4 Pulse shaper: CR-200

Following the preamplifier, a pulse shaper processes the step function into a Gaussian pulse. The shaper circuit is made up of a differentiating circuit followed by two band-pass filters (Sallen Key filters) - a number of integration stages to shape the pulse into a Gaussian distribution [10]. In a long time scale the pulse shape from the preamplifier appears like a step function. The step function, however, does not rise instantaneously, instead it is continuous, so differentiating this signal produces a pulse which rises quickly, then reaches a maximum and quickly returns to baseline. The Sallen Key filters are used to shape this pulse into a more Gaussian like pulse and provide the 10 time gain as well. The shaper essentially removes the long tail from the preamplifier pulse so that the measured signal returns to baseline quickly and another measurement can be made sooner. The pulse shaper is able to handle closer pulses because the preamplifier signal is similar to a step function. Imagine a second pulse following closely behind the first such that the second pulse is sitting on top of the tail of the first pulse. The pulse shaper will only deal with the slope of the incoming pulse and will not care what the relative offset of the pulse is, and so the resulting output of the pulse shaper is two Gaussian pulses separated by as much time as the two incoming pulses. There is an inherent time delay between the peak of the Gaussian pulse and the rising edge of the preamplifier signal due to the processing time required, but this time delay is consistent and as such can be ignored.

The pulse shaper, also manufactured by Cremat, was the CR-200 Gaussian shaping amplifier. There are four models available, offering shaping times of 100ns, 250ns, 1.0 μ s and 4.0 μ s. The amplifier provides an inherent gain of 10 times, and any additional amplification is done prior to the shaping. In our application an additional 10 times gain is provided through the amplification stage on the CR-160 evaluation board. The width of the output pulse is related to the shaping time by a factor of 2.35 (ex. A shaping time of 1 μ s means a FWHM pulse of 2.35 μ s) [10].

In the next chapter, a custom designed acquisition system designed with 'off the shelf' parts to digitize the pulse shaper peaks is discussed at length.

Chapter 4

Multi-Channel Analyzer

4.1 Peak Detection

The pulse generated by the pulse shaper, as discussed in the gamma camera overview (section 2.2.4) passes so quickly that to time the digital to analog conversion to sample at the very peak would be difficult to do reliably. The peak detect and hold circuit utilizes the functionality of the *Analog Devices' Monolithic Peak Detector with Reset-and-Hold Mode* chip. 'Monolithic' refers to the fact that some of the components are fabricated into the substrate which provides inherent advantages for charge injection and droop rates error reduction (to be discussed later) [11]. The chip can be run in three operational modes, 'peak detect', 'peak hold' and 'reset' as described in table 4.1. In peak detect, the output will be equal to the maximum voltage experienced on the input since the last reset. In the Peak hold mode, the current peak will be held (on an external capacitor) such that any higher peaks occurring after the initial peak will be ignored. And finally the reset simply resets the charge in the 'hold' capacitor.

Table 4.1: PKD01 Operation Modes [11]

<i>RST</i>	<i>DET</i>	Operation Mode
0	0	Peak Detect
0	1	Peak Hold
1	1	Reset
1	0	NA

The 'slew rate' is measured in voltage per time unit and is essentially a measure of how quickly a voltage is able to rise or fall. In the

case of a peak detect system, a slew rate equal to or better than the slew rate of the shaper is essential for ensuring accurate voltage sampling. The smaller the capacitor (storing the charge) the faster it can charge up, that is, the faster the slew rate. Table 4.2 shows a few capacitors that were tested in the PKD01 for how quickly the voltage could be reset to baseline - similar values were tested for rise times. In figure 4.1 an input pulse shown by the thick black line rises and falls, and the dashed line represents a tracking signal, limited by a given slew rate, as it attempts to follow the original pulse. As shown in the figure, the tracking signal can match the input signal only when the slew rate of the signal is less than or equal to the limiting slew rate of the tracker, provided it was starting from a matching point prior. That is, if the two signals are already in separation, then a certain amount of time is required for the tracker to catch up - specifically determined by the signal and tracker slew rates.



Figure 4.1: Effects of limiting slew rates

Table 4.2: PKD01 Response times for various capacitors

Capacitance (pF)	Rise time (0 to 3.3V) (μs)	Fall time (3.3V to 0) (μs)
1000	6 - 8	6 - 8
330	2 - 3	2 - 3
56	~ 0.5	~ 0.5




Figure 4.2: Peak Detector: Droop rate

Droop rate is, as depicted in figure 4.2 when a pulse height is detected and held, it will slowly and exponentially lose its voltage level due

to leakage in the capacitor. Smaller capacitors are more prone to leakage, and as such will exhibit an increased droop rate. This being the case, a balance between acceptable droop rate and slew rate must be determined. Testing showed that the slew rate associated with the 56 pF capacitor was too excessive to warrant its selection. Capacitors in the 300 pF range were found to offer an acceptable balance of slew rate and droop rate. Some of the relevant specifications for the peak detect and hold chip are provided in table 4.3 as quoted by the manufacturer.

Table 4.3: PKD01 Specifications - $C_H = 1000\mu F$ [11]

Parameter	Symbol	Min	Typ	Max	Units
Zero-Scale error	V_{ZS}	-	2	4	mV
Slew Rate	SR	-	0.5	-	V/ μs
Droop Rate	V_{DR}	-	0.02	0.15	mV/ms

4.2 Event Trigger

Up to this point, the fact that an array of APD's is required for imaging has not been discussed for simplicity. However, the design of components beyond this point strongly depend on the fact that a number of detectors are operating in parallel. The signals have to be evaluated both individually and also as a single (summed) entity. When a single crystal shares the light among a number of detectors, individual signals will vary depending on the position of the event, but on average the overall energy deposited in the crystal will still be proportional to the incident energy. Therefore, to determine if a photon with an energy we are interested in has been collected we must look at the summed signal from each APD. However, because we are still interested in the individual signals in order to extract position information, it is crucial that we also maintain the integrity of individual signals while summing them.

The signals from the array of APD's are summed together before being fed into a comparator circuit which is provided on the Monolithic Peak Detector chip. However, if the signals were simply physically connected at a node, the resultant signal would be an average of the four signals, instead of a sum. This is less than perfect for setting triggering thresholds. More importantly though, each signal is spread out over the four channels averaging the individual pulse heights and losing valuable position information. Ex. In a two node connection a signal with a pulse

height of 2 (arbitrary units) and a second signal of 8 would typically position the event much closer to the second detector, however if the two are connected at a simple node for summation, the individual signals would be approximately 5 each. To avoid this problem, each channel is designed to be isolated from the summing circuit. A typical summing circuit is shown in figure 4.3.




Figure 4.3: Typical Op amp Summing Circuit

However, the summing circuit on its own is subjected to a node connection, as discussed above, for each of the channels. To avoid this, each channel is passed through a voltage follower circuit. The voltage follower acts as a buffer, providing no amplification, and because of the high impedance of op amp inputs, it prevents the channels from interfering with each other. Once each channel is buffered, the outputs are summed together using the summing circuit as shown in figure 4.4.

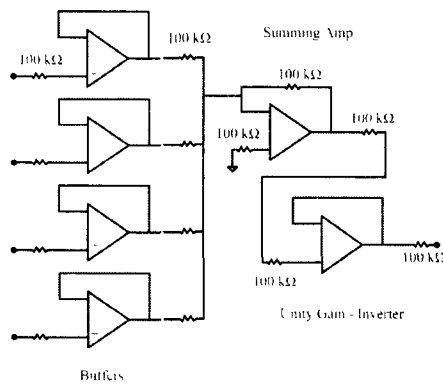


Figure 4.4: Op amp Summing Circuit with voltage follower buffers

The output of the summing circuit is inverted, however, summing it again with 0V re-inverts the voltage to the desired sum, shown in figure 4.4. This signal, finally, can be fed into the comparator. The comparator will

output either a high or low voltage depending on the two input signals. Signal one is a DC value selected as the threshold for which the second signal is compared to. When the comparator sees that the second signal (the summed signal) is greater than the first (DC) it will output a logic-high voltage on its output. The actual voltage is clamped by a resistor network, which in this case is set to be about 3.3V to match the requirements for the Digital Signal Controller (DSC - described next) input. This resistor network can be seen in the actual circuit layout included in figure B.1 in appendix B and B.3.

Once the peak voltage has been tracked and the comparator now set to logic high, it interrupts the DSC, sending it into a routine which will digitize and record the analog value on each channel, described in detail next.

