A MIXED ANALOG AND DIGITAL PIXEL ARRAY DETECTOR FOR SYNCHROTRON X-RAY IMAGING

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Daniel R. Schuette

August 2008

A MIXED ANALOG AND DIGITAL PIXEL ARRAY DETECTOR FOR SYNCHROTRON X-RAY IMAGING

Daniel R. Schuette, Ph.D.

Cornell University 2008

We present description and documentation of the development and first applications of the Mixed–Mode Pixel Array Detector, a new type of imaging detector for synchrotron based x-ray science. Today there exists a great gulf between the intense x-ray fluxes that modern synchrotron light sources are capable of producing and the capabilities of imaging detectors to measure the resulting signal. This detector is intended to help bridge this gulf by offering readout times of less than 1 ms, a dynamic range extending from single x-rays to a full well of more than 2.6×10^7 x-rays/pixel, capable of measuring fluxes up to 10^8 x-rays/pixel/s, with a sub-pixel point spread. These characteristics exceed, by orders of magnitude, the capabilities of the current generation of x-ray imagers. As a consequence this imager is poised to enable a broad range of synchrotron x-ray experiments that were previously not possible.

BIOGRAPHICAL SKETCH

Daniel R. Schuette was born in Iowa in 1977. He attended Washington High School in Cedar Rapids, Iowa from 1992 to 1996. As an undergraduate he attended the University of Chicago from 1996 to 2000, graduating with honors degrees in Mathematics (B.S.) and Physics (B.A.). From 2000 to 2001 he worked as a research associate on the Solar Tower Atmospheric Čherenkov Effect Experiment (S.T.A.C.E.E.); first at the University of Chicago's Enrico Fermi Institute then with the Department of Physics and Astronomy at the University of California, Los Angeles. Following this, in the fall of 2001, he entered graduate studies at Cornell University, receiving his M.S. in Physics in January of 2005 and then his Ph.D. in Physics in August of 2008.

In memory of Raphaël H. Kapfer (1977–2002)

ACKNOWLEDGEMENTS

A thesis in experimental physics is built on the work of many people; some directly involved in the work of the project, some who facilitate this work by providing material support or direction, and some who provide the ideas that the work is built upon. This is evident even in a relatively small project like the Mixed–Mode Pixel Array Detector (PAD). In this project the combined efforts of a large group of people, involved in different aspects of the project's planning and management as well as the detector design, fabrication, packaging, and testing, were required to bring about a functioning detector along with the array of supporting electronics, hardware, and software needed to make it functional. Beyond those who worked directly on the Mixed–Mode PAD there are also many people who deserve my thanks for their help in preparing me for this undertaking and aiding me throughout it. I consider myself fortunate to have been in an environment at Cornell that fostered developing the skills needed for this task by providing a wonderful group of mentors and compatriots. Because of either their direct contributions to developing the Mixed–Mode PAD or their indirect assistance, through the guidance and education they offered me, many people deserve my thanks. My apologies to anyone who might have been left off this list.

First and foremost I must thank my thesis advisor, Sol Gruner, for making this entire undertaking possible. Sol has been my advisor since the day I arrived at Cornell, nearly seven years ago. During this time there have been many successes, a few failures, some hard times, and some great ones. Throughout it all Sol has been invaluable, with his support, direction, patience, and understanding.

My other committee members, Jim Alexander, Alyssa Apsel, and Brad Minch also deserve thanks for watching over my progress at Cornell. In particular I would like to thank Brad, who was willing to devote a significant amount of time in the

early days of my studies at Cornell teaching me the basics of low-noise Mixed-Analog-Digital VLSI design.

Cornell is fortunate to have a small but vibrant detector x-ray development group, which I have been a member of for most of my Ph.D. tour. Thanks to the other members of the Cornell X-Ray Detector Development Group, past and present: Mark Tate, Matt Renzi, Alper Ercan, Lucas Koerner, Hugh Philipp, Marianne Hromalik, and Darol Chamberlain. To these people, I am greatly indebted for their many exchanges of ideas and equipment over the course of this thesis, as well as support in preparation and execution of the PAD CHESS runs. Within this group Mark, Matt, and Lucas deserve special recognition. Mark, for the experience, advice, and assistance that he brings to detector development and his willingness to pass these on to inexperienced graduate students. Matt, who was the senior graduate student when I started working on PADs and who supervised me in my first year. And last, but definitely not least, Lucas whose help has been invaluable in the most recent years of this project, particularly in preparing for and carrying out runs at CHESS.

Thanks also goes out to our collaborators at Area Detector Systems Corporation (ADSC). This group handled most of the packaging issues related to the hybrids as well as working in concert with Cornell on the detector design and testing. specifically Ron Hamlin, Tom Hontz, Wayne Vernon, Matt Allin, Skip Augustine, Don Abbe, and Doan Nyugen with whom I have worked closely over the duration of this project.

Returning to Cornell, I am indebted to the many people who have passed through the Gruner lab during my stay here. Most prominent in this group were Gil Toombes and Thalia Mills who, beyond being good friends and admirable scientists, often helped me with my children; watching over them when I had to

attend meetings or accomplish work at the lab. In addition to Gil and Thalia, there are many people whom I have benefitted from knowing and working with: Pascale Chenevier, Anurag Jain, Raphaël Kapfer, Marcus Collins, Nozomi Ando, Chae Un Kim, Yi-fan Chen, Darren Southworth (who briefly worked on the Mixed–Mode PAD), Dag Anderson, Tom Caswell, and Buz Barstow to offer an extensive, but not all inclusive, list. Finally, in the Gruner lab, our machinist, Marty Novak, deserves special thanks for the wonderful camera housings that he built.

Then, thanks to all the people who have kindly read through and helped me copy edit this thesis; in particular Kate Green, who reviewed this work from cover to cover.

I also wish to thank the various sponsors throughout my graduate career that made this work possible: Cornell University for a G-line fellowship and the U.S. Department of Energy for support of detector work by our Cornell Group (grants DE-FG02-97ER62443, and DE-FG02-97ER14805). The primary support of this work was a subcontract from ADSC's National Institute of Health grant RR014613 (Dr. Ron Hamlin, P.I.). Part of this work was performed at the Cornell High Energy Synchrotron Source (CHESS), which is a national user facility supported by the National Science Foundation and the NIH/NIGMS under Award DMR-0225180.

Although only indirectly related to my dissertation work, Corbin Covault deserves my thanks for introducing me to research in physics while I was an undergraduate at the University of Chicago. I also owe a great debt of thanks to René Ong, my undergraduate advisor at the University of Chicago, for his mentoring, when I was an undergraduate, and friendship, in those years and the years since.

Finally, thanks to my family. To my parents, who were my first teachers. To my wife Dorothy—without all her love, patience, and support (not to mention

editing skills) this thesis would not have been possible. To my elder son, Henri, who made sure I set the thesis aside and went out to play now and then. And to my younger son, Conall, who has been my consistent companion through much of this writing and a persistent reminder of how long this task has taken to complete.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVATIONS

ADC	Analog-to-Digital Converter
ADSC	Area Detector Systems Corporation
AFM	Atomic Force Microscopy
ASIC	Application Specific Integrated Circuit
CCD	Charge-Coupled Device
CDS	Correlated Double Sampling
CHESS	Cornell High Energy Synchrotron Source
CSR	Control Shift Register
CTF	Contrast Transfer Function
ELT	Enclosed Layout Transistor
ESR	Edge Spread Response
ESRF	European Synchrotron Radiation Facility
FPGA	Field-Programmable Gate Array
GISAXS	Grazing Incidence X-Ray Scattering
LCLS	Linear Coherent Light Source
LOCOS	LOCal Oxidation of Silicon
LSR	Line Spread Response
MOSIS	Metal Oxide Semiconductor Implementation Service
MTF	Modulation Transfer Function
NSD	Noise Spectral Density
OTF	Optical Transfer Function
PAD	Pixel Array Detector
PLD	Pulsed Laser Deposition
PSF	Point Spread Function
PSR	Pixel Spot Response

PSRR	Power Supply Rejection Ratio
RHEED	Reflective High–Energy Electron Diffraction
SEM	Scanning Electron Microscopy
SEL	Single Event Latchup
SEU	Single Event Upset
SPICE	Simulation Program with Integrated Circuit Emphasis
STM	Scanning Tunneling Microscopy
STI	Shallow Trench Isolation
TEM	Transmission Electron Microscopy
TSMC	Taiwan Semiconductor Manufacturing Company
TID	Total Ionizing Dose
VLSI	Very Large Scale Integration

LIST OF SYMBOLS

CHAPTER 1

INTRODUCTION

When the microscope was first invented in the early 1600s, it opened up a window onto a new world of scientific discovery, revealing organisms on a scale that had not been considered before. While remarkable, this feat was not unique as the pattern of new scientific instruments spawning new scientific understanding is a general paradigm that is repeated throughout the history of science. Particularly in the 20^{th} and 21^{st} centuries, where access to high levels of technology has become easier and more widespread, instrumentation has been a driving force behind scientific advancement, with new instruments (like the electron microscope, space telescopes, particle colliders, and many others) opening otherwise unaccessible windows onto the universe and resulting in a proliferation of scientific knowledge.

Despite the impact of new scientific instruments, a focus on the development of new instrumentation is not a widely acclaimed path for a physicist—offering little of the exotic appeal of fields like string theory or cosmology. Consequently, one sometimes encounters the bias that instrumentation is not a science but an extension of engineering. This view is rather limited, as it presumes the act of discovery is decoupled from the thing that enabled that act. A well conceived instrument is one that enables new scientific discoveries by virtue of its design; or, in other words, because it enables scientists to look where it was not possible before and where there are interesting things to discover. The latter portion of this point is quite critical in distinguishing scientifically meritous instrument development from more general engineering as it implies that the instrument developer appreciates where opportunities exist for scientific discovery and that their instrument is designed to realize these opportunities. It also implies that the instrument developer looks for a multiplicative effect from the instruments they create, so that their efforts can

enable, through the application of the instrument by many scientists to a broad field of questions, the creation of a body of scientific knowledge that will live long past the instrument's working life.

This thesis offers description and documentation on the development and first applications of a new type of imaging detector for synchrotron based x-ray science, the Mixed–Mode Pixel Array Detector (PAD). As will be argued in the remainder of this chapter and those that follow, this work constitutes scientifically meritous instrument development because it was designed with, and achieves, the objective of opening an otherwise unaccessible window for scientific investigation onto, what is expected to be, a very fertile field for scientific discovery.

1.1 X-Rays & Synchrotron Light Sources

X-rays have proven to be an extremely powerful structural and compositional probe of matter spanning length scales from the macroscopic to the atomic. Evidence of the utility of x-rays to science can be seen in the rapid development of sources and techniques that followed Wilhelm Röntgen's initial description of 'a new kind of radiation' in 1895. In their earliest application, x-rays were used in radiography to gain information about the structure of optically opaque objects. Yet, within a decade after this initial report, it was discovered that they could yield compositional information as well, when it was observed that gases exposed to x-rays emit x-rays at lower energies, characteristic of their elemental composition. And before two decades had elapsed, the first diffraction of x-rays from a crystal was observed, extending the reach of x-ray research into the structure of matter on atomic length scales. Since then, x-rays have played a critical role in some of the most important scientific discoveries of modern times, including contributing significantly to at least sixteen Nobel Prizes spread between the fields of Physics, Chemistry, and

Medicine.[1]

The achievements of x-ray science would not have been possible without steady improvements in x-ray sources and detectors. In terms of x-ray sources, the goal had been to develop successively more intense sources of radiation that allow greater control over the spectral distribution of emitted x-rays. To measure the signals produced by these, a vast array of detection methods have been developed ranging from point detectors to two dimensional imagers, based on technologies including, though not limited to x-ray film, photostimulable phosphors, geiger counters, scintillator, and direct semiconductor detectors.

To appreciate what distinguishes the Mixed–Mode PAD from these previous detectors and why it has the potential to enable a broad range of new x-ray science, one must understand that the Mixed–Mode PAD is an x-ray imager designed for use with a very unique source of x-rays, the synchrotron light source. A synchrotron light source, often shortened to simply synchrotron, is a facility that produces intense beams of x-ray radiation through the motion of ultra-relativistic charged particles. This phenomena occurs when, at these extreme velocities, the standard radiation field of an accelerated particle is Lorenz Transformed into the laboratory frame of reference. Under this transform, radiation from the particle is beamed into the direction of the particle's motion, Doppler Shifting the spectra to higher

[1](1901) W. Roentgen in Physics for the discovery of x-rays. (1914) M. von Laue in Physics for x-ray diffraction from crystals. (1915) W. H. Bragg and W. L. Bragg in Physics for crystal structure derived from x-ray diffraction. (1917) C. Barkla in Physics for characteristic radiation of elements. (1924) K. Siegbahn in Physics for x-ray spectroscopy. (1927) A. Compton in Physics for scattering of x-rays by electrons. (1936) P. Debye in Chemistry for diffraction of x-rays and electrons in gases. (1962) M. Perutz and J. Kendrew in Chemistry for the structure of hemoglobin. (1962) J. Watson, M. Wilkins, and F. Crick in Medicine for the structure of DNA. (1964) D. Hodgkin in Chemistry for the determination of the structure of penicillin and other important biochemical substances. (1976) W. Lipscomb in Chemistry for the determination of boranes. (1979) A. Cormack and G. Hounsfield in Medicine for computed axial tomography. (1981) K. Siegbahn in Physics for high resolution electron spectroscopy. (1985) H. Hauptman and J. Karle in Chemistry for direct methods to determine x-ray structures. (1988) J. Deisenhofer, R. Huber, and H. Michel in Chemistry for the structures of proteins that are crucial to photosynthesis. (2003) P. Arge and R. MacKinnon in Chemistry for discoveries concerning channels in cell membranes. (http://nobelprize.org)

energies creating a radiation source with, as we will see shortly, quite exceptional characteristics [1].

The traditional method of x-ray generation, used since x-rays were first discovered, has been through the vacuum tube. In these structures, electrons are boiled off of a cathode into a vacuum, then accelerate through a strong (many kV) electrical field until they strike the anode target. Within the anode, these highly energetic electrons excite atomic transitions that give rise to characteristic x-ray emission lines, whose energy depends on the material composition of the anode, along with a background of bremsstrahlung radiation. Because x-ray production via this technique is very inefficient (with more than 99% of the incident energy converted to heat) these sources are fundamentally limited in the x-ray brilliance,

$$\text{Brilliance} \equiv \frac{\text{photons/s}}{(\text{mm}^2 \text{ source area})(\text{mrad}^2)(0.1\% \text{ bandwidth})}, \tag{1.1}$$

by the heat load they are capable of imparting on the anode.

Synchrotron sources, on the other hand, exhibit intrinsic characteristics that distinguish them, in terms of their spectra and brilliance, from the x-ray tube. Because of their unique, and highly efficient, x-ray emission mechanism they can provide very intense fluxes of x-rays over a relatively broad spectrum that, itself, may be controlled through the speed and acceleration of the particles. The brilliance of these sources are only limited by the ability of the accelerator physicists to control the beam dynamics and current, a science which they have been improving upon steadily, achieving roughly an order of magnitude increase in photon emission brilliance every four years [106]. In addition, because their high flux extends over a broad spectrum, it is possible to use x-ray optics to isolate the beam within a selected energy bandpass. These two factors of high intensity and energy selectabilty allow experiments to be performed that would be either impossible or take excessive time with conventional x-ray sources [61].

While extremely useful, facilities to produce synchrotron radiation require a large investment to build and maintain. Today, world-wide there exist over 70 synchrotron sources in various stages of planning, construction, or operation; representing a combined investment of ∼$10B [40] where within the US alone the annual operating budget of synchrotron facilities is estimated at over $200M [99]. Yet the utility of these facilities is evident in user base they support, which is $> 10,000$ scientists [40], covering an array of fields including Physics, Biology, Materials Science and Engineering, as well as more unexpected fields, such as Archaeology and Art History. This broad user base provides a field where new detector technology can take root, sprouting new scientific discoveries.

1.2 Need for New Detectors

The investment into synchrotron technology and facilities that has taken place from the 1960s (when synchrotron experiments were performed parasitically, operating off of machines built for, and typically performing, particle physics experiments) to today (when we are looking towards a fourth generation of synchrotron light sources in facilities like x-ray free electron lasers and energy-recovery linear accelerators) has resulted in tremendous increases in available x-ray flux. This phenomenal growth has produced an abundance of x-ray flux, to the extent that fluxes in excess of 10^{13} x-rays/s/0.1% $\frac{\Delta E}{E}$ are typical of modern beamlines. Yet because a commensurate level of investment has not been made into x-ray detector technology, there is today a general concensus in the synchrotron community that there exists such a gulf between the capabilities of synchrotrons to deliver high photon fluxes and the capabilites of detectors to measure the resulting x-ray signals that the detector is the limiting element in many experiments [99, 61].

A clear consequence of this gulf is that it is quite common in current experi-

ments to find experimenters attenuating their beam or a portion of the scattering pattern because of the flux and well depth limitations of available detectors [61, 32]. In imaging detectors, the reason for this is a combination of the extended point spread and limited dynamic range of the current generation of x-ray imagers. The first property, the point spread, is the extent to which a signal incident at a point on the detector is observed at other locations. In the presence of very intense signals on the imager, an extended point spread can make it impossible to resolve weak signals. In the case of the dynamic range, many systems studied at synchrotrons have dynamic ranges that span many orders of magnitude (particularly diffraction experiments). Because of their intensity, synchrotron light sources make accessing signals in the weaker portion of this range possible, in principle, yet this is limited by the dynamic range of which current imagers are capable.

A second, and arguably more serious, consequence of this gulf is that there exists a broad range of time resolved phenomena that could be studied with currently available x-ray fluxes, but are not accessible because of the lack of capable detectors for measuring these fluxes and framing at high rates [99, 61]. Today, dynamic measurements using x-rays are either limited to point or 1D detectors for continuous measurements or to the imaging of repetitive phenomena. In the latter category, one requires cyclic systems or pulse–probe style experiments, where a shuttering mechanism (either electronic or mechanical) is used to gate x-rays onto the detector for times much shorter than the continuous frame rate. In this way, these ensemble time resolved measurements are capable of building up a temporal mosaic of the ensemble response of the system under investigation. This technique suffers from the criticisms that: it is cumbersome and difficult, at least in cases where a mechanical shutter must be used, are therefore is not widely used; and, more seriously, that there exists a large class of dynamic systems that are inac-

cessible by these techniques because they do not conform to the limits outlined above. In this class are many self–assembled growth phenomena where the difficulties of exactly reproducing initial conditions make it impossible to temporally tile multiple data sets.