4.3 Digital Signal Controller (DSC/DSP)

The DSC/DSP is possibly the most complicated piece of electronics in the entire electronic system. The design of the micro processor actually straddles the definition of both a *digital signal processor (DSP)* and a *digital signal controller (DSC)*; as such the chip may be referred to here as either a DSP or a DSC. The Freescale 56F807 DSP is a 16-bit processor running at a clock speed of 40Mhz, pictured on a prototype board in appendix D. Although it was designed for use in modern consumer devices like automobiles and washing machines, and almost anything requiring motor control, it is a very robust and affordable DSP which merges the benefits of a DSP and a microcontroller. The DSC is designed on a Harvard-style architecture containing three parallel execution units which allows up to six operations per instruction cycle [12]. The chip contains a 16 channel Analog to Digital converter, 16 programmable 16-bit timers and 32 general input/output pins. Of the 32 general input/outputs some offer a dual function including a Serial Peripheral Interface Port (SPI) and a Serial Communications Interface Port (SCI) among others, leaving 14 pins solely dedicated to general input/output.

The DSC is setup to act as a slave to the PC which ultimately collects the data, but the DSC is responsible for controlling the operation of the peak detect/hold (PKD01) circuit. As mentioned above, the PKD01 offers a few operating modes, 'peak detect', 'peak hold' and 'reset' which is

determined by the state of two pins on the chip (RST and \overline{DET}). The DSC selects the appropriate mode by setting these bits according to table 4.1 with two output pins that are hardwired to the PKD01 - thus allowing the DSC algorithm to decide when and how to collect data. Two additional pins are connected to the output of the comparator, one to determine when the voltage rises above the threshold and one to determine when it falls below again - allowing for a unique window into the state of the signal without undergoing the rigorous ADC process.

Signal Conversion

The DSP operates its external inputs and outputs on a 3.3V logic high convention, while the peak detector (PKD01) requires a 5V high convention for operation. To mediate the differences in logic protocol, a CMOS hex buffer is used which allows output from the DSC to be converted from 3.3V to 5V logic high values. A schematic of the chip integration is provided in figure B.2.

4.3.1 SCI

The Serial Communication Interface module allows for asynchronous serial communication with the data collecting PC (in this case a laptop). This is done by the serial to universal serial bus (USB) module which will be discussed later. The Serial Communication Interface has been tested to reliably transmit at speeds up to 500,000 baud. The DSC SCI module is capable of faster rates but is limited by baud matching with the serial to USB module. The SCI module sends 8-bit words, so in order to send a 16-bit word, the word is split into two chunks (upper byte and lower byte), sent separately and rejoined on the receiving side. Through this connection instructions can be sent to the running DSC from the PC and data can be sent back to the PC. How this operation is organized will be discussed in section 4.5.

4.3.2 ADC

The DSC has two dual, four-input 12-bit ADCs. That is two modules named ADCA and ADCB which each have two four-input analog to digital converters (see figure 4.5). Each of these converters has up to 12-bit resolution (it can be lowered if desirable) and can complete 8 conversions in $5.6\mu s$. However, the ADC stores results in a result buffer, and this conversion time does not include time to transfer the values from

the result buffer into other working memory, or further processing. The ADC can operate in single input mode or differential input mode - this system utilizes the single input mode to increase the number of available ADC channels, as differential mode requires two pins per channel. If the 12-bit conversion is utilized then the 3.3V range is divided into a total of 4096 different states and is stored in a 16-bit word. The 12-bits are shifted 3 bits to the left, leaving 1 most significant bit for sign and 3 trailing bits which are always zero. To convert from a known digital value to a voltage equation 4.1 should be used.

$$\frac{V_{IN}}{32768} \times V_{REF} \quad (4.1)$$

where, V_{IN} is the measured digital value, and V_{REF} is the reference voltage, in this case 3.3V. Any noise that exists on this reference voltage can propagate into the converted values, provided the noise is great enough to exceed the minimum resolution of the ADC.

The ADC scan sequence can be initiated in software by setting the START bit to true, but also can be set to trigger off an external interrupt as it is in this design. The pin designated to be the external interrupt is hardwired to the output of the comparator in the peak detect and hold circuit. Once a SCAN is initiated the ADC will convert the voltage of each of the channels listed in the list register (ADLST). The list can contain any arrangement of eight channels as long as no channel is listed twice. Additionally, to sample fewer than eight channels, samples can be disabled with the disable register (ADSDIS). Once a SCAN is initiated it can sample up to eight unique channels (four in differential mode). The ADC can be configured to perform a single scan and stop, scan whenever triggered, or scan repeatedly until manually stopped. For example, consider how to setup the ADC to convert two channels when triggered. The disable register (ADSDIS) would be loaded with all ones except the first two (0x00FC) and the list register (ADLST1) 0x0076, which would dictate the ADC to convert channel 6 first and store that in result register 0 (ADRSLT0) and then it would convert channel 7 and store it in result register 1 (ADRSLT1).

Figure 4.5 is a block diagram of the ADC module as depicted by Freescale Semiconductor Inc.. The diagram displays one of the two ADC modules; either ADCA or ADCB. The ADC input channels first go to a multiplexor (MUX) which essentially allows the ADC to select which of the eight channels is passed. Following this is an additional sample/hold mechanism to hold the voltage as it is being converted. After the

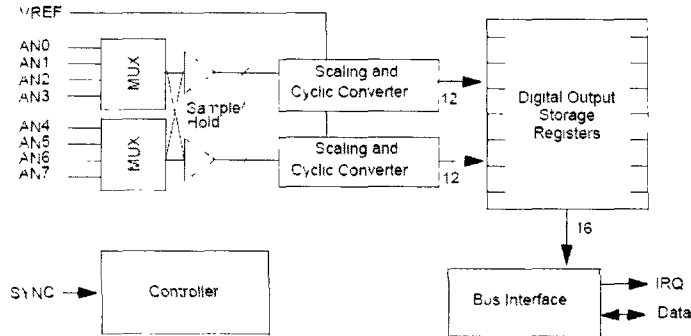


Figure 4.5: Analog to Digital Converter: Functional Block Diagram [12]

sample/hold is the ADC converter which then stores the 12-bit result in the result register as shown. This data can then be moved around via the bus interface.

To understand exactly how the ADC works in time, it is important to review a timing diagram of the ADC as shown in figure 4.6. The ADC CLK is the internal clock that runs the ADC. It is designed to run at a maximum of 5MHz. The timing diagram shows a SYNC pulse which initiates a conversion; the MUX selects the first channel from the list and begins conversion of it. The first conversion takes 8.5 ADC clock cycles (at 5MHz that's $1.7\mu\text{s}$) but each successive conversion requires only 6 clock cycles or $1.2\mu\text{s}$ during conversion operation [12]. That is, if the ADC is allowed to run with out interruption, it will continue to convert values in 6 clock cycles, but if it is stopped and restarted it will take 8.5 cycles to initiate again. Shown in the diagram, the Conversion in process (CIP) bit is set to high once the conversions start and does not switch to false until all requested conversions are complete.

4.4 Master CPU

The DSC is controlled by a laptop in a master/slave configuration, meaning the computer will initiate all actions taken by the DSC. The communication is done via a USB connection which device drivers allow to appear like a normal serial port in Matlab. The laptop is running the Windows XP operating system, with 1 Gb of ram, 60Gb hard drive and parallel port which is used to program the DSC. In Matlab, the graphic

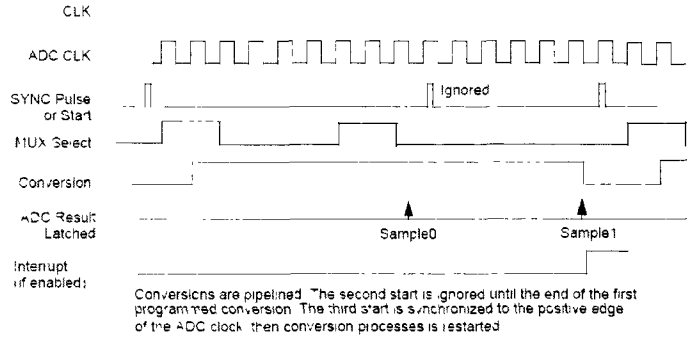


Figure 4.6: Analog to Digital Converter: Timing Diagram [12]

user interface development environment (GUIDE) was utilized to create a display interface which can display live spectra and images as data is collected. The program will collect data from the DSC for a user defined amount of time; when the time expires the data collection is halted and the program is able to save the collected data or clear and sample again.

4.5 System Operation

This section will discuss the overall operation of the system. to supplement the description, a functional diagram is provided in figure 4.9. The overall operation of this system is ultimately user driven. The user initiates a request for data collection and specifies the sampling time via a Matlab GUI based program. The program then starts a timer to keep track of sampling time and subsequently sends a code to the DSC requesting 100 samples to be sent back - this is an arbitrary number of samples and no reason exists that this could not be modified. The Matlab program will wait for the data while the DSC collects it. Due to the design of serial communication interfaces, data is (often) sent in packets of 8 bits and since the data collected on the DSC takes up at least two bytes (12-bits) then the data is sent in two successive bytes which are put back together on the PC. Once the data is received, processed and stored the Matlab program, provided the timer has not yet expired, requests another 100 samples. This continues until time expires, allowing the final transfer to complete itself before terminating the connection and returning control to the user.

Because the PC initiates the transfer of data it is considered the

Master and the DSC is considered the Slave. As a slave the DSC is designed to wait for orders from the PC. A typical process for the DSC to undergo is described in the following. When the DSC boots up it first sets up the ADC registers to operate the ADC in the desired manner and prepares some memory locations for variable storage - then it enters an infinite loop and waits. The DSC will wait until it receives instructions from the PC, and when it does it checks the instructions and performs the appropriate actions. Currently the DSC is setup to only perform a single action, that is, running the ADC and passing the data to the PC, but it is designed to accept other instructions as well. For instance, if the DSC were to control a digital to analog converter for voltage thresholds, the PC (read: user) would be able to instruct the DSC to set the threshold to the desired value. Once the instruction is interpreted the DSC takes action, in this case running the conversion algorithm. The DSC waits for the signal to drop below the threshold, thus indicating a clean starting point, and it enters a loop which will iterate 100 times, each time collecting data-points from the ADC. Currently the algorithm is designed to collect a data point from each of the four channels connected to it. However, this could be modified to accept up to 16 channels and any number of channels in between - for instance 9 signals from a 3×3 array. The first step in the iteration is to set the PKD1 to 'peak detect mode', and wait for the signal to cross the threshold, triggering the DSC to proceed to the next step. After a delay loop which ensures the complete rise of the signal has occurred, the PKD01 is set to 'peak hold mode' allowing the value to be held while the ADC processes it. The ADC is then initiated and the samples are collected and stored in an array. Finally the DSC will wait until the pulse has dropped back down below the threshold before attempting to reset the voltage on the PKD01. This is iterated until 100 data points have been collected - so in the case where four channels are collected simultaneously, then 25 iterations would occur resulting in 100 data points. Once the 100 data points are collected, the DSC enters a loop which breaks the data points up into two 8-bit words which can then be piped to the serial port which of course is received by the PC.

A better understanding of the operation and its timing may be gleaned from reviewing various critical signals as they would appear on an oscilloscope. In figure 4.7 the various signals are presented as they would be viewed while the system is in operation. Trace 1 is the signal received from the preamplifier, and the second trace is the signal from the pulse shaper output. Trace number 3 is the signal experienced on the output of the PKD01, and the two square pulses are the 'hold' signal from the DSC (4)




Figure 4.7: Analog to Digital Converter: Oscilloscope Traces

and the 'reset' signal from the DSC (5). As trace 2 rises, trace 3 tracks it as its maximum slew rate and detects the peak and holds the value there - shortly after the DSC has responded to the interrupt due to crossing the voltage threshold and outputs trace 4 which holds the voltage value from further peaks. In the time between the rise of the fourth trace and the rise of the fifth trace the ADC converts the analog voltage values to digital values and stores them. It then outputs a square pulse (trace 5) which resets the PKD01 - the square pulse is timed to ensure it is held long enough to exceed the tested fall times as presented at the beginning of this chapter.

4.6 System Resolution

To test a data acquisition system without the added variables associated with detector effects it is helpful to have a known input to work with. Fortunately, a reference pulser that operates at 50 Hz with controllable amplitude was available. A reference pulser outputs pulses similar in profile to those expected from a charge sensitive preamplifier. The pulses were fed into the Cremat Gaussian pulse shaper and subsequently passed into the acquisition system. The resultant spectrum collected would ideally be a sharp peak at the bin associated to the voltage the reference pulser is set to. Presented in figure 4.8 is such a spectrum - collected for 5 minutes with the pulse height calibrated for 2 volt peaks. The system was able to resolve the peak with a 1.1% full width half max (FWHM). This peak width is indicative of the best resolution achievable




Figure 4.8: Spectrum collected from 50Hz Reference Pulser

with this system before any geometric effects, noise, interference or other degradations are able to infiltrate the signal.

4.7 Conclusion

Each of the components discussed here and the previous chapter are connected as shown in figure 4.9. To summarize, the signal from the APD or PMT will travel to the preamplifier, which then outputs to the Gaussian pulse shaper. Each channel from the shaper is passed to the summing circuit as well as to individual peak detector circuits. The summed signal is connected to a comparator, the output of which controls an interrupt pin on the DSC. The individual output channels on the peak detect are all connected to corresponding ADC channels on the DSC. Lastly, the DSC communicates with the PC via a serial communication interface and USB bridge. The PC is able to send instructions to and receive data from the DSC, which in turn controls the subsequent circuitry.

Having discussed the design of the data acquisition system and front end electronics, the next section will begin to investigate the operation of the components and the system as a whole. The following chapter will evaluate the system in single channel operation, collecting spectra in various arrangements with a number of signal sources.

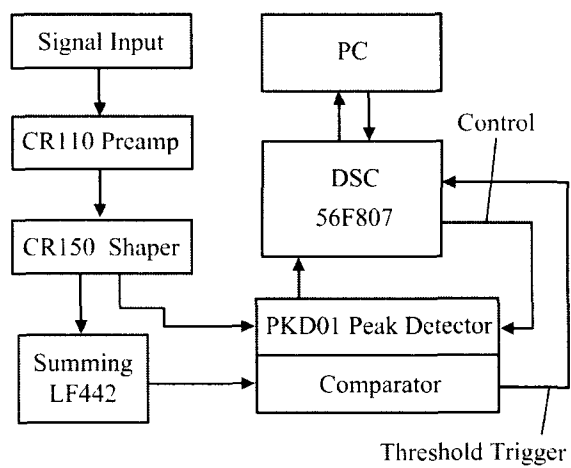


Figure 4.9: Functional Diagram of gamma camera acquisition system

Chapter 5

Single Channel Experiments with CsI and an APD

5.1 Single Crystal - Single APD

The APD's themselves were tested in section 3.2 to identify an ideal gain at which to operate each APD. In these tests the APDs were exposed to a red LED (exact emission spectrum unknown) which approximates the emission spectrum of the CsI crystal. However, the calculated optimum gain for each APD may not be an ideal representation of when they are coupled to a CsI crystal, but, it is assumed that the estimation is sufficiently accurate for our purposes. The same two APD's presented previously were used in the following investigation.

When a single APD is coupled directly to an APD sized crystal block, a clear energy peak can be detected. Due to the one to one coupling of crystal and detector, the maximum amount of energy of each incident photon should be collected by the APD. The crystal is wrapped in Teflon tape to aid in keeping the emission light from escaping the crystal. However, when the tape becomes coated with optical coupling grease which is between the crystal and the APD, the grease seems to have a tendency to degrade the reflectivity of the tape and may reduce its effectiveness.

A small Cobalt-57 (^{57}Co) source is used to approximate a ^{99m}Tc source; ^{57}Co emits mostly a $\sim 122\text{keV}$ gamma ray where ^{99m}Tc primarily emits photons at $\sim 140\text{keV}$. The sealed source is placed approximately one field of view away, directly above the crystal. At the time of measurement the source activity was estimated to be $4.3\mu\text{Ci}$. Figure 5.1 is a spectrum

collected from APD A while biased at its estimated optimal bias voltage (460 V from figure 3.6). The energy data collected has a 0 to 3.3V range and is sorted into 256 bins accordingly. The large peak, centered near bin 64, represents the photo peak for the cobalt source. In a full spectrum, in energy bins below the peak, the counts would rise sharply due to low level noise. It can be seen near bin 32 that this begins to occur, however, because of the presence of a lower level discriminator in the circuitry, energies lower than the threshold will not be sampled - causing the sharp drop off seen near bin 32. If the threshold level was reduced further more of the low level energies would be collected and if lowered enough the low energy counts could dominate the spectrum.