Thus, the imager frame range, dynamic range, and point spread may be identified as areas where improving the imaging detectors available at synchrotron light sources will expand the possibilities for x-ray science. The objective of the Mixed–Mode PAD is to answer these three needs in a single detector by offering a device capable of framing at continuous rates that are nearly three orders of magnitude faster than what is attainable with current detectors, with a dynamic range nearly four orders of magnitude beyond the capabilities of current imagers, and with a point spread that is essentially limited by the pixel size of the detector. These advancements have the potential to not only improve science that is currently being done but to enable new types of x-ray science, specifically continuous time resolved imaging experiments on the ms timescale.

1.3 Document Organization

The body of this work is divided into five parts with the intention to divide the discussion conveniently for different portions of its intended audience. Given the length of this thesis, the author does expect that most readers would want to read it from cover to cover. So the following is offered as a guideline to help readers select what they might find interesting and relevant.

In the first part, encompassing chapters 2 and 3, a general background to pixel array detectors is offered. This background reviews the methodology and physics that underlies the operation of PADs, concluding with a historical review of PADs preceeding and contemporary with the Mixed–Mode PAD.

In the second part, encompassing chapter 4, the design of the final Mixed–Mode PAD pixel is presented and analyzed. This section is intended to help students starting on new PAD designs by laying out the considerations that went into the Mixed–Mode PAD pixel, while also providing a detailed reference on operation and performance expectations of this pixel.

The third part, encompassing chapter 5, discusses the prototype camera built for characterizing the hybrid detector performance. It is included mainly as a background chapter for the characterization and experimentation work that follows it, although it contains some control and readout timing information that should be of interest to anyone working on a controller for this detector.

The fourth part, encompassing chapter 6, reports on the characterization of the detector. This chapter is intended for those interested in understanding the performance capabilities and limitations of the device. It also contains a section discussing image correction algorithms that may be used to improve the quality of the data from the Mixed–Mode PAD along with a discussion of how these calibration terms may be measured.

The final part, chapter 7, presents results from the first experiments performed with the Mixed–Mode PAD and is intended for the polymaths among the reading audience.

CHAPTER 2

PIXEL ARRAY DETECTOR FUNDAMENTALS

Before embarking into the rough waters of the design, characterization, and first experiments of the Mixed–Mode PAD, a basic understanding of what constitutes a Pixel Array Detector is advised. In addition, to appreciate the design decisions made, a familiarity with some fundamental aspects of semiconductor physics is recommended. This chapter provides both of these by offering a discussion of the general Pixel Array Detector methodology along with a comparison to contemporary synchrotron x-ray imagers so as to highlight the distinguishing features of PADs. This is followed by a selective review of basic semiconductor physics with an emphasis on topics relevant to the design and operation of a PAD. The chapter concludes with a discussion of radiation effects on a PAD hybrid.

2.1 PAD Methodology

As discussed in the introduction the name Pixel Array Detector or PAD denotes a broad class of x-ray detectors that incorporate custom signal processing electronics into each individual pixel. For our purposes, PADs are two-layer hybridized devices: one layer acts as a detector, directly converting x-ray photons into an electrical signal, while the other layer contains custom electronics that process this electrical signal. A grid of metallic interconnects, called bump bonds (commonly indium or solder), join the individual pixels of the two layers. This configuration is illustrated in figure 2.1.

This hybrid methodology distinguishes PADs from more conventional synchrotron x-ray imaging devices by allowing them to directly detect x-ray photons and immediately process the resulting signal with custom electronics. To see and understand the impact of these features, it is useful to put our discussion in context

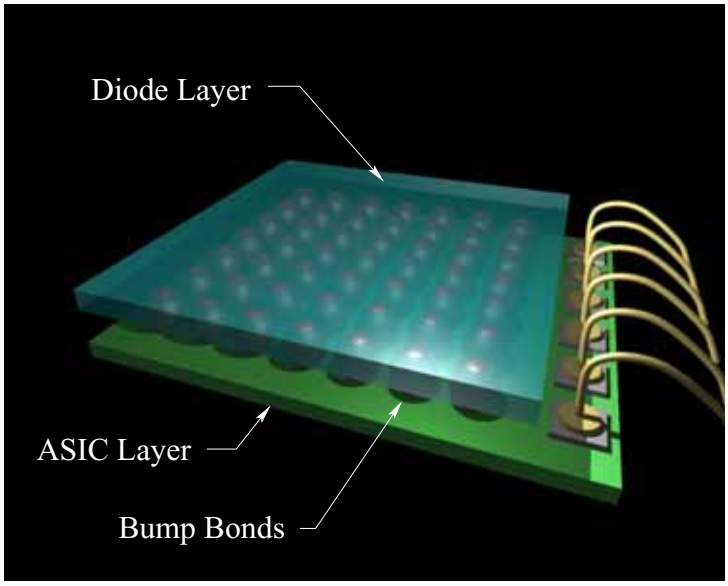

Figure 2.1: Artist's conception of a Pixel Array Detector (PAD) illustrating: the detector diode layer, responsible for converting photons into electrical charge; the signal processing application specific integrated circuit (ASIC) layer, responsible for processing the signal generated by the detector diode; and the array of bump bonds that provide electrical interconnects between corresponding pixels on the diode layer and the ASIC. Thanks to Hugh Philipp for the image.

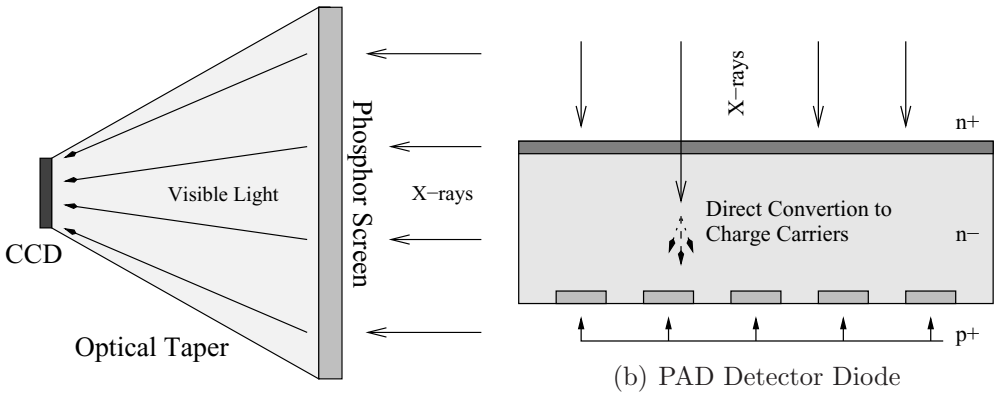

(a) Phosphor Coupled CCD

(b) PAD Detector Diode

Figure 2.2: Comparison of the indirect method of x-ray detection used in phosphor coupled CCDs with the direct detection approach of PADs. Panel (a) offers a cut away of a CCD detector showing: the phosphor screen, that converts x-rays into optical light; the optical taper, that collects and transits this light; and the CCD which quantitatively records the light. Panel (b) describes the detector diode layer of a PAD hybrid, illustrating: the uniform $n+$ region at the diode surface, used to distribute the reverse bias voltage; the $n-$ region in the detector bulk, where x-rays directly convert into photocurrent; and the pixelated $p+$ regions at the base of the detector where the photocurrent is collected. Neither figure is drawn to scale.

by comparing a PAD with the current workhorse of synchrotron x-ray imaging, the phosphor coupled CCD.

Phosphor coupled CCDs were first developed for synchrotron science in the late 1970s and early 1980s [42]. The scientific CCDs used in these devices exhibit exceptional, nearly quantum limited, sensitivity combined with excellent stability and linearity as well as a broad dynamic range [96]. As they would be ineffective for imaging x-ray diffraction patterns directly, because of an x-ray transparent active thickness, a thin phosphor sheet is used to convert the x-rays into visible light. This light is then conveyed via an optical fiber taper onto the CCD, as the cutaway of a typical phosphor coupled CCD shown in panel (a) of figure 2.2 illustrates. Reduction factors as high as 5:1 are commonly used to condense the image of the phosphor screen, allowing a single CCD to span a substantially larger active area while also facilitating the tiling of multiple CCDs together into a large area detector.

While phosphor coupled CCDs are impressive devices, they do have a number of limitations. First, the indirect x-ray detection method of an optically coupled taper and phosphor screen drastically reduces efficiency of the detector, particularly with the image reduction factors of the optical tapers commonly employed. As a result, the signal yield of these systems is typically much less than a hundred e^- per 10 keV x-ray.[1] Modern CCDs are sensitive enough to make up for this low efficiency, though only at the expense of the detector's readout time—which is necessarily long to maintain the required fidelity of the detector's analog signal. As a consequence, these detectors normally require a second or more of dead time to read out, limiting their maximal frame rate to a fraction of a Hz.

In contrast, Pixel Array Detectors do not use an intermediate stage but rather

[1]For example, the ADSC Q270 (www.adsc-xray.com) reports an efficiency of 22 e^- per 12 keV x-ray whereas the MAR USA SX-165 (http://www.mar-usa.com) reports 8 e^- per 8 keV x-ray, with values taken from manufacture published spec. sheets.

collect the charge directly produced by x-ray conversion in the detector diode layer, as illustrated by panel (a) of figure 2.2. This method yields roughly 2,700 charge carriers ($\sim$ 0.5 fC) for a 10 keV x-ray. The much greater charge yield produced by direct detection relaxes the noise performance constraints needed to attain x-ray quantum limited signal-to-noise performance. This in turn gives PADs a level of design flexibility not enjoyed by phosphor coupled CCD systems. A clear illustration of this is in the processing technologies available for fabrication of PAD hybrids in contrast with CCDs. To attain their exceptional level of sensitivity CCDs require dedicated fabrication lines with special processing steps. The PAD signal processing ASIC, on the other hand, may be fabricated on commercial CMOS fabrication lines in technologies with notably poorer noise performance than those used to manufacture scientific CCDs. This offers a significant economic advantage through the relative economies of scale. Specifically, due to the much higher volume of CMOS device fabrication substantial infrastructure exists to support relatively inexpensive prototyping while the variety of vendors and CMOS processes available serves to lower the price of full-lot fabrications. Beyond the economic advantage, this level of signal also provides a design advantage in terms of the types of signal processing that are possible. Unlike conventional x-ray imagers, which are limited to aggregate measurements of the total signal yield, a x-ray detected by a PAD yields sufficient signal to permit processing of individual photons [87]. Not all PADs take this approach, but the fact that it is possible illustrates the great degree of flexibility offered by direct x-ray detection.

In addition to relatively poor efficiency, a second limitation of the indirect detection method is the spatial spreading of the signal from individual x-rays. Sometimes referred to as the point spread function, though this term technically only applies in digital imagers whose analog impulse response is not degraded by

pixelation.[2] Measured phosphor coupled CCD point spread functions typically have an extent on the scale of mm at the 1% level [96]. In the direct detection method of Pixel Array Detectors, the analog impulse response is determined by the spread of charge carriers generated from individual x-ray conversion events. As will be discussed in sections 2.2.1, 2.3, 6.3, and 6.4.4, the precise spreading of charge carriers depends on a number of factors but under typical operating conditions will be less than 50 μm. This tight spatial response makes it possible for the PAD to detect weak signals in much closer proximity to intense signals than would be possible with a phosphor-based synchrotron x-ray imager.

The final distinguishing feature of the PAD methodology that we will discuss here is the capacity these imagers have for integrating pixel-level signal processing electronics. Traditionally, x-ray detectors have been fixed-point analog integrators with no capability to alter their behavior *in-situ* with an exposure [39]. The advent of CCD based x-ray imagers changed this to a limited degree by offering the ability to shift charge, effectively relocating the imager pixels, within an exposure. As most commercially available CCDs are only capable of charge shifts in one dimension, this technique has seen limited applications in x-ray imaging [27, 108, 52]. More exotic CCD architectures allowing two dimensional shifts have been developed and applied in other fields such as astronomy [18] and could find application in synchrotron CCD systems, yet even this degree of functionality pales in comparison to the possibilities of a modern PAD.

As previously mentioned, PADs are capable of integrating custom signal processing electronics into each pixel. This degree of integration permits the creation of smart pixels whose degree of functionality is predominantly limited by only the area available within the pixel and the imagination of the designer. This assertion

[2]The classical definition of the point spread function assumes translational invariance of the detector [54]. This presents a problem for pixelated devices if the pixelation breaks this symmetry.

is borne out by the wide variety of PADs that have been or are in the process of being developed, including: digital counting PADs with energy discrimination capabilities [71, 48]; analog integrating PADs with multi-frame memory integrated into each pixel [85]; and high continuous frame rate PADs incorporating in-pixel full or partial analog-to-digital conversion [76, 4].

Because of their unique hybrid methodology and its resulting distinguishing characteristics of direct x-ray detection and in-pixel signal processing, Pixel Array Detectors represent a new generation of x-ray imagers with a potential to greatly advance synchrotron science. The remainder of this chapter focuses on the physical principles permitting PAD performance, with particular attention given to the effects of subjecting these devices to the intense radiation environment found at synchrotron light sources.

2.2 PAD Semiconductor Physics

Understanding a complicated integrated circuit device, such as a Pixel Array Detector, is difficult without a few basic concepts in semiconductor electronics and an understanding of how they are applied to produce selected semiconductor devices. This foundation is necessary to understand the effects of radiation, both wanted and unwanted, on these devices and to appreciate the steps taken to mitigate damaging effects. Towards this end, this section presents a brief review of semiconductor physics, where the scope has been limited to elements that are directly relevant to understanding PAD performance and radiation hardness. A basic background in semiconductor physics is assumed (*e.g.* band gap, intrinsic material, Fermi Level, conduction band, valance band, etc.). Readers interested in a more thorough discussion of semiconductor physics are referred to [97], [93], and [70].

2.2.1 Charge Concentration

Semiconductors are materials that exist in the gray area between conductors and insulators. These materials exhibit a band gap (*i.e.* the energy separating the conduction band from the valence band) small enough that at room temperature there are a small, relative to a conductor, but appreciable, relative to an insulator, number of carriers with enough thermal energy to enter the conduction band. Control over the concentration of electrons in the conduction band and holes left behind in the valance band is fundamental to semiconductor device physics. It is accomplished through the introduction of ions (*dopants*) into the silicon lattice that either supply electrons to the conduction band (*donors*, commonly antimony, phosphorus, or arsenic), or bind an electron from this band leaving behind a hole in the valance band (*acceptors*, commonly boron, aluminium, or gallium) [93]. These dopants alter the Fermi Level (E_{F}) of the semiconductor, changing the carrier concentrations via:

$$n(x) = n_i \exp\left\{\frac{E_{\mathrm{F}} - E_i}{kT}\right\}, \tag{2.1}$$

$$p(x) = n_i \exp\left\{\frac{E_i - E_{\mathrm{F}}}{kT}\right\}, \tag{2.2}$$

where n_i is the intrinsic carrier concentration, k is the Boltzmann constant, T is temperature, and the Fermi Level (E_{F}) of the intrinsic semiconductor is $E_i = \frac{E_{\mathrm{C}}+E_{\mathrm{V}}}{2}$, where E_{C} is the energy of the conduction band and E_{V} is the energy of the valance band [93]. The difference between the energy of the valance band and the conduction band is the band gap of the material ($E_G = E_{\mathrm{C}} - E_{\mathrm{V}}$).

2.2.2 Charge Transport

Charge transport within semiconductors influences many aspects of Mixed–Mode PAD design and performance. Examples include the relationship between the high

voltage bias on the detector diode and the resolution of the Mixed–Mode PAD (section 6.3) as well as the degradation of voltages held in the analog correlated double sampling and sample and hold circuits (section 4.3.2). A discussion of the basic charge transport properties of semiconductors follows.

The derivation of the charge transport equations within a semiconductor begins with the current density ($\vec{J}_{\rm e}$ and $\vec{J}_{\rm h}$) for electrons and holes, respectively, comprising drift and diffusion components that describe the flow of electrons and holes through the semiconductor [93],

$$\vec{J}_{\rm e} = q\,(\overbrace{\mu_{\rm e}\vec{\mathcal{E}}n}^{\text{drift}} + \overbrace{D_{\rm e}\vec{\nabla}n}^{\text{diffusion}}), \tag{2.3}$$

$$\vec{J}_{\rm h} = q\,(\mu_{\rm h}\vec{\mathcal{E}}p - D_{\rm h}\vec{\nabla}p), \tag{2.4}$$

where q is the fundamental unit of charge, n and p are the electron and hole densities (resp.), $\mu_{\rm e}$ and $\mu_{\rm h}$ are the electron and hole mobilities (resp.), $D_{\rm e}$ and $D_{\rm h}$ are the diffusion constants for electrons and holes (resp.), and $\vec{\mathcal{E}}$ is the electric field. Local continuity requires that

$$\frac{dn}{dt} + \vec{\nabla}\cdot\vec{J}_{\rm e} = \overbrace{G_{\rm e}}^{\text{generation}} - \overbrace{\frac{n}{\tau_{\rm e}}}^{\text{recombination}}, \tag{2.5}$$

$$\frac{dp}{dt} + \vec{\nabla}\cdot\vec{J}_{\rm h} = G_{\rm h} - \frac{p}{\tau_{\rm h}}, \tag{2.6}$$

where $G_{\rm e}$ and $G_{\rm h}$ are the electron and hole generation rates (resp.) and $\frac{n}{\tau_{\rm e}}$ and $\frac{p}{\tau_{\rm h}}$ are the electron and hole recombination rates (resp.), with $\tau_{\rm e}$ and $\tau_{\rm h}$ the electron and hole lifetimes (resp.). Combining these results gives the basic equations governing charge transport in both depleted and undepleted regions of the semiconductor,

$$\frac{dn}{dt} = \overbrace{D_{\rm e}\nabla^2 n}^{\text{diffusion}} + \overbrace{\mu_{\rm e}\vec{\nabla}\cdot\left(\vec{\mathcal{E}}n\right)}^{\text{drift}} + \overbrace{G_{\rm e}}^{\text{generation}} - \overbrace{\frac{n}{\tau_{\rm e}}}^{\text{recombination}}, \tag{2.7}$$

$$\frac{dn}{dt} = D_{\rm h}\nabla^2 p - \mu_{\rm h}\vec{\nabla}\cdot\left(\vec{\mathcal{E}}p\right) + G_{\rm h} - \frac{p}{\tau_{\rm h}}. \tag{2.8}$$

As equation 2.7 indicates, these equations are composed of four terms: a diffusion term, representing the thermal dispersion of non-equilibrium charge concentrations; a drift term, describing the flow of charge under the influence of an external electric field; a recombination term, representing the finite lifetime of these free carriers; and a generation term, describing the spontaneous thermal or photonic generation of free charge carriers.

A general, analytic, closed-form solution to these charge transport equations does not exist. Typically systems with inhomogeneous or time varying electric fields require simulations to be accurately modeled. Detailed simulators have been developed [22] along with analytical approximations [13]. That said, there are a number of illustrative special cases for which this problem may be solved analytically. The remainder of this section considers a set of these that will be useful later in this thesis.

2.2.2.1 Generation and Recombination

If we assume that the generation rates for electrons and holes and the lifetimes of these particles are homogeneous constants, then, if u is a solution to the homogeneous continuity equation, *i.e.* equation 2.5 or 2.6 with generation and recombination terms set to zero, one can show that

$$u' = u \exp\left\{-\frac{t}{\tau_{\mathrm{e/h}}}\right\} + G_{\mathrm{e/h}}\tau_{\mathrm{e/h}}, \tag{2.9}$$

is a solution to the inhomogeneous equation.

The electron and hole generation and recombination terms can vary depending on the temperature and local energy band structure. However, as long as the temperature is held stable and the local energy band structure remains constant, this solution will be locally valid. As these constraints are expected to hold in the cases of interest for this thesis, we will neglect generation and recombination terms

in the remaining examples—with the understanding that they may be reintroduced to any subsequent solution by way of equation 2.9.

2.2.2.2 Pure Diffusion

Setting $\vec{\mathcal{E}} = 0$ in the generation/recombination-free form of the charge transport equations, 2.7 and 2.8, yields the pure diffusion equation,

$$\frac{du}{dt} = D\nabla^2 u, \tag{2.10}$$

where D is the diffusion constant. If we choose our initial condition to be $u(\vec{x}, 0) = \delta(\vec{x})$, the well known solution to this partial differential equation (PDE) is

$$u(\vec{x}, t) = \frac{1}{(\pi 4Dt)^{\frac{3}{2}}} \exp\left\{-\frac{\vec{x}^2}{4Dt}\right\}, \tag{2.11}$$

representing a Gaussian sphere of charge with RMS extent $\sqrt{6Dt}$. As it is possible to write an arbitrary initial distribution of charge $q_0(\vec{x})$ as

$$q_0(\vec{x}) = \int d^3x'\, q_0(\vec{x}')\delta(\vec{x}' - \vec{x}), \tag{2.12}$$

the linearity of the integration and differentiation operators allow us to use our preceding result as a propagator to determine the time evolution of an arbitrary initial state ($q(\vec{x}, t)$). Thus,

$$\begin{aligned} q(\vec{x}, t) &= \int d^3x'\, q_0(\vec{x}')u(\vec{x}' - \vec{x}, t) \\ &= \int d^3x'\, \frac{q_0(\vec{x}')}{(\pi 4Dt)^{\frac{3}{2}}} \exp\left\{-\frac{(\vec{x}' - \vec{x})^2}{4D_{\mathrm{e/h}}t}\right\}. \end{aligned} \tag{2.13}$$

A complete and detailed derivation of this result may be found in numerous sources, such as [55].

2.2.2.3 Diffusion within a Constant Electric Field

This example is arguably the most useful result that we will derive as it offers a good means by which to estimate the movement of non-equilibrium charge

(*e.g.* charge generated by x-ray conversion) in undepleted silicon and provides a basis for approximating the behavior in depleted silicon.

To begin, note that, if we neglect recombination and generation terms, the charge transport equations, 2.7 and 2.8, take the form

$$\frac{dn}{dt} = D_{\mathrm{e}}\nabla^2 n + \mu_{\mathrm{e}}\vec{\nabla}\cdot\left(\vec{\mathcal{E}}n\right), \tag{2.14}$$

where for brevity we only explicitly present the electron results, as the hole result is analogous. The influence of the arbitrary external electric field makes this problem very difficult to generally solve in closed form. However, if we stipulate that $\vec{\nabla}\cdot\vec{\mathcal{E}} = 0$ and $\frac{d\vec{\mathcal{E}}}{dt} = 0$ the problem reduces to

$$\frac{dn}{dt} = D_{\mathrm{e}}\nabla^2 n + \mu_{\mathrm{e}}\vec{\mathcal{E}}\cdot\vec{\nabla}n, \tag{2.15}$$

then we find ourselves presented with a PDE that has the form of the convection-diffusion equation,

$$\frac{du}{dt} = D\nabla^2 u + \vec{c}\cdot\vec{\nabla}u. \tag{2.16}$$

To solve this PDE, let $u(\vec{x},t)$ be a solution to equation 2.11, the pure diffusion problem. Employing the change of variables $\vec{x} \to \vec{x}' = \vec{x} + \vec{c}t$ and expanding the total time derivative of $u(\vec{x}',t)$ into its partials gives

$$\begin{aligned}\frac{d}{dt}u(\vec{x}',t) &= \frac{\partial}{\partial t}u(\vec{x}',t) + \frac{\partial}{\partial \vec{x}'}u(\vec{x}',t)\frac{\partial \vec{x}'}{\partial t} \\ &= \frac{\partial}{\partial t}u(\vec{x}',t) + \vec{c}\cdot\frac{\partial}{\partial \vec{x}'}u(\vec{x}',t) \\ &= \frac{d}{dt}u(\vec{x},t) + \vec{c}\cdot\vec{\nabla}u(\vec{x},t).\end{aligned} \tag{2.17}$$

Since, by assumption, $\frac{d}{dt}u(\vec{x},t) = D\nabla^2 u(\vec{x},t) = D\frac{\partial^2}{\partial \vec{x}^2}u(\vec{x},t) = D\frac{\partial^2}{\partial(\vec{x}')^2}u(\vec{x}',t) = D\nabla^2 u(\vec{x}',t)$, substituting this result into equation 2.16 shows that our change of variables is sufficient to turn a solution of the pure diffusion equation into a solution to the convection-diffusion equation. For the particular problem of electron and

hole transport within a constant electric field, inspection shows us that $\vec{c} = \mu_e \vec{\mathcal{E}}$ and $\vec{c} = -\mu_h \vec{\mathcal{E}}$, respectively. Given an initial change distribution $q_0(\vec{x})$ we may use our previous result, equation 2.13, to determine the time evolution of the distribution,

$$q(\vec{x}, t) = \int d^3x' \, \frac{q_0(\vec{x}')}{(\pi 4 D_e t)^{\frac{3}{2}}} \exp\left\{ -\frac{(\vec{x}' - (\vec{x} + \mu_e \vec{\mathcal{E}} t))^2}{4 D_e t} \right\}, \tag{2.18}$$

which can be interpreted as a collection of Gaussian spheres of charge drifting at a constant rate of $\mu_e \vec{\mathcal{E}} q$ while expanding through diffusion to an RMS size of $\sqrt{6 D_e t}$. An analogous treatment applies to holes.

2.2.2.4 Diffusion in a Linear Electric Field

Unfortunately equation 2.18 is not directly applicable to the problem of transport in the depletion zone of a reverse biased junction diode as the fields in this region do not meet the condition $\vec{\nabla} \cdot \vec{\mathcal{E}} = 0$ due to the presence of ionized dopants exposed by depletion. Within this region, under the uniform doping approximation which we will discuss in section 2.2.3.1, the electric field increases linearly with depth into the depletion layer [93]. A complete analytical evaluation of the charge transport equations for this case is quite difficult and arguably unnecessary for the case of greatest interest—the detector diode. Here our main concern is the spatial distribution of charge carriers at the pixelated side of the diode (since integration removes all temporal information). The results from the preceeding section are still valid in the horizontal plane of the detector diode, perpendicular to the field lines, so we may estimate the yield profile in these dimensions as a 2D Gaussian where t in equation 2.18 is the mean transit time for a charge carrier through the diffusion region under the influence of drift alone, given by:

$$t_{\text{drift}} \approx \frac{w_{\text{dio}}^2}{\mu V_{\text{dio}}} \ln\left\{ \frac{w_{\text{dio}}}{d} \right\}, \tag{2.19}$$

where w_{dio} is the width of the depletion region, V_{dio} is the potential drop across this region, μ is the mobility of the carrier type, and d is the depth into the region

at which the charge is released.

2.2.3 Basic Semiconductor Devices

There are two fundamental devices that need to be described in order to understand the design of the Mixed–Mode PAD and the effects, both wanted and unwanted, of radiation on this detector. These are:

- The P/N or diode junction.
- The MOS (Metal-Oxide-Semiconductor) capacitor.

Through the combination of these two basic components, one is able to build the CMOS transistor which, in turn, is used to construct the complex electronics, such as op-amps, comparators, counters and registers, that make up the Mixed–Mode PAD.

2.2.3.1 P/N Junction Diode

The operation of the P/N junction diode is well known and detailed discussions can be found in introductory semiconductor physics and analog electronics texts [14, 97, 93, 38]. Here, therefore, we present only a brief review of the P/N junction diode—oriented towards topics that will be needed later. In particular, this discussion is limited to the reverse bias mode of diode operation only, as this is the primary case of importance for the Mixed–Mode PAD. Readers interested in a more complete discussion of this device are directed to the references noted above.

A P/N junction diode is created when two semiconductor regions of differing type share a common boundary. Considered separately, the Fermi Level (E_{F}) of each region is located in a different portion of the band gap, on opposing sides of the intrinsic energy level (E_i). When a junction exists, however, a condition for

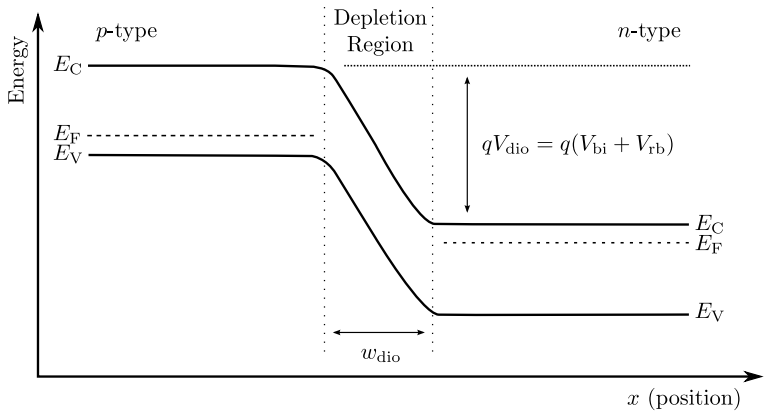

Figure 2.3: Energy band diagram of a reverse biased P/N junction.

static equilibrium is that the Fermi Level must be flat throughout the material. To accomplish this, first, in the absence of an applied voltage, majority carriers within each region diffuse towards zones of smaller concentration in the opposing region, leaving behind space charge in the form of the immobile ionized dopants. Following Poisson's Equation this space charge will result in a static electric field that, assuming that the x axis is normal to the boundary separating the two semiconductor regions, is given by

$$\frac{\partial^2 \psi(x)}{\partial x^2} = \frac{\partial \mathcal{E}(x)}{\partial x} = -\frac{\rho_{\text{free}}(x)}{\epsilon_{\text{Si}}}, \tag{2.20}$$

where $\psi(x)$ is the potential accompanying the field that leads to a bending of the bands within the junction region, as illustrated in figure 2.3. Because this region has a smaller majority carrier concentration than far away from the junction it is called the depletion region of the diode. The total potential shift induced by the junction is known as the built-in potential of the diode (V_{bi}). If an external potential (V_{rb}) is applied in a manner so as to add to the potential difference in the direction of the diode's built-in potential, then this will result in further enlargement of the space charge region and further band bending. Effectively, this potential increases the separation of the bands on either side of the junction by

qV_{rb} so that the total band separation becomes $qV_{\mathrm{dio}} = q(V_{\mathrm{bi}} + V_{\mathrm{rb}})$.

A parameter that is of great interest to PAD designers is the width of the depletion region within the detector diode layer, as this value is needed to know how to bias the detector diode to get full depletion with minimal parasitic leakage. Under the typical diode slab model, where the doping concentration is assumed to be uniform within the region on either side of the junction and undergo a discontinuous change at the junction, the width of the depletion region (w_{dio}) may be approximated as [14]

$$w_{\mathrm{dio}} = \left[\left\{ \frac{2\epsilon_{\mathrm{Si}} V_{\mathrm{dio}}}{q} \right\} \left\{ \frac{N_{\mathrm{A}} + N_{\mathrm{D}}}{N_{\mathrm{A}} N_{\mathrm{D}}} \right\} \right]^{\frac{1}{2}}, \tag{2.21}$$

where ϵ_{Si} is the permitivity of silicon, N_{A} is the density of acceptor dopants, and N_{D} is the density of donor dopants. For most diodes, in particular those used in the Mixed–Mode PAD detector layer, one side of the junction receives a substantially higher doping than the other. Supposing $N_{\mathrm{D}} \ll N_{\mathrm{A}}$ the formula for the width simplifies to

$$\begin{aligned} w_{\mathrm{dio}} &\cong \sqrt{\frac{2\epsilon_{\mathrm{Si}} V_{\mathrm{dio}}}{q N_{\mathrm{D}}}} \\ &= \sqrt{2\epsilon_{\mathrm{Si}} \mu_{\mathrm{e}} \rho_{\mathrm{e}} V_{\mathrm{dio}}}, \end{aligned} \tag{2.22}$$

where ρ_{e} is the resistivity of the n-type region.[3]

From the width of the diode region it is possible to estimate the capacitance of the diode (C_{dio}) via the simple formula

$$\begin{aligned} C_{\mathrm{dio}} &\cong \frac{\epsilon_{\mathrm{Si}} A_{\mathrm{dio}}}{w_{\mathrm{dio}}} \\ &= A_{\mathrm{dio}} \sqrt{\frac{\epsilon_{\mathrm{Si}}}{2\mu_{\mathrm{e}} \rho_{\mathrm{e}} V_{\mathrm{dio}}}}, \end{aligned} \tag{2.23}$$

where A_{dio} is the cross-sectional area of the diode. For integrating PADs, which will be described in section 3.2, this capacitance is a crucial parameter as it provides

[3]The Mixed–Mode PAD detector diode is fabricated on high resistivity silicon wafers, reported by the manufacturer to be 5 kΩ-cm to 10 kΩ-cm.

a path for the bias voltage of the detector to couple directly into the integration node of the pixel, with a coupling magnitude given by

$$\delta V_{\text{outp}} = \frac{C_{\text{dio}}}{C_{\text{int}}} \delta V_{\text{dio}}, \tag{2.24}$$

where C_{int} is the integration capacitance, explained in detail in section 4.2.1. Given C_{int}, C_{dio}, and a noise figure for global fluctuation in the array, this formula sets a limit on the fidelity required from the detector diode bias voltage, relative to the reference voltage supplied by the pixel.

2.2.3.2 The MOS Capacitor

Because the gate of each CMOS transistor is ostensibly a MOS capacitor, understanding how this device operates is a prerequisite to understanding the operation of the CMOS transistor. In addition, and more relevant to our application, the explanation of this device's operation underlies the explanation of how charge trapping within the ASIC surface oxide, the dominant x-ray induced long-term radiation damage mechanism, degrades the performance of CMOS devices. As this topic is very thoroughly covered in many other sources ([97, 93]), this section is limited to a review, highlighting important results that will be useful elsewhere within this thesis. Also, to simplify the discussion we will assume a p-type substrate, as the results discussed here are analogous for n-type substrates, with appropriate changes in the sign of relative potentials and charge.