Figure 5.1: Spectrum - APD A - Bias Voltage 460V

The second APD (APD B) is operated at 420V, its estimated optimal bias voltage - deduced from the results of figure 3.6. The spectrum, figure 5.2, shows a similar peak, as expected. All other settings were kept constant, only the APD was swapped out and the bias adjusted to match the optimal voltage for the specific APD. In both spectra presented thus far, it is possible to see small but sharp fluctuations in certain places along the curve of the primary peak. It is believed that this is caused by combination of the non-linearity of the ADC and the binning process in generating the spectrum.

To exemplify the importance of operating an APD at its optimal




Figure 5.2: Spectrum - APD B - Bias Voltage 120V

voltage, a spectrum was collected operating APD A at APD B's optimal bias voltage of 420 V. The result is the spectrum presented in figure 5.3(a). As expected, the gain for APD A is significantly lower when operated at the bias voltage of 420 V. In comparison to figure 5.3(b) which is a background count taken with the same detector while biased at its optimal bias voltage (160 V), the spectrum looks much like the background spectrum, save a small pile of counts at the left perhaps representing the photo peak. When there is no significant photopeak to detect, the counts will be dominated by noise generated from various sources in the system, so when an APD is under biased the events produce electrical pulses that are hidden within the energy spectrum of the noise. When a spectrum with a source in place resembles the background signal (with no source in place), it implies that the pulses generated by the under-biased APD are not sufficiently large enough to stand out against the noise.

In each of these collected spectra presented here, the amplification stages on the pulse shaper's prototype board was held constant. However the spectrum can look very different by simply adjusting the amplification, bias voltage and/or triggering threshold. It is important to keep in mind that while the actual signal voltages can exist at any value, the ADC will only discriminate values that lie between 0 and 3.3V. Any signal lower than zero will be recorded as zero and likewise any value that exceeds 3.3 V will be recorded as the maximum recordable value (32768) ultimately landing in




Figure 5.3: a) Spectrum for APD A at Bias Voltage 420V. b) Background Spectrum - APD A - Bias Voltage 460V

the uppermost bin in the spectrum. The details of the operation of the Analog to Digital converter will be discussed at length in the following chapter. Increasing the amplification of the signal will essentially spread the spectrum out and any values that are amplified to a value outside of the 3.3V will fall into the last bin. Subsequently, decreasing the amplification will squeeze the spectrum toward the lowest energy bin, and any peaks that may have been lost above the 3.3V threshold may appear from the right side. By increasing the bias voltage the APD gain is increased, the result is very similar to increasing the amplification stage, until the bias voltage is brought too near the breakdown and the signal is dominated by noise. Finally, the threshold can also have a large impact on the spectrum that is collected. If the threshold is set too low, small fluctuations in the signal, from thermal noise for example, will be triggered on and sampled. This can result in saturating the detector with the low energy 'events' thus reducing the amount of time it spends detecting true events. This occurs because the low energy 'events' have a much higher frequency than true events. More intuitive perhaps is the consequence of raising the lower level discriminator too high - this results in a detector which ignores many, if not all, of the true events, depending on just how high it is set.

If each detector is operating at its optimal bias voltage, then a single threshold and similar amplification will suffice for all APD's operating in an array. To investigate the consequences of not operating the APD's each at their own 'optimal bias voltage' let's consider what would happen in a 2×2 array. In the ideal case, a flood source will produce a uniformly distributed (flat) image - but before correction, it will tend to

look much more like 4 evenly spaced hot spots forming a square. However, now consider one APD in the array operating at a voltage much below its ideal voltage - for a given energy, it would produce a much smaller pulse of electrons. This would be interpreted as the event not depositing all of its energy in that APD - thus the event position would be calculated to be further into the center than where it really occurred. Likewise, if one APD is operating at a higher gain than the others, it will tend to 'pull' all of the events toward it. This, of course, is not desirable for creating a uniform response.

Additionally, there are other drawbacks that plague the quality of the spectrum that should be discussed. Crystal geometry is crucial to **effective collection of gamma ray energy**. The crystal that was used in collecting every spectrum in this chapter was cut from an existing crystal. As a result, two sides of the crystal do not have a clean, smooth surface. Moreover, due to the 'softness' of the CsI, the surface has become scratched and marred in the process of cutting it from the original. These shortcomings in combination with the degradation of the reflectivity of the Teflon wrapping (as mentioned above) will inevitably lead to a broader spectrum as light will be more scattered, reflected and more likely to escape.

5.1.1 Isotope Energy Spectrum

In order to investigate the energy range that the system is able to detect, spectra from two sources of different energies will be collected. From their approximate peak location it will be possible to discern the energy values that correspond to the spectrum 'bins'. The spectrum for a $\sim 4.3\mu\text{Ci } ^{57}\text{Co}$ source was collected for 20 minutes; the spectrum for a $\sim 4\mu\text{Ci } ^{99m}\text{Tc}$ source was collected separately, for the same amount of time. In figure 5.4 the spectrum for each of these is presented. The thin solid peak represents the $\sim 122\text{keV}$ from the cobalt source and the dotted curve is the spectrum from the $\sim 140\text{keV}$ from the technetium.

From this plot, the centroid of the peak corresponding to the cobalt-57 ($\sim 122\text{keV}$) is centered at bin 58 ± 2 and the centroid from the technetium-99m is centered about bin 64 ± 2 . Assuming a linear energy response, we can calculate the location of a peak from any energy. However, because the CsI crystal is only 6mm thick (nominally) higher energy gamma rays may not deposit all of their energy thus resulting in an inaccurate representation of those peaks.




Figure 5.4: Technetium-99m (dotted) and Cobalt-57 (solid), 140 and 122 keV primary emission energies respectively.

Figure 5.5(a) below is a spectrum gathered from a $\sim 60\mu\text{Ci } ^{67}\text{Ga}$ source for twenty minutes. To compensate for the higher activity level, the source was placed approximately 15 field of views away. However, to do this, the source needed to pass through a cast aluminium lid; the increased scatter may have been sufficient to affect the spectrum. Gallium-67 has three prominent peaks, 93 keV, 185 keV and 300 keV (approximately) as indicated by the arrows on the plot. There are a number of other energies that also exist and these may be part of the cause for the poor resolution of the higher energy peaks. Also, it is possible that the CsI(Tl) crystal was not able to entirely stop the 300 keV (and above) gamma rays, thus leading to further spreading. The various discontinuities in the spectra (at channels 45, 55 and 85) are a result of non-linearity in the ADC as discussed previously.

If we were to apply what we learned from the technetium and cobalt sources, we would expect the 93 keV peak to appear around bin 48, the 185 keV peak around bin 79 and the 300 keV peak around bin 117. There also exists a 393 keV gamma ray that occurs with a quarter of the relative intensity of the 300 keV gamma ray; this may be the primary

source of events at and below the 150 bin.




Figure 5.5: a) Gallium-67 Spectrum with an APD detector. b) Ga-67, Co-57 and Tc-99m spectrum collected with a PMT detector

It is also interesting to compare the spectrum collected with the APD to that of a PMT detector. In figure 5.5(b), the three spectra are plotted on the same energy axis. The spectra were, however, collected independently and the sources had varying activity levels. Ga-67 was estimated at $7\mu Ci$, the Co-57 source at $4\mu Ci$ and the Tc-99m source was much hotter at $70\mu Ci$. The individual spectra were normalized to unity and when plotted together, the 93 keV peak and 180 keV peak in the Gallium, the 122 keV peak indicative of the Cobalt source and the 140 keV expected from the Technetium are all evident. Gallium exhibits a large peak at the edge of the spectrum: this can be partly attributed to the wide array of high energy gamma rays it is known to emit. A peak at bin 280 has been noted when room light leaks into the detector box and is expected to be the cause of the peak at 280 present here.

5.2 Double Sized Crystal - Single APD

In order to extract position information from the events, it will be essential to share the light across an array of APD's. To move from a single APD with a single crystal (of matching size) towards an array of APD's with a larger, shared crystal, it is important to investigate how sharing light will effect the output of the APD's. In the experimental setup with

one crystal coupled directly to one APD all of the light generated in the crystal will find its way to the APD (ideally), whereas if the light is shared among an array of APD's then only a fraction of the total gamma ray energy is collected by each APD (even in the ideal case).

To investigate the effects of light spreading in a larger crystal, a crystal block twice as large as an APD (in one direction) is coupled to a single APD, as shown in figure 5.6.