The MOS capacitor may be likened to Neapolitan Ice Cream, consisting of three stacked layer: at its base is the substrate wafer; above this is a silicon-oxide passivation layer; and with the final layer, the gate, being either metallic (typically Al) or highly doped polysilicon. Phenomenologically, changing the gate potential alters the electric fields in the silicon di-oxide and the bulk silicon along with a depth dependent local potential $\psi(x)$, resulting in a redistribution of charge within

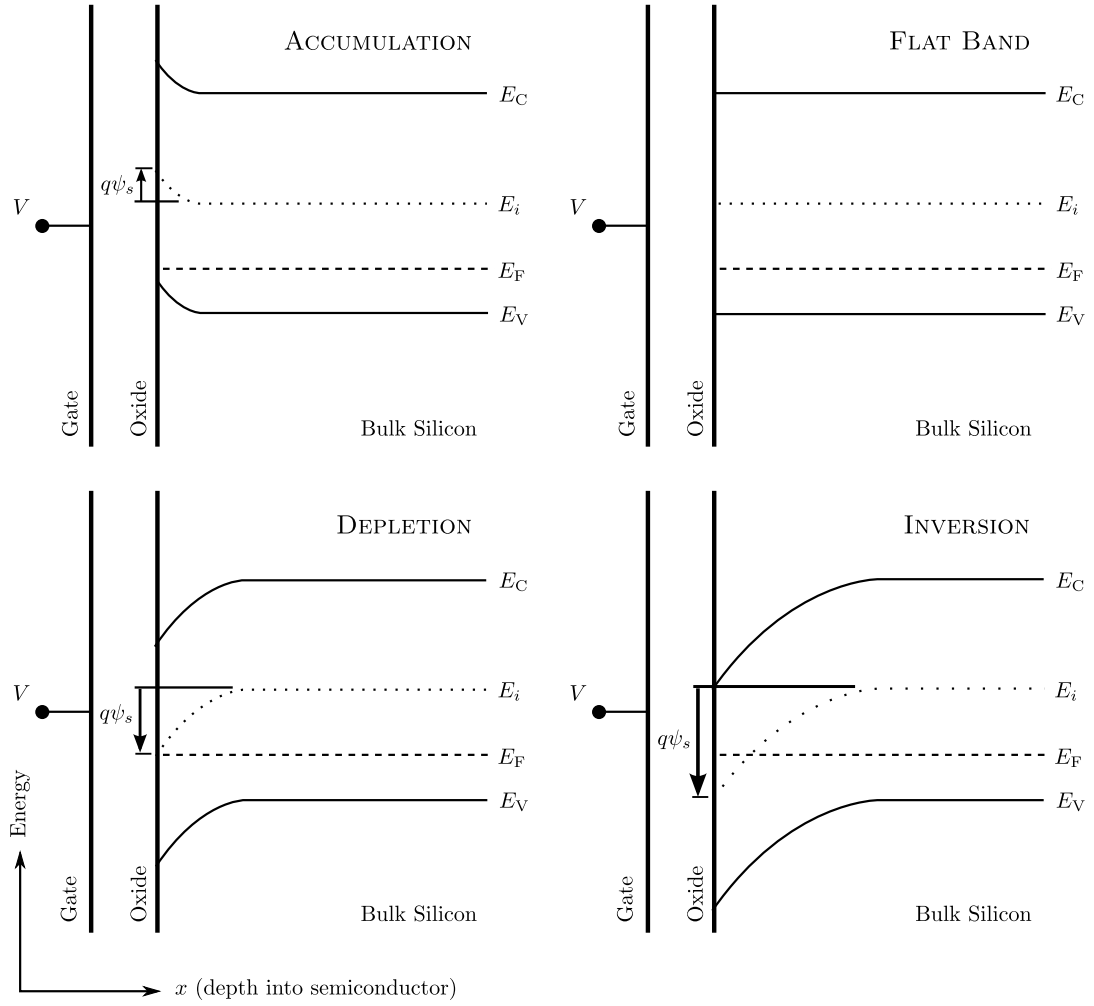

Figure 2.4: Semiconductor band diagrams depicting the accumulation, flat band, depletion, and inversion states of a p-type substrate. The parameter ψ_s represents the surface potential induced by the applied voltage.

the semiconductor. If the applied voltage (V) is written relative to the potential of the bulk silicon, then four different charge distributions, illustrated in figure 2.4, are distinguished:

Accumulation : $V < 0$, an excess of majority carriers (holes) is drawn to the region of bulk silicon beneath the silicon di-oxide passivation layer.

Flat Band : $V = 0$, an aptly named distribution as there is no deformation of the bands.

Depletion : $0 < V < V_{\mathrm{T}}$ (where $V_{\mathrm{T}} = \frac{E_{\mathrm{F}}}{q}$ is the threshold voltage of the device), the applied potential drives down the concentration of the majority carriers (holes), depleting it below the thermal equilibrium level.

Inversion : $V_{\mathrm{T}} < V$, once the silicon/silicon di-oxide interface and adjacent bulk regions are fully depleted, then an excess of minority charges are drawn into this region. When the concentration of native minority carriers exceeds the concentration of native minority carriers, which will occur when $E_{\mathrm{F}} < q\psi(x)$, then the region is said to be inverted.

More rigorously, following the treatment given in [93], the local potential ($\psi(x)$) at a depth x into the semiconductor is given by equation 2.20 where the local charge density ($\rho_{\mathrm{free}}(x)$) is given by the sum of the local free and space charge,

$$\rho_{\mathrm{free}}(x) = q\left[n(x) + N_{\mathrm{A}} - (p(x) + N_{\mathrm{D}})\right]. \tag{2.25}$$

The local free electron density ($n(x)$) and free hole density ($p(x)$) are in-turn given by:

$$n(x) = n_i \exp\left\{\frac{q(\psi(x) - \psi_{\mathrm{B}})}{kT}\right\}, \tag{2.26}$$

$$p(x) = n_i \exp\left\{\frac{q(\psi_{\mathrm{B}} - \psi(x))}{kT}\right\}, \tag{2.27}$$

where $\psi_{\mathrm{B}} = \lim_{x\to\infty} \psi(x)$ is the bulk potential. These equations (2.26 and 2.27) are slight modifications of the standard charge carrier concentration forms (equations 2.1 and 2.2), with $q\psi_{\mathrm{B}}$ replacing $E_{\mathrm{F}} - E_i$. This gives us a basic ordinary differential equation (ODE) whose boundary conditions, stipulating that the capacitor gate voltage (V) is referenced to the potential of the bulk silicon, are $\psi(x = 0) = V$ and $\lim_{x\to\infty} \psi(x) = 0$. These equations show us how the voltage applied to the MOS capacitor influences concentration of charge in the underlying silicon.

2.3 Radiation Effects in Silicon

For an x-ray detector, it goes without saying that one must understand the effect of radiation to properly characterize and use the device. Generally, there are three ways in which x-rays interact with matter [70]: 1) at the lowest energies, up to a few keV, interactions are predominantly through the photoelectric effect. In brief, this is when an x-ray interacts with an electron bound to an inner atomic orbital, resulting in the absorption of the x-ray, the creation of a free electron with kinetic energy nearly that of the incident x-ray as well as a vacancy (hole) in the electron's original binding location. 2) At higher energies Compton scattering begins to dominate, this is when the x-ray elastically scatters from a free or loosely bound electron, transferring a portion of its energy into kinetic energy of the electron and recoiling at a lower energy. 3) Finally, at the highest energies, above 1.022 MeV, production of electron–positron pairs becomes important. Over most of the spectral range where the Mixed–Mode PAD is designed to operate the photoelectric effect is the dominant interaction mechanism, though at the highest portion of this range the contribution from Compton scattering begins to be noticeable.

Both the photoelectric effect and Compton scattering produce highly energetic δ-electrons [56]. Depending on where within the hybrid the interaction occurs, the

effect of the resulting δ-electron will vary. Broadly, we can distinguish three cases:

- X-ray absorption within the depletion region of a P/N junction.
- X-ray absorption within undepleted silicon.
- X-ray absorption within an oxide passivation layer.

In the first two cases, interaction in depleted and undepleted silicon, the δ-electron produces a cloud of free charge carriers[4] through ionization of local silicon atoms. The quantity of charge yielded will be proportional to the x-ray energy (E_x) deposited in the detector as given by:

$$Q = qN = q\frac{E_x}{\xi}, \tag{2.28}$$

where Q is the generated charge, $N = \frac{E_x}{\xi}$ is the number of generated carriers, E_x is the deposited energy, and ξ is the material ionization energy (for silicon $\xi_{\mathrm{Si}} \approx 3.6$ eV) [56]. The statistics of this generation process have been found to be better that the $\sqrt{N}$ expected of a pure Poisson process, presumably due to correlations in the generation process [29]. To account for this, a term known as the Fano Factor ($\mathcal{F}_{\mathrm{a}}$) is incorporated into the Poisson noise formula (RMS) so that

$$\delta Q = q\sqrt{\mathcal{F}_{\mathrm{a}} N}, \tag{2.29}$$

with $\mathcal{F}_{\mathrm{a}} \approx 0.1$ in silicon [51, 101].

While the production of free carriers is very similar in depleted and undepleted silicon, the impact of these carriers is markedly different. Two points in particular distinguish the two cases. First is the matter of majority carriers. In undepleted silicon, doping concentrations typically raise the level of free carriers of the dominant species (*i.e.* holes in p-type material or electrons in n-type) to a

[4]The term free charge carriers is being used somewhat loosely here as shorthand to denote a conduction band electron or a valance band hole, as opposed to a carrier with kinetic energy greater than the work function of the material.

point that additional photo-ionization charge is insignificant by comparison. As a result, only the minority carriers, which are suppressed under these circumstances, may be detected and measured. In depleted silicon, there are effectively no free carriers so that photo-signal from both species may, in principle, be observed. The second point involves the charge transport mechanism dominant in these two cases. In undepleted silicon, the electric fields are typically weak or non-existent so the dominant charge transport mechanism is diffusion. In contrast, depleting silicon requires electric fields to sweep away free charge carriers. These fields may either occur naturally as with the built in field of an unbiased P/N junction or be imposed externally. The presence of these fields shifts the predominant charge transport mechanism from carrier diffusion to carrier drift. The phenomenological difference between these two charge transport mechanisms was discussed in section 2.2.2.

When a δ-electron is produced within the oxide passivation layer of the hybrid the result is very different. Here, the substantially greater separation of the conduction and valance bands effectively eliminates thermally generated free carriers. As the mobility of electrons is relatively large (20 $cm^2/V{\cdot}s$ at room temperature) in comparison with that of holes (1.6×10^{-5} $cm^2/V{\cdot}s$ at room temperature), the presence of an electric field will rapidly sweep free electrons out of the oxide. Free holes, on the other hand, will undergo much more gradual drift, often becoming trapped at defect sites either in the oxide bulk or at the oxide/silicon interface [16]. As the lifetime of these trapped states can be quite long, though this is strongly dependent on the temperature [8], over time this can lead to a charging of the oxide that detrimentally affects the performance CMOS devices [3].

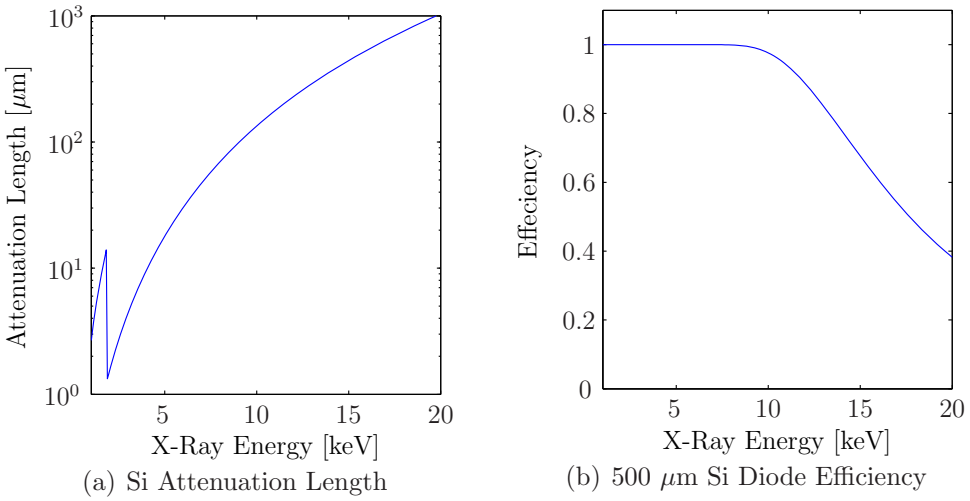

(a) Si Attenuation Length

(b) 500 μm Si Diode Efficiency

Figure 2.5: Absorption properties of silicon. Panel (a) shows absorption length as a function of energy. Panel (b) show the relative absorption efficiency of a 500 μm detector diode layer to normally incident x-rays. Values for these plots were obtained from the Berkeley Lab, Center for X-Ray Optics web site (www.cxro.lbl.gov), which in turn cites [46].

2.3.1 X-Ray Detection in the Mixed–Mode PAD

Hereto we have spoken about the effects of an x-ray conversion within various areas of the PAD hybrid. Now, we shift our focus to the likelihood of this conversion and, if it does, its likely location. To offer context for this discussion, we consider the problem in terms of the detector diode. To an x-ray the diode is a block of silicon, effectively indistinguishable from the rest of the hybrid, but from our perspective this block of silicon has a specific purpose—to convert the x-ray into charge and convey the charge to the signal processing electronics.

The detector layer of the Mixed-Mode PAD is effectively a monolithic P/N diode. It is made from a 500 μmthick silicon diode with a light n- doping ($\sim 10^{11}$ to 10^{12} donors per cm^3 [26]). The face of the detector, towards the x-ray source, receives a heavy $n+$ doping along with aluminization to provide an evenly distributed bias voltage, while the back of the detector receives pixelated $p+$ doping, correspondent with pixels on the Application Specific Integrated Circuit (ASIC).

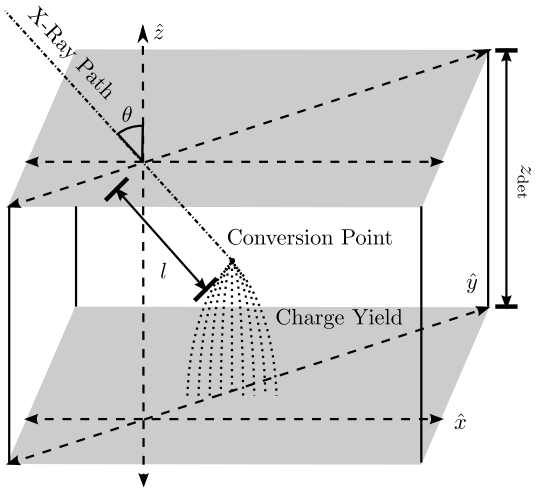

Figure 2.6: Model used for calculating the charge yield profiles of a monochromatic x-ray beam incident on a fully depleted silicon diode. The $y = 0$ plane is defined by the vertical plane containing path of the x-ray, while the the $x = 0$ plane is defined to be the vertical plane perpendicular to the $y = 0$ plane, containing the point where the x-ray enters the diode.

When fully depleted, this architecture produces a vertical electric field that sweeps charge carriers generated by x-ray conversion to the pixel integration stage.

The fraction of an x-ray beam of flux Φ absorbed in the Mixed–Mode PAD detector diode layer is dependent on the energy and angle of incidence (θ) relative to the surface normal. Generally, the probability that an x-ray will convert a distance l into the detector layer is given by

$$P(l, E) = \frac{dl}{\lambda(E)} \exp\left\{-\frac{l}{\lambda(E)}\right\}, \tag{2.30}$$

where $\lambda(E)$ is the x-ray absorption depth at the energy E. Figure 2.5 summarizes the x-ray absorption properties of the Mixed–Mode PAD detector diode layer and their dependence on energy.

As mentioned, the purpose of the detector diode is not only to absorb x-rays but to convey the resulting charge carriers to the pixels for measurement. For a monochromatic beam of x-rays, the profile of charge measured by the detector will depend on the incident angle of the beam and the energy of x-rays within it

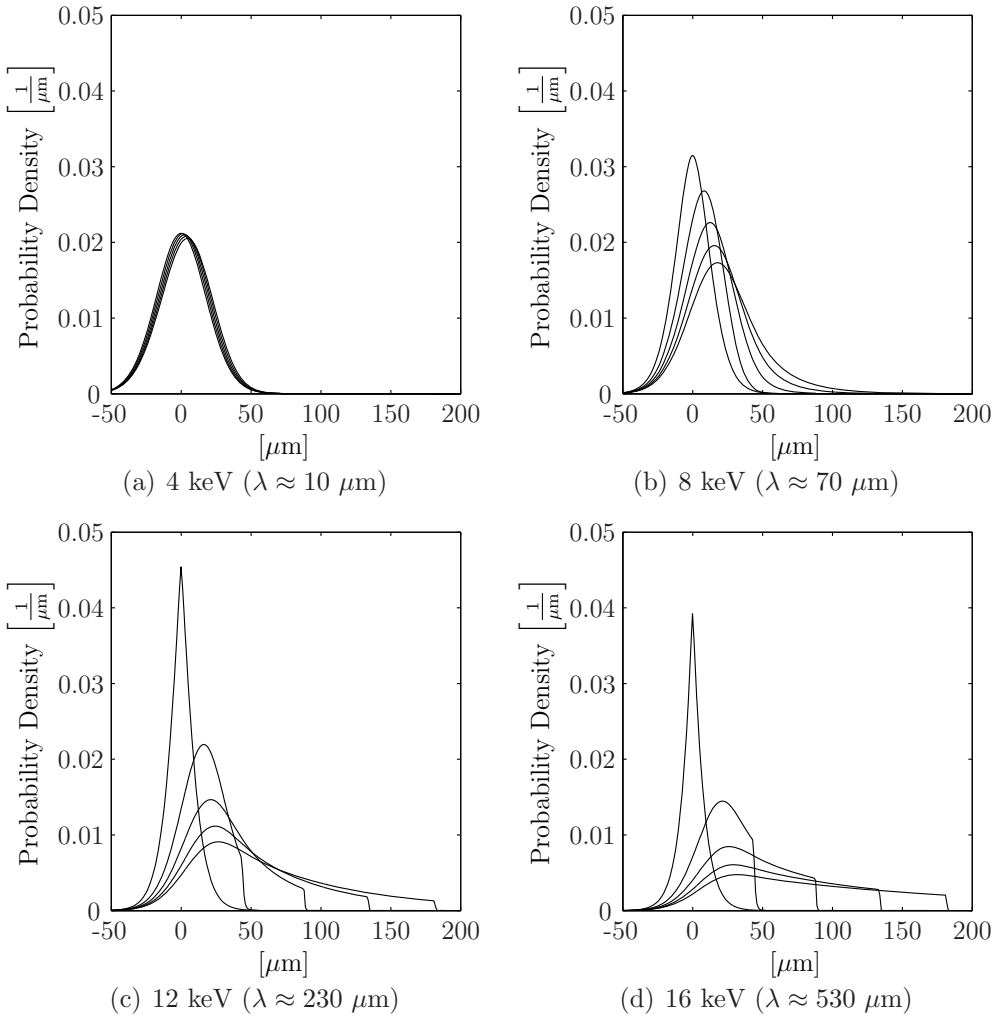

(a) 4 keV ($\lambda \approx 10$ μm)

(b) 8 keV ($\lambda \approx 70$ μm)

(c) 12 keV ($\lambda \approx 230$ μm)

(d) 16 keV ($\lambda \approx 530$ μm)

Figure 2.7: Charge yield profiles of monochromatic x-ray beams of differing energies at incidence angles of 0, 5, 10, 15, and 20 deg. from the surface normal of a 500 μm detector diode. The cutoffs exhibited in the 12 keV and 16 keV plots, panels (c) and (d), at high incidence angles are due to x-rays passing completely through the 500 μm thick diode layer.

in a manner that is relatively straightforward to calculate. From equation 2.30, we know the probability of an x-ray converting a distance l into the detector diode. As discussed earlier, this conversion results in a δ-electron that initiates a complex cascade of secondary ionizations. It is generally assumed that the resulting distribution of charge carriers may be described as a three-dimensional Gaussian,

$$q_0(\vec{x}) = \frac{Q_0}{\pi^{\frac{3}{2}}\sigma_r^3} \exp\left\{-\frac{(\vec{x}-\vec{x}_0)^2}{\sigma_r^2}\right\}, \tag{2.31}$$

where Q_0 is the total charge yield, $\vec{x}_0$ is the coordinate of the distribution center of mass, and $\sigma_r \approx 0.012(E/1\text{ keV})^{1.75}$ μm is a characteristic initial width of the cloud [75]. From the discussion in section 2.2.2 we have a means to estimate the time evolution of this charge profile as it moves through the depletion zone of the detector diode. Combining equations 2.31, 2.30, and 2.18 gives:

$$p(x,y,E,\theta) = \int_0^{\frac{z_{\text{det}}}{\cos(\theta)}} \frac{q_0(\vec{x}_0)dl}{\lambda(E)\pi\sigma(l,\theta)} \exp\left\{-\frac{(x-l\sin\theta)^2+y^2}{[\sigma(l,\theta)]^2} - \frac{l}{\lambda(E)}\right\}, \tag{2.32}$$

where

$$\sigma(l,\theta) \cong \sqrt{4kT\left(\frac{z_{\text{det}}^2}{V_{\text{dio}}}\right)\ln\left[\frac{z_{\text{det}}}{l\cos(\theta)}\right]}, \tag{2.33}$$

with all parameters as defined in figure 2.6 and equation 2.31. In figure 2.7, we use this result to depict the normalized charge profile generated by monochromatic beams of x-rays incident on the detector in the same location at varying angles of incidence.

2.3.2 Radiation Damage

Conversion of x-ray photons into electrical charge within the detector diode is a basic principle allowing PADs to operate. Conversion in other regions of the hybrid, *e.g.* the diode depletions surrounding each transistor diffusion or within the undepleted bulk substrate, tend rather to have detrimental effects on the detector. Regarding photon induced radiation damage, two categories are typically

distinguished: single event effects, resulting from unintended x-ray conversion in sensitive portions of the ASIC layer; and dose dependent effects that are the result of the long term accumulation of damage within the detector [70].

2.3.2.1 Single Event Effects

Single event effects are exhibited as a change of state within the detector electronics as a result of the burst of photocurrent that accompanies x-ray conversions within the ASIC layer. Two types of effects are commonly distinguished: Single Event Latchup (SEL) where a highly ionizing particle deposits enough charge in a small volume to activate a parasitic thyristor [3]; and Single Event Upset (SEU) which is manifest through an abrupt increase in the voltage of a node, or possibly nodes, in the circuit resulting in a change of state within the device. SEL errors are quite dramatic, resulting in the circuit entering into an inoperable state of high current draw, potentially destroying the device [70]. Fortunately, if guard rings and substrate contacts are used extensively throughout the ASIC layout then the threshold for these errors becomes quite high, with tolerance up to 89 MeV-cm^2/mg reported by some sources in a 0.25 μm process [19]. SEU errors are more prevalent and problematic with, as we will discuss, the potential to affect both digital and analog circuitry.

SEU errors occur in digital circuits when an x-ray converts sufficiently close to a transistor diffusion that the voltage on the node, at least temporarily, rises high enough to change the logic state. This form of radiation damage can rewrite registers, open gates, initiate logic sequences, and generally wreak havoc in a digital design [70]. As with most radiation damage mitigation, deciding how to suppress this damage involves a complicated matrix of design considerations. On one hand, there are digital logic circuits that are designed specifically to perform robustly in high-radiation environments. These architectures, however, entail a significant

cost in terms of area and design complexity [3, 70]. At a lower level of robustness, there are certain logic families that offer improved protection (*i.e.* static logic, where active elements retain logic states, as opposed to dynamic logic, where logic states are held on capacitive elements that are periodically refreshed). While much more compact than rad-hard logic circuits, there is still a trade off in terms of area and power consumption for improved radiation tolerance. Particularly, due to the limited area available within each pixel, care must be taken to weigh the costs and benefits of these measures.

Regarding analog components the impact of single events is more subtle due to the substantially larger capacitance typically found within these circuits and the absence of a binary state. The result of these effects is seen primarily in the capacitive storage elements within the pixel and comes in the form of increased leakage from switches connected to these nodes where charge is stored. The significance of this current depends on the rate of x-ray conversion within the bulk silicon, the area over which this charge may be collected, and the capacitance of these nodes. While there is little one can do to affect the rate of x-ray conversion, apart from operating at lower x-ray energies or finding a more efficient diode material, there are steps that may be taken to mitigate the other effects.

The simplest mitigation measure is to increase the capacitance associated with the analog storage elements. As previously noted, a 10 keV x-ray will yield roughly 0.5 fC of charge. If a 50 fF storage element is used and all this charge is collected then the effect on the output is a 1% shift. Increasing the storage element capacitance 500 fF reduces this to a 0.1% effect. Where area and circuit bandwidth limitations allow, increasing the capacitance of critical nodes is an effective, though brute force, means to reduce single event effects. A much more elegant approach that works well in conjunction with increasing the capacitance is to limit the col-

lection range for single event effects through the use of $p-$ or n-wells. The fields within the reverse-diode junction between the well and the substrate permits only carriers of the substrate minority type to cross into the well. These carriers are, for the same reason, prevented from passing through the reverse-diode junction between well and transistor diffusion connected to the analog storage element. As a result, the effective collection area for single event effects is reduced to the area of the well. Some processes, such as the TSMC 0.25 μm, offer a deep n-well option that allows a designer to imbed a large n-well into a p-type substrate in which further p-wells may placed for fabrication of isolated n-type devices. Where this option is not available, designers must rely on single type switches.

2.3.2.2 Long-Term Damage

Photons are, fortunately, much more forgiving, in terms of the types of radiation damage that they induce, than massive particles. Unlike particles with mass, x-rays do not induce persistent changes in the bulk silicon [70]. Instead, their dominant long-term radiation damage mechanisms involve a gradual charging of ASIC oxide layers through the creation and trapping of holes [70] and increases in the leakage from the detector diode [14, 26]. These radiation damage effects cannot be neglected, as they have the potential to seriously degrade the detector performance if not mitigated through design [3]. Designing for radiation tolerance, however, must be undertaken with care as the most robust techniques to mitigate damage incur substantial costs in terms of area and may affect circuit performance in ways that are difficult to model.

The trapping of holes in the oxide layer has an effect analogous to changing the potential on the gate element of a CMOS capacitor, namely the trapped holes induce electrical fields, which, in turn, alters the charge distribution in the underlying silicon via the mechanism discussed in section 2.2.3.2. Depending on the

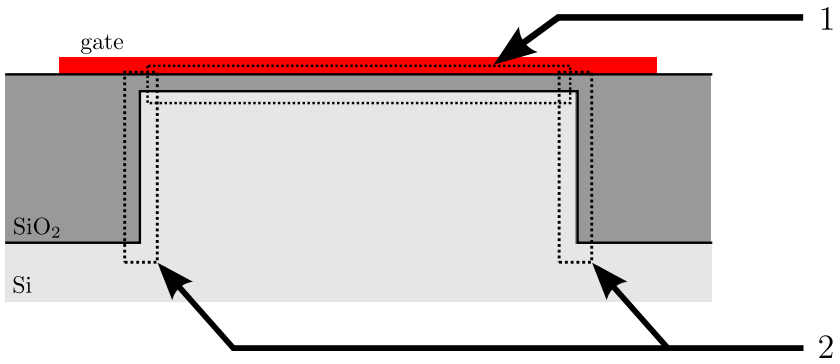

(a) Transistor Cross Section—*taken along dashed line in panel (b)*

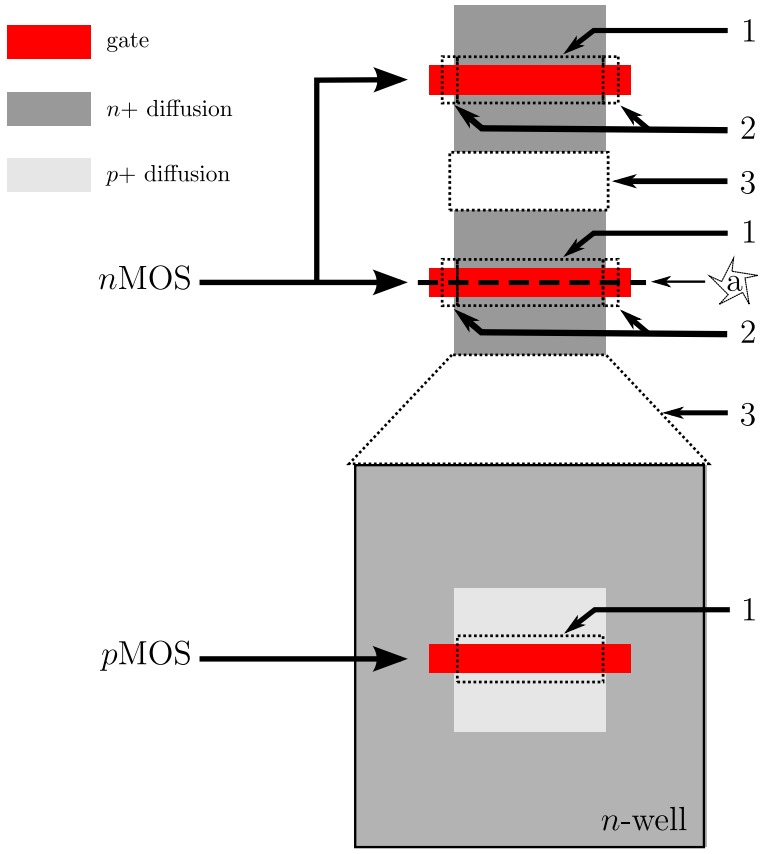

(b) Transistor Layout Top View

Figure 2.8: Illustration of a sub-μm CMOS layout (shallow trench isolation) indicating regions susceptible to radiation damage. Panel (a) depicts a transistor cross section taken along the dashed line indicated by the star encircled 'a' in panel (b), which, in turn, depicts the top view of a transistor layout. In both panels, region 1 denotes where ionization induces transistor threshold voltage shifts, region 2 denotes where ionization induces the formation of parasitic transistors between the source and drain diffusions of a nMOS device, and region 3 denotes where ionization induces the formation of parasitic transistors in the field oxide.

location of this oxidation damage, these holes can either act as an effective gate, applied commensurate with the actual gate of a transistor, or induce the formation of parasitic transistors.

The flavor of silicon substrate (*i.e* n-type vs. p-type) also has a significant effect on the consequences of long term damage. In n-type silicon, accumulation of holes trapped in the oxide draws majority carriers, electrons, to the surface of the substrate. With a p-type substrate, the situation is reversed: here accumulating damage repels majority carriers, driving the substrate under the field oxide towards depletion and inversion. This discussion is expanded upon in figure 2.8, which distinguishes three oxide regions susceptible to radiation damage in a deep sub-micron process, such as the 0.25 μm process from Taiwan Semiconductor Manufacturing Company (TSMC) used to fabricate the Mixed–Mode PAD: 1) the oxide separating the gate from the channel in both nMOS and pMOS devices; 2) the oxide forming the passivation boundary at the edge of each nMOS transistor, beneath the gate; and 3) the field oxide over the p-type substrate between n-type diffusions from nMOS devices or n-wells.

In the first region, the effects of radiation damage are typically characterized as shifts in the threshold voltage (V_{T}) of the transistor. For nMOS devices, this shift effectively lowers the threshold, resulting in increased channel current for a given applied gate voltage. For pMOS devices this shift effectively raises the threshold, leading to decreased channel current for a given applied gate voltage. As will be discussed in more detail in section 6.7, the relationship between ionizing dose and consequent threshold shift is non-trivial and strongly influenced by the presence and strength of fields within the oxide, the thickness of the gate oxide, as well as the temperature of the transistor.

In the second region, nMOS transistors are profoundly affected by the passiva-

tion technique used to isolate the deep sub-micron transistor structures. Known as Shallow Trench Isolation (STI), this technique is commonly used in processes with 0.25 μm or smaller feature size. It differs markedly from the LOCal Oxidation of Silicon (LOCOS) isolation used in larger-feature-size technologies and discussed at length in previous PAD theses [83, 26]. In this technique, passivation trenches descending 300-500 nm, with nearly vertical sides, are carved into the silicon to form a boundary around each transistor, as depicted in the transistor cross section depicted in panel (a) of figure 2.8. Along this surface, electric fields develop as dose accumulates in the oxide, inducing the formation of parasitic channels between the transistor source and drain. This effect is most significant at the corner where the boundary oxide of the transistor meets the transistor gate oxide, because the increased oxide volume relative to silicon volume, in this area, results in higher fields with accumulating dose. These parasitic channels, that form along the edge of the main channel between its source and drain, are typically modeled as independent transistors, in parallel with the main transistor. Unfortunately, there are no direct means to control these parasitic transistors with the gate voltage of the main device. Consequently, as we will see later in this section, they can substantially degrade the performance of, or even render useless, common circuit elements, such as the simple current mirror.

In the third and final region, parasitic transistors may form in the field oxide over any oxide covered p-type material, most commonly the substrate, between any n-type diffusions (*i.e.* the source or drain diffusions of nMOS devices or an n-well) of differing potentials. These parasitic transistors are known as field-oxide transistors or FOX transistors and tend to be less significant than the parasitic edge transistors that form in the second region. There are multiple reasons for this. First, the effect of these parasitic devices can be partially or completely mitigated

through the placement of substrate contacts within a layout, so that basic practices of good mixed-signal layout will protect against this problem. Second, minimal transistor spacing rules cause these devices to be relatively long. Finally, the structure of the STI passivation results in a somewhat circuitous path that current must follow in these parasitic devices. Despite the typically lower significance, consideration of the effect of parasitic field-oxide transistors should not be neglected as it is easy to lay out structures in which these parasitic devices can become quite problematic.

As with transient radiation damage, the methods commonly used to mitigate long-term damage span a spectrum of trade-offs between effectiveness and cost, in terms of area, power, and complexity. At one extreme, it is possible to use special, rad-hard foundries such as those operated by Lockheed-Martin in Manassas, VA or the Honeywell Solid State Electronics Center in Plymouth, MN. These foundries, though, tend to be hard to gain access to and produce devices with lower yield and performance characteristics roughly one generation behind the current generation of commercial CMOS electronics [19]. Fortunately, it has been shown that a large degree of radiation tolerance may be attained by selecting an appropriate process for the design [58, 92]. Studies of commercial CMOS processes carried out as part of the Large Hadron Collider (LHC) detector development effort observed that radiation tolerance effects first seen in thin MOS capacitors [43] extended to processes with smaller feature sizes [19]. The effect is explained by noting that, as feature size decreases, the thickness of the gate oxide decreases as well. Decreasing oxide thickness will improve device radiation tolerance as there will be less material for x-rays to convert in, with the shift in threshold voltage per Mrad deposited in the oxide falling off roughly proportional to the square of the oxide thickness. However, below a gate oxide thickness of $\sim$ 10 nm an additional

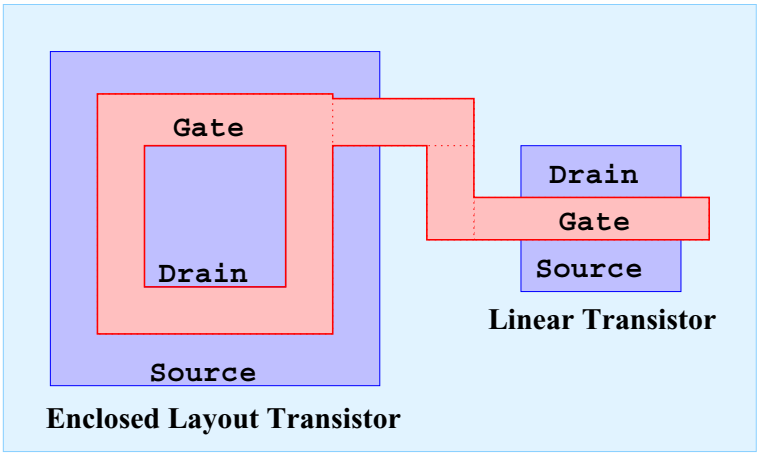

Figure 2.9: Illustration of an Enclosed Layout Transistor (ELT) in contrast to the traditional linear transistor.

reduction is gained [19].[5] This additional radiation hardness is thought to be due a reduction in the mean lifetime of holes trapped in the silicon di-oxide as a result of recombination with electrons tunneling into the oxide from either the gate or channel regions.

While process choice is very effective in mitigating radiation induced transistor threshold shifts, it is less effective in preventing the formation of parasitic transistors. Other methods must be employed to deal with these. The most costly, in terms of area consumed, and brute force methods available for standard CMOS technologies involves transistor layout techniques [92]. Most commonly a layout approach known as enclosed layout transistors (ELT), illustrated in figure 2.9, is used to prevent the formation of parasitic transistors by eliminating any paths between an nMOS device source and drain that are not controlled by a section of the transistor gate [2]. Use of these devices is not without cost. ELT devices tend to take up more area, require larger effective W/L ratios, and are more difficult to model than standard linear devices. The cost in terms of area of this layout approach along with the uncertainty it introduces into circuit modeling means that

[5]For context the gate oxide of a 0.25 μm process, as used in the Mixed–Mode PAD, is 5-7 nm.

its use within a PAD pixel design should be judicious.

More recently a new, linear, radiation-hard structure was reported in [92]. Transistors based on this report were designed and characterized by Dr. Alper Ercan, a former member of the Cornell PAD development group, and are presented in detail in his thesis [26]. As Dr. Ercan's vetting of the linear, radiation-hard transistor was not completed in time to incorporate them into any Mixed–Mode PAD submission, we will only review the basic idea of this transistor structure here. Readers interested in a more thorough discussion are referred to the two sources above.