Figure 5.6: 1x2 Crystal Experimental Setup

The APD is located at one end of the crystal and the other end is left empty, not sampled at all. A spectrum was collected for two cases: a source directly over the APD and a source over the 'empty space'. This will simulate what will happen when events occur over neighbouring APD's in the array and was implemented by placing an uncollimated source (^{57}Co) over the desired location.




Figure 5.7: 1x2 Crystal Spectrum - Dark line: Source over crystal area occupied by an APD, Light line: Source over crystal area not occupied by an APD

The result is a pulse which is shifted to the left on an energy spectrum plot. The consequence is that the signal becomes closer to the low energy noise thus reducing the signal to noise ratio. If the signal to noise ratio is degraded too far the signal may not be recoverable and the system would become inoperable. In figure 5.7, the area of the crystal is twice that of the area of an APD, as shown in diagram 5.6. In figure 5.7, the bold curve, the source is placed above the crystal over the half that covers the APD. For the thin curve, the source is placed on the side that does not cover an APD. It is suspected that what is occurring here is that when the source is over the APD most events will occur directly above the APD thus depositing most of their energy into the APD - producing that larger peak on the bold curve. While the events that impact on the far side of the crystal will be few, they will also produce less light that will find the APD (i.e. most will just exit the back of the crystal into the empty space). In the thin curve, where the source is over the empty portion of the crystal, the majority of events occur in the empty space, shedding less energy onto the APD, while a smaller amount of gamma rays may still occur directly over the APD, thus producing that small rise in curve on the right side of the peak.

5.3 Effects of Magnetic field on APD Spectrum

The APD's were tested in a small magnetic field provided by a rare earth magnet. The magnet successfully rendered a photomultiplier-based detector useless by bending all of the electrons away from their targets, but the same magnetic field seems to exhibit zero effect on the APD detectors.

Shown in figure 5.8 is a spectrum collected for 5 minutes with the magnet placed directly beside the APD and a second spectrum without the magnet in the immediate vicinity. The solid line represents the spectrum collected with the magnet in place and the dotted line for the spectrum without the magnet in place. There is a slight discrepancy in the peak counts. To remove the magnet the case containing the APD needed to be opened, and the source moved out of the way. The source was more than likely placed in a slightly different orientation with respect to the detector than the first spectrum, and it is suspected this is the cause of the small discrepancy.




Figure 5.8: APD - Co-57 Spectrum in Magnetic Field (solid) and without (dotted) collected for 5 minutes

Chapter 6

APD Gamma Camera

The next step after collecting a spectrum from a single APD coupled to a single crystal is to expand this to an array of APD's and a larger crystal to accommodate them. This was accomplished by using a crystal of CsI(Tl) measuring 3 cm by 3 cm by 0.6cm - from the same original crystal that the single crystal was carved from to collect spectra. The four APD's, covered by a single crystal of CsI(Tl) were each connected to their own preamplifier and pulse shapers. Each of those were connected to a peak detect and hold circuit (PKD01) - the same as that used in the collection of spectrum. Each PKD01 is controlled by the same signals from the DSC such that every PKD01 is always in the same state as the next. The four outputs from the PKD01 are connected to the analog to digital converter (ADC) on the Digital Signal Controller (DSC). Two channels are connected to ADCA and two to ADCB - to allow all four conversions to occur simultaneously. The following sections present some of the images that were collected using the MCA system and various detectors. To ensure that actual events are triggered on, each signal is passed through a buffer to preserve the signal and then each buffer output is summed and passed to the threshold on the PKD01. Even if each individual signal is not strong enough to cross the threshold, if the summed value (indicative of the total photon energy) is sufficient it will cross the threshold and the ADC can still trigger.

6.1 Images

Once the system has collected and stored groups of data points in sets of four, relating to the four APD's in the array, then the construction of

an image can begin on the computer. Each event is localized by evaluating the four APD energies with respect to the total energy as described in section 2.2.7. Essentially an image is a two dimensional histogram. For example, a pixel, starting with no events would appear black (representing zero), but with each event that is placed in that pixel location the pixel colour would lighten. As this happens with each event and each pixel, a two dimensional histogram, what we call an image, is formed.

6.1.1 PET Module

A PET module is both readily available and professionally manufactured, making an ideal place to start testing because the signals should have fewer degradations within the detector than a similar system assembled in our lab. Also, the PET module is PMT based, providing an improved SNR over the APD's, thus making it, even further, a natural choice for initial image testing. The BGO crystal in the module is pixilated into an 8 by 8 grid. The pixels are cut most of the way through the crystal, but some space near the PMT's is left to allow for light spreading. We would expect to see these pixels in an image collected on this type of system before any image corrections are applied.




Figure 6.1: Image from PET module - 5 minute scan

Figure 6.1 is an image collected from a PET module with no image

corrections applied. It was collected for five minutes and clearly shows the 8 rows and 8 columns generated by the pixilated crystal. The module was biased with 1600 volts and signals were passed through a Cremat CR-113 CSPs designed for PMT signals and as always through the Gaussian pulse shaper and into the acquisition system. The source was a Na-22 source, with an emission energy of 511 keV. The low spatial frequency distortion evident in this image is a consequence of the Anger logic event localization, and can be corrected for with image correction techniques.

6.1.2 PMT on NaI(Tl)

Four PMT's were removed from a PET module and coupled to a large $10\text{cm} \times 10\text{cm}$ NaI(Tl) crystal from Saint Gobain. The PMT's are one inch square, nominally, and so only covered a small fraction of the area of the crystal. The PMT's were biased at 1600 volts and were coupled to the crystal with optical coupling grease. The crystal and PMT's were enclosed in a solid aluminium case manufactured to fit the crystal. The box was then covered with cloth to block any ambient light. The spectrum was collected for 5 minutes.




Figure 6.2: Image from four PMT coupled to NaI crystal.

The image collected is presented in figure 6.2. The four hot spots on the image correspond to the four PM tubes, and is apparently typical of

uncorrected images. Because the crystal is much larger than the PMT's there is a large opportunity for light to not all land on the PMT's but instead be lost out the sides of the crystal. This would result in a noisier, blurrier image as events will be placed incorrectly when light is lost. Also of note is the slight skew of the image. It appears to be stretched in one direction (or pinched in the other) along the diagonal. This occurs when the response of the PMT's is not perfectly matched. Previous images collected exhibited even stronger skewing until an attempt at balancing the gain on the APD's was made. This improved the skewing, however, further matching would, of course, be required if the system were to be used in studies.

The aluminium box has been fitted with braces to support a total of 9 PMT's. Once this has been assembled and the acquisition system upgraded to support 9 or more channels, improved imaging with the PMT's and NaI(Tl) will be possible.

6.1.3 APD on CsI(Tl)

Photomultiplier tubes provide large gains, on the order of millions of times, whereas avalanche photodiodes operate on lower gains, on the order of a hundred times. The poor amplification capabilities of the APD's make them much more difficult to work with. This is especially true when they are in an array and expected to share the light among a number of detectors. When this light is shared each signal is reduced, on average, by 4 times the amplitude as that experience in the single APD configuration used to collect the spectrum previously. What can happen if the system is not robust enough, is the actual signals can be pushed detrimentally close to the low energy noise. When this occurs it is difficult to set a triggering threshold that does not collect predominantly noise.

A 2×2 array of APD's was constructed on a circuit board and an unpixelated CsI(Tl) crystal was coupled to the array. The crystal was cut from an existing crystal, just as the single square was cut. And just like the single square crystal used to collect spectrum, the 2×2 crystal suffers from the same poor edges on two sides. This inevitably will lead to image degradation as light scatters unusually off these edges or even escapes despite the Teflon tape covering.

The array and crystal was placed inside a plastic box which was




Figure 6.3: Image grid collected with collimated source placed at multiple locations over an array of APD's coupled to CsI

then placed inside a cast aluminium box to which BNC jacks were connected. The primary goal of this arrangement was to shield the sensitive APD's from both light and RF interference. In an attempt to test the localization response of the 2×2 array a Na-22 source (511 keV) was collimated (with $\sim 7mm$ of lead with a single, small hole) and placed over the array, outside the plastic box - about one field of view away.

The collimated source was placed at various points above the array in a grid pattern and images were collected for 5 minutes at each location. Figure 6.3 shows the grid of images generated. The upper left image in the grid represents the image that was collected when the source was placed over the upper left corner of the APD array and the image in the center represents the image collected when the source was placed over the middle of the APD array, and so on. It is evident from this grid of images, that the array of APD's does exhibit a position sensitive response to a collimated source. However, without image correction it is expected that the localization will be particularly 'tubby' as demonstrated in figure 6.2 with four PMT's on NaI(Tl). Additionally, the greyscale image provided does

not do justice to the contrast range in the image - instead tending to make the images looked burned out, this is in fact not the case.

Furthermore, because the source is a high energy, Na-22 source, the CsI(Tl) would be unable to stop the entire gamma ray and this would lead to noisier, blurred images. However, in these tests, the higher energy source was needed to generate signals large enough to rise above the noise.

Taking a closer look at two different images collected with the same detector array described above, figure 6.4 displays both the image and the associated spectra below. The image shown in both figure 6.4(a) and figure 6.4(b) were collected with a bias voltage of 430 V and were taken for 20 minutes each. There are four spectra, one for each APD channel. The two images were taken with the source more or less centered over the array. However, in image 6.4(b), the source was in a location that seems to be slightly preferentially over one APD. This accounts for the change in weighting in the four 'hot spots' between the two images, but also, is exemplified in the spectrum as there is a clear increase in one channel in 6.4(d) when compared to the spectra in 6.4(c).

Recalling the sensitivity of the gain of each APD to bias voltage. As discussed in chapter 5, it would be ideal if each of the four APDs in the array could share identical characteristics. However, this is not what happens in reality. Instead each APD has its own photo current and dark current responses to bias voltage. As a result, each APD will have a **different ideal bias voltage at which to operate. The present design does not accommodate multiple bias voltages, and there is no mechanism to adjust individual gains as there were on the PMTs. As a result, no APD is necessarily operating at its ideal bias voltage - instead a voltage was found that worked well for the array as a whole. This value was found to be 430 volts and was found by scanning the range between 420 and 440V and the images were objectively reviewed to determine a voltage that provides a good compromise. The effect of bias voltage on the APD's, even when off by a few voltage can have a significant affect on its gain. This can result in particular APD's with higher gain essentially pulling events towards their centers, away from the other APD's.**




Figure 6.4: 4 APD array - Images and associated spectrum collected for 5 minutes with bias voltage of 130V

Chapter 7

Conclusions

This chapter will review the objectives and achievements of the project and outline future work that is required to move the project into the next stage of development. The areas of potential improvement will be identified and specific improvements discussed.

The project objectives were to design a gamma camera prototype that can operate on a general purpose acquisition board. The APD-based camera is not as limited by cost as much as how the camera will function in magnetic fields. As a consequence the PMT-based camera requires an acquisition system that is able to collect data at sufficiently high count rates while attempting to minimize cost. By selecting 'off the shelf' parts and selecting components that narrowly meet the operational requirements the price can be greatly reduced from the cost of those acquisition systems available commercially and those systems developed specifically for one-off applications. The APD camera limits the construction of the acquisition system in its physical construction, that is, certain components must be separated from the detector and the MRI. Specifically, two primary requirements need to be met for this camera. First the preamplifiers must be separated from the Gaussian pulse shapers to avoid interference, likewise the preamplifiers should be thermally insulated from the APDs to avoid gain fluctuations due to heat generated by the preamplifiers. Second, the digital electronics should be separable from the front end analog amplifiers and shapers. Digital electronics have a tendency to emit RF interference which can negatively affect the operation of the MRI camera.

These requirements have been met with the current design and the system has been tested in a controlled environment, collecting spectra while operating a single channel and for some configurations, images operating a

multi-channel array. Spectra have been collected and analyzed utilizing four detector system configurations: i) a PET module, utilizing PMTs coupled to a pixilated BGO crystal with a high energy Na-22 source; ii) a single PMT extracted from a PET module akin to the one mentioned previously, and coupled to NaI(Tl) with a Co-57 source; iii) the same PMT coupled to a CsI(Tl) crystal with a variety of sources, including Tc-99m, Ga-67 and Co-57 and iv) an APD coupled to a CsI(Tl) crystal with a variety of sources including Tc-99m and Co-57.

Images have been collected from multiple camera configurations, all in a 2 by 2 array. APDs with CsI(Tl), PMTs on NaI(Tl) and the PET module, and PMTs on a BGO crystal were all tested. The PET module was most successful in producing recognizable images, due to its pixilated crystal and the high gain of the PMTs. The PMTs on the NaI(Tl) crystal produced very 'tubby' images but did exhibit position sensitivity. With the APD camera the lack of matching operating gains on each of the APDs and the relatively low overall gains produced by APDs proved to make it difficult to produce reliable images. However, it was shown that the array was position sensitive and with improved crystal geometries it is expected a much more desirable result can be achieved.

Future work on the APD detector array is required to improve performance. As such, an improved, pixilated CsI(Tl) crystal is ready to be tested on the system to investigate further the geometric effects believed to be degrading the detectors spatial sensitivity. Future plans include advancing from a 4 input acquisition system to a 9 input system, allowing for a 3 by 3 array to operate with the pixilated crystal described above. Porting the prototype acquisition from the current arrangement to integration on a printed circuit board will allow for quick tests and adjustments while reducing troubleshooting time.

Since the inception of this project further designs of pixilated APDs have reached the market. Intrinsically, they operate in a slightly different manner than the current APDs, but also offer significant improvements in gain, making them comparable to PMT gains. Additionally, they are available in large arrays and operate with much lower bias voltages (10s of volts). By replacing the current APDs with an array module it is reasonable to expect significant improvement in SNR and low energy resolution.

When the system reaches the stage, in the near future, where reliable images are able to be collected it would be beneficial to apply the

image corrections as discussed in 2.2.7. These corrections can make a significant improvement in the image quality produced by a camera.

The system may benefit from calculating the position information for each event directly on the DSC before transferring the data - thus reducing the volume of data being transferred between the two devices. Also to circumvent the current bottleneck presented by the transfer speed available between the DSC and the PC via the USB bridge, it may be desirable to replace the current USB bridge with an upgraded version that allows for faster transfers.

The ultimate objective is, once a gamma camera module has been developed to the desired operation specifications, to create an array of these modules which can operate in concert. In this type of system it would be desirable to have a separate device dedicated to managing the communication between the gamma cameras and the PC. To connect each DSC module to the PC via a USB connection would require as many as eight simulated communication ports open on the PC simultaneously. To simplify this operation a single device can be used to negotiate data transfer between the cameras and the PC, while at the same time perhaps, offering a faster, inherent USB communication port to connect to the PC, thus further improving the transfer speeds.

Appendices

Appendix A

APD Pulse Size Sample Calculation

To estimate the pulse size of an event exiting the photomultiplication stage we assume that the entire photon is collected in the crystal. The general relationship is.

$$Q = E_{abs} \epsilon G$$

where Q is the pulse size (in electrons) E_{abs} is the absorbed energy per photon, ϵ is the quantum efficiency of the detector, and G is the gain of the detector [9].

For the camera in question let us consider a ^{99m}Tc photon ($140keV$) is fully absorbed in the CsI(Tl) crystal. The scintillation efficiency of CsI(Tl) is 64800 photons per MeV (from figure 2.6). The quantum efficiency of the S8664-1010 avalanche photodiode at the CsI(Tl) peak emission wavelength (540 nm) is about 86% [7]. The gain of the APD can be adjusted with bias voltage, so let us consider two different gain situations, 50 and 100. While we are at it, lets consider what would happen if we swapped the APD out for a PMT with a quantum efficiency of 8% (for the $CsI(Tl)$ spectrum) and a typical gain of 3.0×10^6 .

The calculation for the first case of an APD with 50 gain would be as follows.

$$\begin{aligned} N_p &= 64800 \text{ (photons/MeV)} \times 0.140 \text{ (MeV)} & (A-1) \\ &= 9072 \text{ (photons)} \end{aligned}$$

where N_p is the number of photons generated in CsI(Tl). Those photons from equation A-1 then are assumed to all interact with the APD or PMT with its given quantum efficiency,

$$N_e = 0.86 (Q.E.) \times 9072 \text{ (photons)} \quad (A-2)$$

$$\approx 7802 \text{ (electrons)}$$

where N_e is the number of electrons generated from N_p photons. Those electrons initiate an avalanche effect, whether it be an APD or PMT, which results in a gain which in this case is a gain of 50 times.

$$Q \text{ (electrons)} = 7802 \text{ (electrons)} \times 50 \text{ (gain)} \quad (A-3)$$

$$\approx 390 \times 10^3 \text{ (electrons)}$$

In Coulombs that is,

$$Q \text{ (Coulombs)} = 390 \times 10^3 (e^-) \times 1.6 \times 10^{-19} (e^- \text{ per Coulomb}) \quad (A-4)$$

$$\approx 0.062 \text{ picoCoulombs}$$

where Q is the size of the pulse.