The idea behind the linear, radiation-hard transistor is to remove the oxide at the edge of the nMOS device by encircling it with bulk silicon. This is done with standard fabrication design tools by extending the gate and active area to cover the region around the edge of the transistor, including surrounding the source and drain diffusions. Now, this structure is effectively a larger transistor than we desire, formed from the linear transistor directly between the source and drain diffusions, and secondary transistors at the edges of this device, due to the gate over the active boundary region. One gets back the performance of a linear transistor (*i.e.* performance that can be characterized by the W/L ratio between the source and drain) by using a $n+$ select masks in the desired active transistor region and a $p+$ select mask in the boundary region. For our purposes, the effect of the $n+$ select mask is to cause the polysilicon gate in the region it covers to receive a $n+$ doping. This doing induces a small positive gate voltage, effectively reducing the threshold voltage in the region directly between the source and drain diffusions. Similarly the $p+$ select mask causes the polysilicon in the region bounding the transistor to receive a $p+$ doping that acts as a small negative gate voltage, effectively increasing the threshold of the secondary devices in the boundary region. The difference between

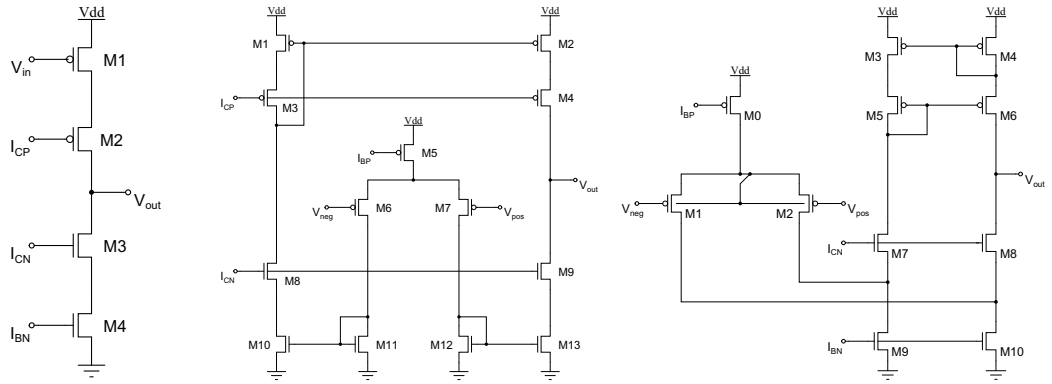

(a) Cascode (b) Mirror Cascode (c) Folded Cascode

Figure 2.10: Three amplifier architectures offering similar performance characteristics but drastically different levels of radiation tolerance. Amplifiers are ordered from left to right in order of increasing radiation tolerance.

these effective gate voltages, ∼1 V (the Si bandgap) assuming both the n-type and p-type polysilicon are heavily doped, acts to suppress channel formation outside of the region directly between the source and drain diffusions. As a result, for low gate voltages, this technique effectively yields a linear device that is radiation hard to 100s of Gy TID(SiO_2).[6] This device is not without its drawbacks, though. Foremost among these is the fact that, if the gate is raised high enough, then these secondary transistors in the boundary will become active at a significant level. This can present a leakage problem unless the region around the boundary gate is protected by a guard ring. Consequently, this layout technique enlarges the area required for the device by the bounding gate and guard ring at the perimeter of the device. In addition, the larger gate area increases the capacitance of the gate node while the secondary transistors distort the channel current vs. gate voltage curve of the device. Because of these effects, this layout technique is certainly not a silver bullet for the problem of radiation damage; however, as we will discuss in more detail in chapter 4, there are areas where it could find useful application.

An alternative approach to radiation hardness involves choosing circuit archi-

[6]See the discussion on units and dose in section 6.7 for an explanation of this terminology.

tectures that are inherently more resilient to long term radiation damage. This concept has long been used in radiation hard digital design, where combinatoric logic circuit elements (*e.g.* redundant latches with a majority voter logic output) are commonly used in radiation hardened designs to protect against single event upsets [70, 102]. While discussion is essentially absent in the literature, the idea of radiation tolerance through circuit architecture extends to analog devices as well. To illustrate this point, consider the three amplifier topologies shown in figure 2.10, which represent three generations of the front-end amplifier used in PADs produced by the Cornell PAD development group.[7]

The cascode structure shown in panel (a) of figure 2.10 is, by far, the most susceptible to radiation damage. Specifically, any threshold increase induced by radiation in transistor M1 will translate directly into a decrease in the quiescent voltage of the amplifier output. In addition, and probably more significant, increases in the channel current of transistor M4, the dual effect of radiation-induced lowering of the threshold of this device and parasitic channels formation around the edges of this device, will lead to greater current flow through the device at a given bias voltage, resulting in an additional lowering of the quiescent voltage of the amplifier output.

The mirror cascode architecture shown in panel (b) of figure 2.10 is much more radiation tolerant due to its use of a differential topology. Unlike the cascode architecture, the differential pairings in this device, and consequent common mode cancellation, means that only differences in accumulated dose effects will shift the quiescent output voltage of the amplifier. This particular architecture, however, suffers from its use of nMOS current mirroring stages (transistors M10, M11, M12, and M13). These current mirrors become, gradually, less effective as parasitic tran-

[7]The cascode architecture was used in the integrator for the original 100×92 μs imager [83], while the mirror cascode was used in the larger area redesign of the μs imager in the TSMC 0.25 μm process [26], while the folded cascode is used in the Mixed–Mode PAD.

sistors form along their edges, stealing current from the gate controlled portion of the device. This current loss degrades the transconductance of the diode connected component of the mirror, ultimately degrading the frequency response of the circuit and potentially leading to device instability under some feedback configurations. In principle this effect could be reduced by increasing the length of the transistors in the mirror. However, this will come at the expense of the amplifier frequency response, as any increase in L will either reduce the transconductance of transistors M11 and M12 or require a proportional increase in W thus increasing the capacitance of the gate node. Consequently, if circuit performance is not to be sacrificed, this design requires an ELT layout for true radiation hardness.

The final design, the folded cascode shown in panel (c) of figure 2.10, is the most intrinsically radiation tolerant as a result of its differential architecture and the fact that the two sets of nMOS devices present (transistors M7 and M8, along with M9 and M10) are both very tolerant to the formation of parasitic transistors. In the case of the pair of bias devices (transistors M9 and M10) they operate with a constant gate voltage and therefore may be lengthened to minimize the impact of parasitic devices without degrading the performance of the circuit. In addition, to the extent that parasitic devices do form along the edges of these transistors, their effect is minimal as, as with the Mirror Cascode, only the difference in parasitic leakage will affect the quiescent output voltage of the amplifier. The cascode devices (transistors M7 and M8) are both source driven so they do not exhibit the same loss of control seen in gate driven devices. Hence, the effects of radiation on this architecture are minimal compared to the other architectures presented.

To achieve a robust yet compact pixel, the Mixed–Mode PAD incorporates a combination of these radiation damage mitigation techniques. In chapter 4, where the design of the Mixed–Mode pixel is presented in detail, we will offer a discussion

of the particular steps taken to protect individual components of the pixel against radiation damage. Finally, in chapter 6, we will present measurements of the radiation tolerance of this imager.

CHAPTER 3

MIXED–MODE PAD PREHISTORY

For a new graduate student starting on a detector design project like the Mixed–Mode PAD, it is easy get the impression that this project is a unique radical departure from what has gone on before. In reality, though, this project is better described as a continuation and extension of a much larger body of work. Before delving into the details of the Mixed–Mode PAD architecture, it is useful, then, to spend a few pages looking at where this design came from. This prehistory offers a context for the design decisions that went into the Mixed–Mode PAD as well as credit to those individuals whose hard work made it possible for this project to achieve what it has.

In the early days of Pixel Array Detector development, the x-ray detector R&D community divided itself along two distinct paths: Analog Integrating PADs, which operate by collecting charge on an analog storage node—conceptually similar to a Charge-Coupled Device (CCD) imager; and photon counting Digital PADs, which use a discrimination circuit to count individual x-rays. These approaches each entail unique challenges and offer distinct advantages. It was the desire to combine the advantages of these two approaches that led in 2003 to the genesis of the Mixed–Mode PAD project, a collaboration that grew from the multi-framing Analog PAD project at Cornell and an attempt to develop a pure Digital PAD at ADSC [4].

More broadly, though, the Mixed–Mode PAD project finds its origin in the development of Scientific CCD imagers and Silicon Pixel Detectors for Particle Physics that occurred in the later part of the 20th century. By offering two possible starting points for developing x-ray imagers, these sources can be seen as the cause of the division between Digital and Analog PADs. As the following sections will

illustrate, the capabilities of the detectors that sprang from these two origins are quite distinct, to the point that both have broad and separate ranges of applications at which they excel. Because of this, the confluence of these two approaches in the Mixed–Mode PAD has yielded a new type of detector with capabilities that are distinct from these predecessors.

3.1 Digital Pixel Array Detectors

Digital Pixel Array Detectors or photon counting Pixel Array Detectors are so called because their data is quantitized into a digital format immediately upon the acquisition of each x-ray photon. Their circuit architecture typically comprises three primary stages:

- An analog pulse shaping stage that acts to bandwidth limit and amplify the detector diode signal.
- A discrimination or thresholding stage that determines when a sufficient level of signal is observed.
- A recording stage that counts the number of discriminator triggers.

This approach is, by far, the most popular, having been pursued by numerous groups [71, 22, 12, 48]. Perhaps the reason for this popularity is rooted in the fact that the first hybridized detectors followed this model. These predecessors of the x-ray photon counting Pixel Array Detectors were Silicon Pixel Detectors developed for the vertex tracking chambers of large particle physics experiments, *e.g.* ATLAS, CMS, and ALICE at the LHC [104]. Since beginning in the 1980s, research into Silicon Pixel Detectors has proven very successful in overcoming rate limitations and improving spatial resolution in increasingly energetic and complex particle collider events [87]. The research that came out of these projects can be

credited with creating the hybrid detector concept and inspiring the Digital PAD development that began in the 1990s.

Yet, while Digital PADs were modeled after Silicon Pixel Detectors, the experimental constraints and detector requirements of vertex tracking in a collider-based particle physics experiment are very different from those of synchrotron x-ray experiments. In a vertex tracking chamber, the choice of a digital over analog approach is one necessitated by the large quantities of data produced by every event and the necessity to sparse this data at the earliest possible stage. Particle physics experiments assume that, in each event, only a small number of particles, typically zero, will deposit signal into a given pixel. At the same time, the amount of deposited signal varies depending on the particle's type and path as it passes through the pixel, so the intensity of this signal needs to be recorded [87]. In many synchrotron x-ray experiments, a monochromatic x-ray beam is used so the deposited signal per-x-ray is known to within a small range of variation, as explained in section 2.3, and the number of x-rays observed during an exposure by a pixel can be in the millions.

Because of these different experimental requirements, adapting the Silicon Pixel Detector approach for Digital PADs has involved evolving the pixel back-end into a counter along with a circuit to prevent retriggering on the same photons. The resulting Digital PAD circuitry, specifically the front-end pulse shaping and back-end retrigger prevention electronics, entail a processing time for each x-ray photon that creates dead time in each pixel, suppressing the observed flux relative to the actual flux. This dead time is one of the main limiting factors for Digital PADs as it sets an effective flux limit on the device. Typically, Digital PAD papers report a flux limit of $\sim 10^7$ x-rays per pixel per second [12, 48]; however, this value is somewhat misleading, as it reflects a theoretical flux limit calculated

based on the shaping time and processing time of the pixel with assumed source statistics. As x-rays are not uniformly distributed in time, the observed number of events needs significant correction at fluxes within a few orders of magnitude of the quoted maximum. As the actual flux approaches the maximal flux, the dead time correction factor increases rapidly, typically exponentially, along with the uncertainty introduced by this dead time correction. Accurate dead time corrections within an order of magnitude of the flux limit require extremely careful characterization of the x-ray source, particularly in the case of synchrotrons where the x-rays do not obey a true Poisson distribution due to the underlying structure of electron bunches within storage ring.[1] From an experimental standpoint, the flux limit translates into a restriction on the detector's dynamic range for a given time scale. As an example, if we assume a practical flux of 10^6 x-rays per pixel per second and require a measurement with 1% statistics, then a 100 ms integration would only yield one decade of dynamic range. These considerations do not arise in event-driven particle collider experiments since the per-pixel signal levels are so much smaller.

Despite these limitations there are a number of distinct advantages that the photon counting Pixel Array Detectors offer over their analog integrating counterparts. Foremost, as these are discriminating devices they possess an energy threshold associated with their discrimination level. This gives them a natural ability to reject x-ray fluorescence as well as any Compton background present in the beam. Additionally, in these devices, the well depth grows exponentially with the amount of pixel area devoted to storage, as opposed to the linear growth one finds in most analog integrating systems, which is particularly notable because this

[1] As an example to illustrate this point consider the typical bunch timing at the Cornell High Energy Synchrotron Source. Here bunches are ~ 56 ns in duration with either 280 ns or 294 ns between the start of successive bunches. Because of this timing structure a Digital PAD should always see two photons in successive bunches but never see two photons in a single bunch.

increase in well depth does not come at the expense of precision, as the front-end gain remains fixed. Also, once digitized, the data from the pixel is effectively immune to noise yielding a very high signal to noise ratio that is only very weakly dependent on integration time.[2] Because the data is already in digital form, it may be read off the detector very rapidly. Finally, Digital PADs typically do not need to be cooled for operation at low count rates.

3.2 Analog Pixel Array Detectors

Analog PADs can be seen as the indirect descendents of the scientific CCDs that were adapted for use with x-ray sources in the 1980s and have since become the standard in imaging at synchrotron x-ray sources. Primarily advocated and developed by the Gruner Detector Group at Cornell University, a group which originally specialized in x-ray CCD development, Analog PADs adopt the fundamental methodology of CCDs. As descendants of CCDs, they collect photocurrent as a quantity of charge on an analog storage element, retaining this analog quantity until it is digitized after the exposure ends. However, the CMOS Very Large Scale Integration (VLSI) based signal processing electronics available to Analog PADs offer functionality that goes far beyond what is possible in CCDs. Rather than the complicated clocking schemes used to shift charge through and off of CCDs, Analog PADs integrate signal processing electronics at the pixel level. This integration has allowed such variety as detectors with multiple memory elements in each pixel for microsecond framing [83], massively parallel in-pixel analog-to-digital conversion for high frame rate operation [76], and nearly continuous framing though the addition of a sample-and-hold stage operating independently of the signal integration electronics [25]. Still, what all Analog PAD architectures possess in common is an

[2]With the exception of the possibility of single event upsets, as discussed in section 2.3.2.

integrating front end and, as such, they exhibit performance characteristics that differ significantly from the Digital PADs, discussed in the previous section.

Perhaps the most dramatic difference between Analog and Digital PAD methodologies is the maximum flux per pixel these devices can tolerate. Because Analog PADs process current, rather than discriminating photons, their behavior is determined almost completely by the response of the pixel front-end electronics. As the Analog PAD front end integrates currents most signal transients are smoothed out as long as their duration is short relative to the length of the exposure. Hence, generally, the maximum flux per-pixel is limited only by the bias current of the integration stage amplifier, allowing the Analog PAD methodology to exhibit very high per-pixel flux tolerance. This tolerance has been experimentally verified with groups reporting using these detectors at per-pixel fluxes as high as 10^{12} x-rays per second [86], five orders of magnitude higher than the optimal performance of Digital PADs.

While the flux capabilities of Analog PADs are impressive there are also significant drawbacks to this detector methodology. First, converting the acquired signal to a digital value and recording this takes a certain amount of time. With the exception of analog pixel architectures which incorporate very high degrees of parallel digitization (*e.g.* the Analog PADs, mentioned earlier, that incorporate Analog-to-Digital Converters (ADCs) into each pixel have been designed for a frame rates of 120 Hz [76]) there is, in general, a fundamental trade-off between the detector's speed, well depth, and precision in an optimized analog detector. This trade-off arises from the fact that detectors storing acquired signal in an analog element are limited by the span of that element, *e.g.* the depletion voltage of a CCD, so that setting a front-end gain sets the well depth. To achieve a larger well depth one can reduce the gain of the front end; however, to keep the same sensitivity, the

precision of the readout electronics must increase, which, in turn, requires longer settling times and, therefore, a slower readout. Because of the precision required in their readout and their large number of pixels, typical readout time for Analog PADs are on the order of a second—making high continuous frame rate operation difficult to achieve.

A second significant drawback, alluded to earlier, is the relation between well depth and pixel area. Generally Analog PADs use a linear capacitive charge storage element to set the system gain.[3] Because of this the well depth, and consequently the absolute dynamic range of the pixel, is strongly limited by the size of the pixel. As an example, typical sub-micron processes offer capacitors with capacitance of 1 fF/μm^2 with 3.3 V of power supply headroom. In a 150 μm^2 pixel, if one quarter of the area is devoted to the storage element ($\sim$ 5.6 pF), this capacitance and power supply could contain nearly 3.7×10^4 10 keV x-rays. Applying the same dynamic range criteria of 1% statistics per pixel, used earlier in the Digital PAD example, this gives only an absolute dynamic range of less than a decade. For repetitive experiments with a well defined reference time, these limitations are not a serious obstacle, as multiple frames may be combined to yield the necessary statistics; however, imaging systems without these characteristics pose a serious problem.

A final consideration with Analog PADs is the effect the leakage current from the detector diode layer has on the pixel well depth. Over the course of an integration this parasitic current fills a portion of the avaliable well depth of each pixel, dependant on the exposure duration and the detector temperature. Given that, for a room temperature device, the level of the leakage current can easily be equivalent to tens of thousands of x-rays per second, it is frequently necessary to cool Analog PADs to a point where the rate of dark current accumulation becomes negligible.

[3]Logarithmic and other non-linear charge storage elements have been discussed but are considered too susceptible to process and array wide variations as well as radiation damage to be practical.

This greatly complicates the design of the detector housing as it typically requires either a dry gas or vacuum environment to prevent condensation, which in turn creates problems for distributing and recording high speed, high fidelity electrical signals to and from the detector.