Table A.1: Size of Pulse from CsI(Tl) for various configurations

System	Gain	Electrons	picoCoulombs
APD	5.0×10^1	390×10^3	0.062
APD	1.0×10^2	780×10^3	0.12
PMT	3.0×10^6	2.18×10^9	348

Appendix B

Peak Detection Circuits

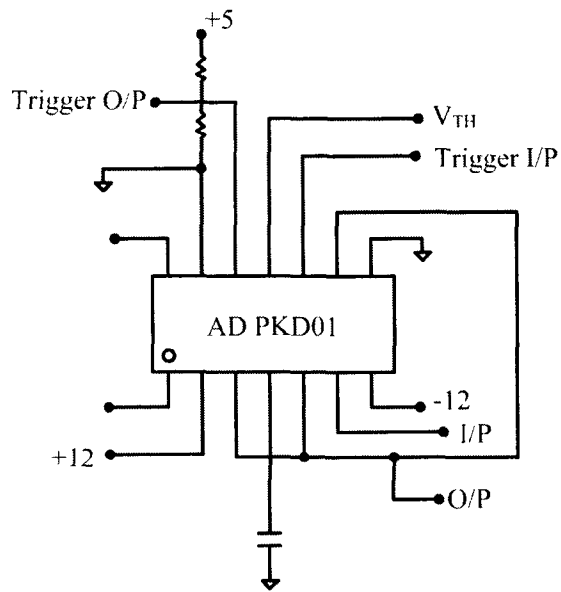


Figure B.1: PKD01 Circuit Diagram with Trigger threshold




Figure B.2: Hex CMOS Inverter used to convert 3.3V logic to 5V logic to operate the PKD01




Figure B.3: Summing Circuit - output connects to Trigger Threshold Input

Appendix C

Digital Signal Controller Code

```
/* Including used modules for compilling procedure */
#include "Cpu.h"
#include "Events.h"
#include "AS1.h"
#include "PTB0.h"
#include "EventTrigger.h"
#include "PTB2.h"
#include "EventFall.h"
#include "PTB4.h"

/* Include shared modules, which are used for whole project */
#include "PE_Types.h"
#include "PE_Error.h"
#include "PE_Const.h"
#include "IO_Map.h"

/* My Modules */

/* My Globals */
volatile int Rd_ready;
volatile int Event_ready;
volatile int Event_fall;
volatile unsigned int ADC_data[400]; // 2 bytes per data point.
volatile unsigned int ADC_index;

/* My Functions */
```

```
void putChar(unsigned int c);
void SendResults(void);
unsigned int AD1_GetOne(void);
void wait(unsigned int time);

/* Main Module */
void main(void)
{
/* Main variables */
unsigned char sciData; // used for Rx Char pointer.
unsigned int ad_data;
byte channel = 7;
unsigned int i;

    /*** Processor Expert initialization. DON'T REMOVE THIS CODE!!! ***/
    PE_low_level_init();
    /*** End of Processor Expert internal initialization.      ***/

// ADC A Register Settings //
ADCA_ADCCR1 = 0x0000; // clear stop bit and set
// SCANMODE to 'Single Sequential'
ADCA_ADCCR2 = 0x0000; // clock speed
ADCA_ADSDIS = 0x00FC; // CONVERT ONLY 'SAMPLE0 & SAMPLE1 & SAMPLE2'
ADCA_ADLST1 = 0x0070; /***0111 0000* SAMPLE 1 set to AN7
// and SAMPLE0 set to AN0 pin.

// ADC B Register Settings //
ADCB_ADCCR1 = 0x0000;
ADCB_ADCCR2 = 0x0000;
ADCB_ADSDIS = 0x00FC; // SAMPLE0 & SAMPLE1 to be sampled.
ADCB_ADLST1 = 0x0060; // SAMPLE1:BN6 SAMPLE0:BNO

// GLOBAL VARIABLE INITIALIZATION //
Rd_ready = 0; // serial Rx flag
Event_ready = 0; // eventInt flag
Event_fall = 0;
ADC_index = 0; // ADC data array index

// INITIALIZE PKD01 //
```

```
// hardwire connections
// PTB0 -> RESET
// PTB2 -> NOT_DETECT
PTB0_PutVal(1); // note, signal goes through the HEX inverter
PTB2_PutVal(1); // so a 1 o/p is actually a 0 at the PKD01.
for(;;)
{
while(Rd_ready == 0)
{
// wait for serial instruction //
}

Rd_ready = 0;
AS1_RecvChar(&sciData); //Get Rx byte and put it in sciData.

if(sciData == 65) // if 'A' was sent, then do the following.
{
// Reset PKD01 o/p value.
PTB0_PutVal(0);
PTB2_PutVal(0);

while(Event_fall == 0)
{
//wait for falling edge before starting next group
}
Event_fall=0;

// Prepare the external Equipment for sampling.

do
{
Event_ready = 0; // eventInt flag cleared - incase it
// was set before sampling was requested.
PTB0_PutVal(1);
PTB2_PutVal(1);

while(Event_ready == 0)
{
// wait for pulse event to convert
}
Event_fall=0;
```

```

//%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
// Wait 5 uSec ~ Rise time of the Pulse Shaper.
i=0;
while(i<16) // i=10 ==> 5us pause
{
asm
{
nop;
}
i++;
}
//#####

PTB2_PutVal(0); // HOLD value after 5uSec so peak is reached -
// but waiting longer may result in higher than
// normal 'pileup' as other events come in,
// they can bump the voltage up
// (if of higher E) and thus a preferential
// sampling of higher events will occur.
// By setting B2
// (NOT_DET) to 1 (HEXINV) we
// stop the value from increasing further.

//%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
PTB4_PutVal(1);

// set START bit - starts conversion //

ADCA_ADCCR1 |= (0x2000);
ADCB_ADCCR1 |= (0x2000);

// GRAB SAMPLES A & B WHEN READY
while((ADCA_ADSTAT&0x0001)==0)
{
// wait for ADCA SAMPLE0 to be ready.
}
ADC_data[ADC_index] = ADCA_ADRSLT0;
ADC_index++;
ADC_data[ADC_index] = ADCA_ADRSLT1;
ADC_index++;

```

```

while((ADCB_ADSTAT&0x0001)==0)
{
// wait for ADCB SAMPLE0
}
ADC_data[ADC_index] = ADCB_ADRSLT0;
ADC_index++;
ADC_data[ADC_index] = ADCB_ADRSLT1;
ADC_index++;

//#####
while(Event_fall == 0)
{
//wait for signal in to drop below threshold.
//Once it's low you can reset the PKD.
}
Event_fall = 0;
PTB4_PutVal(0);

//#####
// Reset PKD01 o/p value.
PTB0_PutVal(0);
PTB2_PutVal(0);
wait(4); // wait of 4 is enough time to bring signal
// from 3v to 0v. ~ 0.25 uSec

//#####

}while(ADC_index < 100); // 100 data pnts collected.
// ADC_index is GLOBAL and counts twice
// for each data point - one for the
// HiByte and one for the LoByte.

//#####
// Send Results to CPU //
SendResults();
// clear ADC index. //
ADC_index = 0;

```



```
#####  
} // end if 'A' is SCI_RX'd  
} // end for-ever loop.  
  
} // end MAIN  
  
/* putChar:  
* Takes: 8bit word to be sent  
* Action: sends a 8bit word out on AS1 SCIO  
* Returns: nothing  
*/  
void putChar(unsigned int c)  
{  
  unsigned int txfull;  
  do  
  {  
    txfull = AS1_GetCharsInTxBuf(); // AS1 bean code.  
  } while(txfull==1); // wait for TxBuffer to be empty  
  AS1_SendChar(c);  
}  
  
void SendResults(void)  
{  
  int i=0;  
  unsigned int h_byte, l_byte;  
  // send out all existing datapoints //  
  do{  
    h_byte = ADC_data[i]/256;  
    l_byte = ADC_data[i];  
    putChar(h_byte);  
    putChar(l_byte);  
    i++;  
  } while(i < ADC_index);  
}  
  
void wait(unsigned int time)  
{  
  unsigned int i;  
  //time = time ; // int time == 1 --> delay = 1us
```

```
for(i=0;i<time;i++)  
{  
asm  
{  
nop;  
}  
}  
}
```

Appendix D

Digital Signal Controller Photo




Figure D.1: Digital Signal Controller Photo

Bibliography

- [1] Aarsvold, J., Wernick M. *Emission Tomography*. San Diego, CA: Elsevier Academic Press, 2004.
- [2] Higgins, W., Glodo, J., Van Loef, E., et al., *Bridgman growth of LaBr₃:Ce and LaCl₃:Ce crystals for high-resolution gamma-ray spectrometers*. Journal of Crystal Growth 287 (pp 239-242), 2006.
- [3] Balcerzyk, M., Moszynski, M., Kapusta, M., *Comparison of LaCl₃:Ce and NaI(Tl) scintillators in γ -ray spectrometry*. Nuclear Instruments and Methods in Physics Research A 537 (pp 50-56), 2005.
- [4] Seguinot, J., Braem, A., Chesi, E., et al., *Novel Geometrical Concept of a High Performance Brain PET Scanner - Principle, Design and Performance Estimates*. European Organization for Nuclear Research, CERN PH-EP, 2004.
- [5] *Avalanche photodiode - A User Guide*. Fremont, CA: PerkinElmer Inc., 2006.
- [6] *Characteristics and use of Si APD (Avalanche Photodiode)*. Hamamatsu City, Japan: Hamamatsu Photonics K.K., 2004.
- [7] *Si APD S8664 Series - Short wavelength type APD*. Hamamatsu Photonics K.K., Hamamatsu City, Japan, 2005.
- [8] *CsI(Tl), CsI(Na)- Cesium Iodide - Scintillation Material*. Newbury, OH: Saint-Gobain Ceramics and Plastics Inc., 2007.
- [9] *CR-110 charge sensitive preamplifier: application guide*. Watertown, MA: Cremat Inc., 2006.
- [10] *CR-200 Gaussian shaping amplifier (Rev.2): application guide*. Watertown, MA: Cremat Inc., 2006.

- [11] *Monolithic Peak Detector with Reset-and-Hold Mode - PKD01, Rev. A.* Norwood, MA: Analog Devices, 2001.
- [12] *DSP56F800 User Manual: DSP56F801-7UM Rev. 6.* Denver, CO: Freescale Semiconductor, 2005.
- [13] Hasegawaa, B., Tanga, H., Da Silva, A., et al., *Dual-modality imaging.* Nuclear Instruments and Methods in Physics Research. A 471 (pp 140-144), 2001.
- [14] Gombiaa, M., Tartaric, A., Gambaccini, M., et al., *Attenuation compensation for breast tissue in combined CT/SPECT devices dedicated to mammography,* Nuclear Instruments and Methods in Physics Research. A 497 (pp 150-156), 2003.
- [15] Barbera, W., Iwataa, K., Hasegawa, B., et al., *A single CdZnTe detector for simultaneous CT/SPECT imaging.* Nuclear Instruments and Methods in Physics Research. A 505 (pp 595-598), 2003.
- [16] Hsua, W., Chenb, C., Wang Z., et al., *Feasibility study of small animal imaging using clinical PET/CT scanner.* Nuclear Instruments and Methods in Physics Research. A 571 (pp 457-460), 2007.
- [17] Moses, W., *Trends in PET imaging.* Nuclear Instruments and Methods in Physics Research. A 471 (pp 209-214), 2001.
- [18] Lai, J.S., Chen, D., *Design consideration for power factor correction boost converteroperating at the boundary of continuous conduction mode anddiscontinuous conduction mode.* Applied Power Electronics Conference and Exposition, 1993.
- [19] Kataoka, J., Sato, R., Ikagawa, T., et al., *An active gain-control system for Avalanche photo-diodes under moderate temperature variations.* Nuclear Instruments and Methods in Physics Reseach. Sec. A 564 (pp 300-307), 2006.
- [20] Sato, R., Kataoka, J., Kanai, Y., et al., *Development of 2cm-square Hamamatsu avalanche photodiodes for high-resolution X-rays and γ -rays detection.* Nuclear Instruments and Methods in Physics Reseach. A 556 (pp 535-542), 2006.
- [21] Carrier, C., Lecomte, R., *Recent Results in Scintillation Detection with Silicon Avalanche Photodiodes.* IEEE Transactions on Nuclear Science, Vol. 37 No. 2 (pp 209-214), 1990.

- [22] Britvitch, I., Deiters, K., Ingram Q., et al., *Avalanche photodiodes now and possible developments*, Nuclear Instruments and Methods in Physics Research. A 535 (pp 523-527), 2004.
- [23] Stapels, C., Squillante M., Lawrence W., et al., *CMOS-based avalanche photodiodes for direct particle detection*, Nuclear Instruments and Methods in Physics Research. A 579 (pp 94-98), 2007.
- [24] Shah, K., Farrell, R., Grazioso, R., et al., *Large-Area APDs and Monolithic APD Arrays*, IEEE Transactions On Nuclear Science, Vol. 48, No. 6 (pp 2352-2356) 2001.
- [25] Renker, D., *New trends on photodetectors*, Nuclear Instruments and Methods in Physics Research. A 571 (pp 1-6) 2007.
- [26] McClisha, M., Farrella, R., Myersa, R., *Recent advances of planar silicon APD technology*, Nuclear Instruments and Methods in Physics Research. A 567 (pp 36-40) 2006.
- [27] Lecomte, R., et al. *Design and Engineering Aspects of a High Resolution Position Tomograph for Small Animal Imaging*, IEEE Transactions On Nuclear Science, Vol. 41 No. 4 (pp 1446-1452), 1994.
- [28] Grazioso, R., Aykac, M., Casey, M., et al., *APD Performance in Light Sharing PET Applications*, IEEE Transactions On Nuclear Science, Vol. 52, No. 5 (pp 1413-1416), 2005.
- [29] Shao, Y., Silverman, R.W., Farrell, R., et al. *Design Studies of A High Resolution PET Detector Using APD Arrays*, IEEE Transactions On Nuclear Science, Vol. 47 No. 3 (pp 1446-1452), 2000.
- [30] Ziegler, S., *Development of a Small-Animal PET System*, Proceedings Of The IEEE, Vol. 93, No. 4 (pp 763-770), 2005.
- [31] Ziegler, S., Pichler, B., Boening, G., Rafecas, M., et al., *A prototype high-resolution animal positron tomograph with avalanche photodiode arrays and LSO crystals*, European Journal of Nuclear Medicine Vol. 28, No. 2 (pp 136-143), 2001.
- [32] Shah, K., Grazioso, R., Farrell, R., Cherry, S., et al., *Position Sensitive APDs for Small Animal PET Imaging*, IEEE Transactions On Nuclear Science, Vol. 51, No. 1 (pp 91-95), 2004.

- [33] Shao, Y., Cherry, S., Slates, R., et al., *Development of a PET detector System Compatible with MRI/NMR Systems*. IEEE Transactions on Nuclear Science, Vol. 44, No. 3 (pp 1167-1171), 1997.
- [34] Slates, R., Cherry, S., Shao, Y., et al., *Design of a Small Animal MR Compatible PET Scanner*. IEEE Transactions on Nuclear Science, Vol. 46, No. 3 (pp 565-570), 1999.
- [35] Raylman, R., Majewski, S., Velan, S., et al., *Simultaneous acquisition of magnetic resonance spectroscopy (MRS) data and positron emission tomography (PET) images with a prototype MR-compatible, small animal PET imager*. Journal of Magnetic Resonance 186 (pp 305-310), 2007.
- [36] Woodya, C., Schlyera, D., Vaska, P., et al., *Preliminary studies of a simultaneous PET/MRI scanner based on the RatCAP small animal tomograph*. Nuclear Instruments and Methods in Physics Research A 571 (pp 102-105), 2007.
- [37] Pichler, B., Judenhofer, M., Nutt, R., Cherry, S., et al., *Performance test of an Iso-apid detector in a 7-T MRI scanner for simultaneous PET/MRI*. Journal of Nuclear Medicine, 4 (pp 639-647), 2006.
- [38] Health Canada, *Arthritis in Canada. An ongoing challenge* Ottawa: Health Canada, 2003.
- [39] Centers for Disease Control and Prevention, *National and State Medical Expenditures and Lost Earnings Attributable to Arthritis and Other Rheumatic Conditions*. United States, 2003.
- [40] St. Clair, E., Pisetsky, D., Haynes, B., et al., *Rheumatoid Arthritis*. Lippincott Williams and Wilkins, Philadelphia, Pa, USA, 2004.