3.3 Contemporary PAD Projects

As we have mentioned, development of Pixel Array Detectors began with Digital PADs in the early 1990s, followed by Analog PADs later in that decade. As a result of years of dedicated effort a few of these projects (Digital PADs only, to-date) have been built into large format imagers. Of these one, the Pilatus imager developed by Dr. Christian Broennimann's group at the Swiss Light Source, transitioned in October of 2006 from a publicly supported R&D project to a product available from a private company.

A general summary of the salient traits of contemporary Analog and Digital PAD projects is offered in table 3.1. The reader should be cautioned not to draw conclusions that are too broad from this simple comparison of PAD detectors as there are caveats to each of these detectors that this presentation masks; caveats that are a direct consequence of the combination of flexibility, offered by custom designed in-pixel signal processing electronics, and constraint, imposed by the limited pixel area, that PAD developers are faced with. Despite this disclaimer, some very important points are evident in this table, echoing the discussion from sections 3.1 and 3.2. Most notable is the substantially larger maximum flux of analog over digital PADs, along with the larger well depth and generally shorter readout time of the digital over the analog PADs.[4]

[4]The LCLS PAD is somewhat of an exception to the general rule that analog PADs require long readout times because of the time they require for each digitization operation. This is because this particular PAD performs massively parallel digitization by incorporating a precision

Table 3.1: Examples of prominent, worldwide, Pixel Array Detector projects. All x-ray referenced parameters assume an x-ray energy of 10 keV.

	Medipix–2	Pilatus	XPAD3	LAD	Cornell 100×92	LCLS
PAD Method	Digital	Digital	Digital	Digital	Analog	Analog
Pixel Size [μm^2]	55×55	172×172	130×130	150×150	150×150	110×110
Read Noise [x-rays/pixel]	< 1	< 1	< 1	< 1	~ 2.6	0.2/1.3
Read Time [ms]	3	3	2	0.4	$\sim 10^3$	< 6
Well-Depth [x-rays/pixel]	2^{14}	2^{20}	2^{12}	2^{15}	$\sim 1.7 \times 10^4$	300/2,250
Max. Flux [x-rays/pixel/s]	2×10^6	$< 8 \times 10^6$	$\sim 10^6$	$\sim 10^6$	10^{12}	$> 1 \times 10^{15}$
References	[48]	[17, 71, 22]	[73, 12]	[30]	[83]	[76]

3.4 Mixed–Mode Pixel Array Detector

While the developments in Analog and Digital Pixel Array Detectors that have taken place since the early 1990s are a clear step towards meeting the detector needs of modern synchrotron science, both these approaches are marked by distinct advantages and disadvantages for synchrotron applications. These strengths and weaknesses are particularly evident when one considers how these detectors would perform in a continuous time resolved experiment. As we have discussed earlier in this chapter, Digital PADs lack the high flux performance to yield broad dynamic range with good per-pixel statistics for short integration times ($\ll 0.1$ s). Analog PADs, while capable of handling the high fluxes needed for time resolved imaging, are severely limited in terms of well depth as well as, typically, limited in terms of frame rate, with notable exceptions. Because of these considerations, neither of these detectors is well suited for continuous imaging of dynamic systems on ms time scales.

While this is only one example of the limitations of the PADs that have been developed or which are under development to-date, it is a scientifically important one. As we will discuss in more detail in section 7.5, studying the dynamics of irreversible, time-evolving systems has become increasingly important as self assembly on molecular and atomic scales plays a more important role in industrial and scientific research. While x-rays offer an attractive structural probe on these length scales, these systems present challenges for *in-situ* measurement on time scales below $\sim$1 s. This is because their growth is often irreversible and dependent on initial conditions that are either difficult or time-consuming to reproduce. While novel methods for studying particular systems have been developed [77], what

ADC into every pixel. Thus, although it requires ms to digitize the signal from a single pixel, because all pixels are digitized at once this long digitization time is does not significantly increase the readout time of the detector.

would substantially advance the field, as we argued in the introductory chapter, is an imager capable of investigating these time scales.

Table 3.2: Imager specifications for the Mixed–Mode PAD. All x-ray referenced parameters assume an x-ray energy of 10 keV. The detector frame rate given here is indicative of what is attainable from a Mixed–Mode PAD hybrid; a camera implementation will be limited by the rate at with the large quantity of data produced by the Mixed–Mode PAD can be processed and stored.

Detector Format	Multi-Mosaic of PAD Hybrids
Single PAD Hybrid Format (ASIC)	128×128 pixels
Pixel Size	150 μm × 150 μm
Framing Rate	1,000 Hz
Readout Time	< 1 ms
Read Noise	0.4 x-rays/pixel
Well Capacity	2.6×10^7 x-rays/pixel
Maximum Flux	$> 1.0 \times 10^8$ x-rays/s/pixel

Because of this and other areas where pure Digital and Analog PADs do not meet the needs of synchrotron science, an attempt to merge the most attractive parts of these two detector methodologies was undertaken in the Mixed–Mode PAD. The concept that was the genesis of this project[5] was based on the observation that, for any imaging detector, there is a fundamental accuracy limit set by the relative calibration of pixels within the device. This limit is a simple consequence of the presence of fixed pattern noise in the detector, which will be discussed in greater detail in section 6.5. Rarely is any commercial x-ray imager calibrated better than 0.5% in this regard. This is typically not a problem, however, because, in most practical applications, what the x-ray scientist needs is, effectively, a value that can be expressed in normalized scientific notation with mantissa that extends to three or four significant figures of precision. In principle, a detector that exhibited this response could offer a dynamic range that extended beyond the capabilities of "fixed-precision" detectors like Digital and Analog PADs, while

[5]Attributable to Dr. Sol Gruner, Cornell University.

offering a very high readout rate. In an area of 150 μm $\times$ 150 μm, it is difficult to see how one could implement the "floating-point-precision" circuit this concept describes. However, a limited form of this functionality can be achieved with an architecture that more directly draws from its Digital and Analog PAD predecessors, where a small well depth analog integrating front end (*i.e.* the Analog element) is coupled to an overflow counter (*i.e.* the Digital element).

While this was the origin of the Mixed–Mode PAD architecture, the merits of this architecture extend much deeper than this description implies. As will be explained in the presentation of the pixel design in chapter 4, this front end architecture breaks the interdependence of frame rate, well depth, and precision that a conventional analog imager is faced with. Consequently, the specifications that the Mixed–Mode PAD is able to attain, summarized in table 3.2, are greater than the sum of the Digital and Analog PAD paradigms that it mixes.

CHAPTER 4

MIXED–MODE PAD PIXEL DESIGN

The signal processing electronics built into each Mixed–Mode PAD pixel lie at the heart of the capabilities of this imager. On the surface the components that make up the Mixed–Mode PAD (*e.g.* integrator, comparator, switched capacitor, counter, etc.) may seem somewhat basic. However, the degree of precision required from these components, compounded with the fact that they must fit into an area of 150 μm $\times$ 150 μm and they must work reliably over an array spanning 19.2 mm $\times$ 19.2 mm, all the while performing in the demanding and harsh environment of the synchrotron, makes this task very challenging. Consequently, the pixel included in the final ASIC is the quite non-trivial result of a substantial development effort involving a group of people spread between the Cornell PAD development group and our collaborators at Area Detector Systems Corporation (ADSC).

Before delving too deeply into a discussion of the design of the Mixed–Mode PAD pixel, it is appropriate to offer a few remarks to acknowledge the contributions of various people to the development of the final Mixed–Mode PAD ASIC. The current Mixed–Mode PAD pixel and on-chip support electronics are the result of a design process that spanned many MOSIS submissions; from the first 16 $\times$ 16 pixel prototype, in the fall of 2003, to the final 128 $\times$ 128 pixel hybrid imager, completed at the end of 2006. Following a standard development model of design $\rightarrow$ implementation $\rightarrow$ testing $\rightarrow$ redesign, the signal processing ASIC evolved considerably through these submissions.

In the design phase, we can distinguish between architectural level design and circuit level design. The former occurred during regular design meetings and reviews that took place over the extent of the project, involving collaborators from

ADSC: Tom Hontz, Wayne Vernon, and Skip Augustine;[1] and from Cornell: Sol Gruner, Mark Tate, and myself.[2] Design, simulation, and layout of the circuits for this ASIC fell to Skip Augustine and myself. In this task, my focus was primarily on developing a high-performance, low-noise pixel front end, while Skip was responsible for most of the peripheral circuitry and combining the individual circuit layouts into the full submission layout. The boundaries between these tasks were, however, maleable, adjusting depending on the needs of a particular stage of the project.

The critical task of testing and verifying the results of these submissions was divided between Wayne Vernon, assisted for much of the project by Doan Nyguen, at ADSC and myself at Cornell. In the testing, a distinction developed after the basic functionality of design was verified, as the ADSC group moved towards developing a wafer level ability to verify ASIC functionality, a step necessary for ultimate commercialization, and design of custom off-chip support electronics for the final detector, while at Cornell the focus was placed on detailed evaluation of the device performance. None of this work could have been accomplished without the efforts of Matt Allin, at ADSC, who managed the initial packaging of these devices and Tom Hontz, who coordinated the flip-chip bump bonding.

In the remainder of this chapter, we will go through the final design of the Mixed–Mode PAD pixel, breaking it down into its individual components and discussing the subtleties associated with each. This final design changed considerably from the initial implementation as a result of the iterative, evolutionary design process. Consequently, although the pixel was my focus in the ASIC design, this

[1]Of Augustine Engineering, San Diego, CA; a VLSI circuit designer contracted by ADSC to work on the Mixed–Mode PAD project.

[2]Matt Renzi and Alper Ercan, former member of the Cornell Detector Development Group, also deserve recognition for their foundational work on the Mixed–Mode PAD project. While this work did not directly propagate into the Mixed–Mode PAD design, it set the framework for many of the early discussion of the Mixed–Mode PAD architecture.

design is, in fairness, a combination of effort, ideas, and learning drawn from across our collaboration. In the conclusion of this chapter, we will review a number of ideas, developed over the course of this project, about how this design could be extended to create an even more capable imager in the future.

4.1 Mixed–Mode PAD Pixel Architecture

Because of the Gruner group's history with Analog PADs, the approach taken to this merger of methodologies began from the integrator based front end of the Analog PAD. From a long history with x-ray CCD detectors, within both the Gruner groups at Cornell and our collaborators at ADSC, the fundamental interdependence of frame rate, well depth, and precision present in an optimized analog detector, discussed in the previous section, was well understood. The underlying concept of the Mixed-Mode PAD is to break this interdependence by performing a coarse analog-to-digital conversion in-pixel during the exposure. Rather than digitizing and counting single photons, as done by Digital PADs, the Mixed-Mode PAD digitizes and counts blocks of photons, typically the equivalent of 100 10 keV x-rays, leaving a small residual charge for post exposure processing. This allows for a precision, high-gain front end without sacrificing well depth. It also permits high speed readout as only a relatively small, in terms of total number of x-rays, residual analog signal remains in each pixel to be digitized in the period between exposures.

Figure 4.1 depicts the basic pixel architecture. Over time, the details of this architecture have changed with the evolution of the pixel, but the basic structure and essential component operations have remained the same. The basic architecture and operation begin with an integrator that accumulates photocurrent from the detector diode (I_{sig}). Charge accumulating on the integrator ($Q_{\mathrm{int}} = C_{\mathrm{int}}(V_{\mathrm{pix}} - V_{\mathrm{outp}})$, as defined in figure 4.1) causes the integrator output (V_{outp}) to slew towards ground.

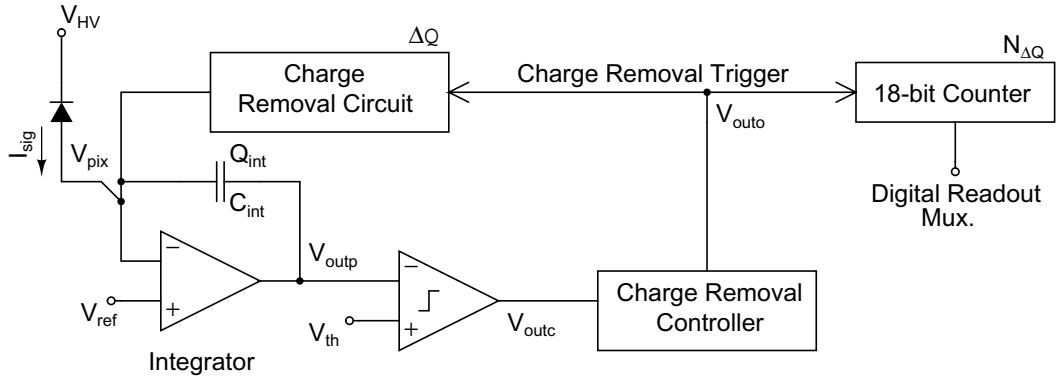

Figure 4.1: High level description of the Mixed-Mode PAD pixel architecture.

When the output voltage falls below the comparator's threshold ($V_{\rm th}$) the charge removal controller is activated. This circuit sends out fixed length pulses which cause charge (ΔQ) to be removed from the integrator and increment an in-pixel counter ($N_{\Delta Q}$). At the end of the exposure, residual voltage at the integrator output ($V_{\rm res}$) and the number of charge removals performed are recorded so that the total charge accumulated ($Q_{\rm tot}$) may be reconstructed by:

$$Q_{\rm tot} = C_{\rm int}\left(\frac{dV_{\rm eqv}}{dN_{\Delta Q}}N_{\Delta Q} + V_{\rm outp}\right), \tag{4.1}$$

where $\frac{dV_{\rm eqv}}{dN_{\Delta Q}}$ is a scaling constant that converts the number of charge removal operations into an equivalent shift of the integrator output voltage.

The initial Cornell and ADSC Mixed-Mode design utilized a reset-to-zero architecture wherein the integrator is reset each time the comparator triggers. This architecture benefits from its electrical simplicity; however, it substantially complicates the data analysis by introducing a signal-dependent dead time. In order for this design to work well, it is necessary to keep the reset switch open long enough for the front end to settle after reset. For this, a triggered fixed-length pulse generator is needed. If the signal is constant during the reset period and the reset period is of constant duration, then the dead time correction is straightforward. Practically, however, these assumptions are not valid: synchrotron sources have

a signal structure that varies depending on the bunch structure and it is difficult to build a fixed-length pulse generator whose pulse length is repetitive over long periods.

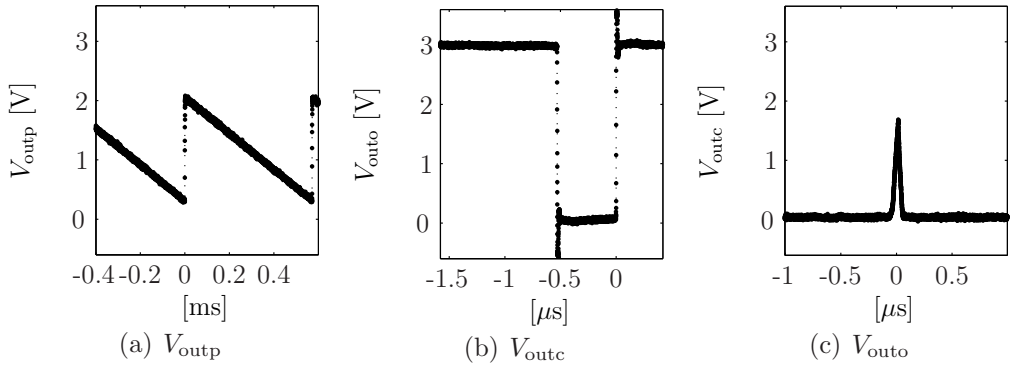

(a) V_{outp} (b) V_{outc} (c) V_{outo}

Figure 4.2: Voltage traces acquired from active nodes within the pixel (AE176 submission), labeled as in figure 4.1, illustrating operation with a constant test current source.

To avoid these problems, the architecture was changed to a design based on the concept of a ΣΔ-Analog-to-Digital Converter (ΣΔ-ADC). A ΣΔ-ADC operates by accumulating signal, typically current, in an integration stage while comparing the accumulated signal to a set threshold. Whenever the threshold is passed, a fixed quantity of the accumulated signal is removed. The ΣΔ operation is illustrated using the Mixed-Mode PAD in figure 4.2, which shows the response of selected nodes within the pixel to a constant test current source. Panel (a) depicts the output from the integrator (V_{outp}). Integration begins when the pixel reset switch is opened and V_{outp} begins to slew towards ground. The lowest voltage reached by V_{outp} corresponds to the threshold voltage (V_{th}) of the comparator. As V_{outp} crosses V_{th} the comparator output (V_{outc}, shown in panel (b)) briefly rises, activating the charge removal process, then a rapidly returns to ground as the removal of charge from the front-end integrator draws V_{outp} once again above V_{th}, as shown in figure 4.2, panel (a). Because a fixed quantity of charge (ΔQ) is removed, V_{outp} does not

necessarily return to V_{ref}; instead, immediately after the charge removal

$$V_{\mathrm{outp}} \approx V_{\mathrm{th}} + \frac{\Delta Q}{C_{\mathrm{int}}}, \tag{4.2}$$

where this result is only approximate because the actual voltage will vary slightly depending on the amount of signal acquired during the charge removal. The ability to acquire signal during the charge removal operation is a very important difference between the ΣΔ architecture and the reset-to-zero architecture. A reset-to-zero pixel throws away signal during the *in-situ* resets, so that the designer must trade off between this reset dead time and time allowed to settle the front-end following the abrupt reset impulse. In a ΣΔ pixel, this interdependence is broken because signal acquired during the charge removal operation in retained. The control signal (V_{outo}) for a charge removal operation is shown in panel (c). Comparing the length of this pulse with the scope traces for V_{outp} or V_{outc}, one can see that the full duration of the charge removal operation is many times longer than the perceivable disturbance from the charge removal on these scope traces. By making the amount of signal removed first-order independent[3] of the duration of the removal operation and signal intensity, this ΣΔ approach avoids the problems discussed with the reset-to-zero architecture.

The Mixed-Mode PAD is not the first case of a ΣΔ style ADC integrated into the pixels of an area imager. The first known example of an imaging detector of this type was developed at Stanford University [107] for optical applications. This imager, however, utilizes a more canonical ΣΔ-ADC architecture than does the Mixed–Mode PAD. In the Mixed–Mode PAD, only the total number of charge removal operations are recorded, as the charge removal is an asynchronous, triggered event. In the Stanford detector, as well as most ΣΔ-ADCs charge removal occurs

[3]There is still a dependence on the front end settling time and input signal intensity; however, this is greatly suppressed compared to the reset-to-zero architecture. A detailed discussion of this dependence can be founds in the detailed discussion of the front end circuit given in section 4.2.2.

synchronous with a sampling clock edge; *i.e.,* a clock edge is used to activate the comparison operation and the results of the comparison are not stored in the pixel but rather streamed out of the detector. While this sort of temporal information would be very desirable in some applications, it creates practical problems for high flux-per-pixel measurements. The maximum flux-per-pixel ($\Phi_{\max}$) tolerable by this design is

$$\Phi_{\max} = \frac{1}{c_{xq}} \cdot \Delta Q \cdot f_{\text{sync}}, \tag{4.3}$$

where c_{xq} is the x-ray charge yield, ΔQ is the quantity of charge removed in each charge removal operation, and f_{sync} is the synchronous sampling frequency. The specifications for our detector require that $\Phi_{\max} \geq 10^8$ 10 keV x-rays per second. If the analog well depth of the detector is 100 10 keV x-rays, then this implies f_{sync} must be at least 1 MHz. While technically feasible, distributing the required clock across the large area of the Mixed–Mode PAD pixel array was deemed very risky due to the potential of noise coupling into the pixel's analog front end.

Instead, an asynchronously triggered charge removal circuit was decided upon, wherein the comparator output acts as a trigger to the charge removal controller. This circuit, which will be discussed in detail in section 4.2.2.3, is responsible for controlling the frequency and duration of the charge removal operation by creating a waveform that controls the charge removal operation and increments the in-pixel counter.

4.2 Primary Pixel Components

The remainder of this chapter is divided into two parts. The first part reviews the primary components of the pixel, those shown in figure 4.1. The second and final part reviews the functions of the peripheral components of the pixel.

4.2.1 Pixel Integrator

The integration stage is the foremost signal processing element on the Mixed-Mode PAD pixel. As such, its performance characteristics strongly affect the pixel's behavior, and, therefore, it is important to thoroughly analyze this circuit to understand how it impacts the performance of the device as a whole. We begin this analysis with a simple derivation of the relations underlying the operation of the integrator.

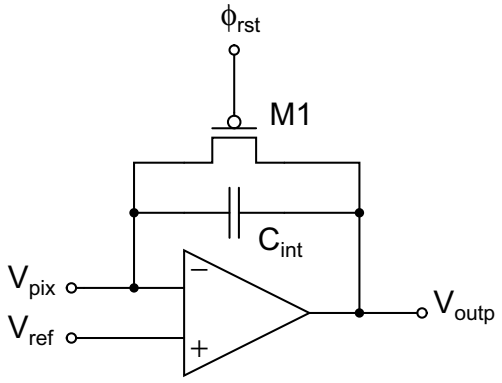

Figure 4.3: Schematic of the pixel integrator.

An integrator is a relatively simple circuit whose basic schematic is shown in figure 4.3. It comprises an operational amplifier and a capacitor (C_{int}), called the feedback capacitor, connected between the amplifier output (V_{outp}) and the amplifier's inverting input (V_{pix}), which is also the integration node of the pixel. The non-inverting input of the amplifier (V_{pos}) is connected to a fixed reference voltage (V_{ref}). This configuration creates a virtual ground at the integration node, whereby the amplifier slews its output as necessary to keep the voltage at this node fixed at V_{ref}. Finally, a reset switch, the pMOS transistor M1, allows for cancellation of the charge stored on the integration capacitor (Q_{int}) prior to each integration.

In its quiescent configuration, $V_{\text{pix}} = V_{\text{ref}}$ and $Q_{\text{int}} = C_{\text{int}}(V_{\text{pix}} - V_{\text{outp}})$. Applying

a small, periodic test signal ($\delta V_{\rm pix}$) of angular frequency ω to the input node, *i.e.* $V_{\rm pix} \rightarrow V_{\rm pix}+\delta V_{\rm pix}$, one finds that the effective capacitance ($C_{\rm eff}$) of the amplifier boosted integration capacitor is

$$\begin{aligned} \delta V_{\rm outp} &= -A(\omega) \cdot \delta V_{\rm pix} \\ \Rightarrow \delta Q_{\rm int} &= C_{\rm int}(1 + A(\omega)) \cdot \delta V_{\rm pix}, \\ \Rightarrow C_{\rm eff} &= \frac{\delta Q_{\rm int}}{\delta V_{\rm pix}} \\ &= C_{\rm int}(1 + A(\omega)), \end{aligned} \tag{4.4}$$

where $A(\omega)$ is the frequency dependent gain of the amplifier and we have assumed, for simplicity, that $Q_{\rm int} = 0$ prior to the application of $\delta V_{\rm pix}$. This result shows that the amplifier acts to boost the effective size of the integration capacitor. The importance of this amplification arises when we consider parasitic capacitance between the integration node and ground ($C_{\rm pix}$). This causes charge sharing between this parasitic capacitance and integration capacitor proportional to the ratio of the of the parasitic capacitance to the effective size of the integration capacitor. Measurements of the Mixed–Mode PAD pixel parasitic capacitance indicate that it is ~ 200 fF,[4] four times the nominal size of the integration capacitor. As a result only 20% of the photocurrent will be drawn onto the integration capacitor through charge distribution. The remaining 80% depend on the amplifier's slewing of $V_{\rm outp}$ to be drawn onto the integration capacitor. In this way, the performance of the integrator strongly depends on the performance of the amplifier. This also illustrates the importance of using a high gain amplifier, *e.g.* a large $A(\omega)$ in equation 4.4.

[4]With this architecture, this measurement may be made easily noting that when the reset switch open $C_{\rm pix}$ creates capacitive feedback network with $C_{\rm int}$. The gain of the amplifier with this network should be $A = (C_{\rm pix}+C_{\rm int})/C_{\rm pix}$. Thus, by modulating $V_{\rm ref}$ and measuring resulting the amplitude at $V_{\rm outp}$, one can estimate $C_{\rm pix}$ for a given pixel.

4.2.1.1 Integrator Amplifier – Performance Specifications

The front-end amplifier is arguably the most significant active circuit in the Mixed–Mode PAD ASIC. Because of its location at the beginning of the analog signal processing chain, its performance strongly affects the quality of the data. The previous section showed the importance of the amplifier gain in collecting the full photocurrent signal. However, beyond the amplifier gain, other parameters, such as transconductance, slew rate, and power consumption, also have significant effects on the pixel performance. In this section, we begin by laying out the performance specifications required of the integrator amplifier. The remainder of the section then describes the structure of the amplifier, followed by a combination of analytical and simulated[5] performance characteristics. Finally, as the amplifier provides a path for noise to enter into the beginning of the signal processing chain, both from sources within the amplifier and from external sources, such as the power supply or the control and bias lines, an analysis of these amplifier aspects is also presented.

Table 4.1 provides a summary of the design specifications for the Mixed-Mode PAD integrator amplifier, along with the expected performance of the amplifier as derived from SPICE simulations and analytical calculations. The design specifications listed are based on the required performance of the detector and were chosen conservatively to ensure design robustness over a wide range of manufacturing variations and operating conditions.

4.2.1.2 Integrator Amplifier – Architecture and Analytical Analysis

The Mixed–Mode PAD integrator amplifier uses a folded cascode topology; the schematic of which is shown in figure 4.4 with transistor sizing given in table 4.2.

[5]Unless noted otherwise circuit simulations were performed with T-Spice, by Tanner Research Inc., using transistor models for the TSMC 0.25 μm process provided by our ASIC manufacturer.

Table 4.1: Summary listing of design specifications and expected performance for the pixel front-end amplifier.

Design Parameter	Min. Value	Max. Value	Expected Value
Differential Gain	60 [dB]	—	91 [dB]
Gain-bandwidth Product	20 [MHz]	—	37 [MHz]
Phase Margin	45°	—	57°
Transconductance	25 [μA/V]	—	47 [μA/V]
Slew Rate (@ 250 fF load)	10 [V/μs]	—	20 [V/μs]
Total Output Referred Noise	—	850 [μV]	$<$ 750 [μV]
PSRR (unity gain feedback)	40 [dB]	—	90.9 [dB]
PSRR (capacitive feedback)	40 [dB]	—	69.1 [dB]
CMRR	40 [dB]	—	60.0 [dB]
Power Consumption	—	50 [μW]	33 [μW]

This topology has merits of simplicity, wide bandwidth, and high gain potential. It consists of a differential input stage, transistors M1 and M2 in the schematic, which channel current generated by the nMOS bias transistors between the input and output stages. Current channeled into the output by the M2 branch is mirrored by the pMOS Wilson Current Mirror to subtract from the current channeled into the output by the M1 branch. Because there is no current amplification, current is only redirected or mirrored, the amplifier's transconductance (G_m) is essentially determined by the sizing and bias of the input stage. Assuming this differential pair is well matched, so that the transconductance of the input pair is roughly equal ($g_{m,\mathrm{M1}} = g_m = g_{m,\mathrm{M2}}$), then the transconductance of the amplifier may be approximated as

$$G_m \approx g_m. \tag{4.5}$$

The open loop output impedance of the amplifier (Z_{out}) is given by the parallel combination of the output impedance ($1/g_{ds}$) of the nMOS bias transistor in the output branch (M10), boosted by the effect of cascoding with transistor M8, and

the output impedance of the pMOS Wilson Current Mirror,

$$Z_{\text{out}} = \frac{g_{m,\text{M8}}}{g_{ds,\text{M8}} \cdot g_{ds,\text{M10}}} \Big|\Big| \frac{g_{m,\text{M3}}}{g_{ds,\text{M6}} \cdot (g_{ds,\text{M3}} + g_{ds,\text{M7}})}, \tag{4.6}$$

where the double lines ($||$) indicate the parrallel operator; *i.e.,* $A||B = (A^{-1} + B^{-1})^{-1}$. This yields a total DC gain (A_{DC}) of

$$A_{\text{DC}} = G_m Z_{\text{out}} \tag{4.7}$$

$$= g_m \cdot \left(\frac{g_{m,\text{M8}}}{g_{ds,\text{M8}} \cdot g_{ds,\text{M10}}} \Big|\Big| \frac{g_{m,\text{M3}}}{g_{ds,\text{M6}} \cdot (g_{ds,\text{M3}} + g_{ds,\text{M7}})} \right). \tag{4.8}$$

The open loop frequency response is expected to have a single pole transfer function given by

$$H(\omega) = \frac{A_{\text{DC}}}{1 + \omega Z_{\text{out}} C_{\text{out}}}, \tag{4.9}$$

where C_{out} is the output, load, capacitance. Solving for $|H(\omega)| = 1$ shows us that the unity gain frequency (*i.e.* the gain bandwidth product (GBW)) is determined by

$$\omega_1 = \text{GBW} \approx \frac{G_m}{C_{\text{out}}}. \tag{4.10}$$

The biasing network matches the current generated by the pMOS bias transistor (M0) and each of the nMOS bias transistors (M9 and M10). Letting I_{ioa} denote this common bias current, then the maximum output current is expected to be $\pm I_{\text{ioa}}$. Based on this, the output slew rate is expected to be given by

$$\text{Slew Rate} = \left| \frac{I_{\text{ioa}}}{C_{\text{out}}} \right|. \tag{4.11}$$

Previous Analog PAD designs by the Gruner Group fabricated in the 0.25 μm TSMC process ([83, 26]) have all utilized a mirrored cascode architecture[6]. However, the Mixed–Mode PAD's need for a lower power and wide bandwidth front end (> 10 MHz, for reasons we will discuss in section 4.2.2.2) prompted a move away

[6] A discussion of the mirrored cascode architecture may be found in [67].

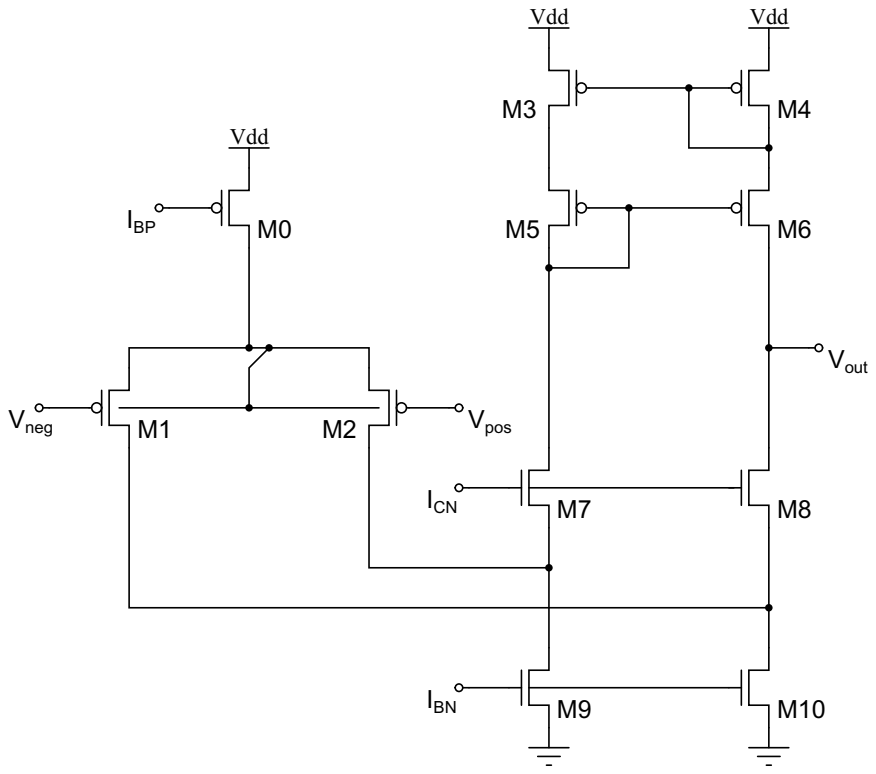

Figure 4.4: Schematic of Mixed–Mode PAD integrator amplifier. Transistor sizing and multiplicity are given in table 4.2. The bulk of transistors M1 and M2 are connected to their common source. All other bulks are connected to the analog supply (V_{DDA}) or analog ground (V_{GNDA}) as is appropriate by type. No stabilization capacitor is needed.

Table 4.2: Transistor sizing for the Mixed–Mode PAD integrator amplifier described in figure 4.4. The length unit of λ is a common VLSI scaling parameter intended to allow design to be easily migrated between different processes. For the TSMC 0.25 μm process $\lambda = 0.12$ μm.

Transistor	W [λ]	L [λ]	W Multiplicity
M0	40	40	4
M1	40	5	8
M2	40	5	8
M3	40	10	4
M4	40	10	4
M5	80	3	8
M6	80	3	8
M7	80	5	8
M8	80	5	8
M9	20	80	4
M10	20	80	4

from this canonical design. While the first order performance characteristics, derived above, are identical for both mirrored cascode and folded cascode amplifiers, there are subtle differences between these architectures that allow the folded cascode to operate at the same bandwidth, but with improved phase margin,[7] noise performance, and device matching. To understand this performance difference, one has to look at the effect of device sizing on conductance and capacitance at internal nodes of the amplifier while considering the noise generation and matching of the devices connected to these nodes.

Both the mirrored cascode and the folded cascode have internal nodes that affect the high frequency performance of the amplifier. In the mirrored cascode, the significant internal nodes occur at the gates of the three current mirrors. The frequencies at which these internal nodes begin to affect the phase response of the amplifier may be approximated by calculating the ratio of the conductance of the

[7]For our purposes, the phase margin is the difference, measured in degrees, between the phase of the output signal of the amplifier and -180 deg., at the unity gain frequency of the amplifier. It is important in determining the stability of the amplifier, as insufficient phase margin can lead to instability [38].

node to its capacitance. Typically, the conductance of the node will be dominated by the g_m of the diode-connected mirror transistor. This transistor is also one of the dominant noise sources in the mirrored cascode architecture. Increasing its transconductance increases its thermal and flicker noise. It is possible to reduce the flicker noise, to an extent, by enlarging the gate area. However, enlarging the gate area increases the node capacitance, lowering the frequencies at which these nodes affect the amplifier frequency response. This interdependence of device transconductance, device noise, and node capacitance links the bandwidth, phase margin, and noise performance of the amplifier in a way that limits the performance attainable with the mirror cascode architecture.

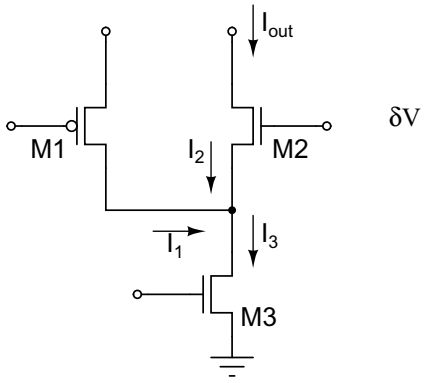

Figure 4.5: Model circuit used to analyze the effective transconductance of the nMOS folded cascode.

The situation is notably different in the folded cascode architecture. In the Mixed–Mode PAD amplifier design, there are four internal nodes that contribute significantly to the high frequency pole of the amplifier. Of these, only one has the same bandwidth, noise, and phase margin trade-offs that restrict the mirror cascode architecture.

Two of the other significant internal nodes occur at the current branching nodes of figure 4.4, equivalent to the branching node shown in figure 4.5, where current

generated by the nMOS bias transistors is split between the input and output stages. The conductance of these nodes is dominated by the source conductance ($g_s = g_m + g_{mb}$) of the nMOS folded cascode transistors. These cascode devices show only a weak dependence between their transconductance and size and the noise they contribute to the amplifier. To see this, consider the effective transconductance of the gate of these nMOS cascodes. Using the notation shown in figure 4.5, a small test voltage (δV) applied to the gate of transistor M2 will create a current (δI) flowing in the channel. At the source of this transistor this current will see a current splitter comprising the drain to source conductances of transistors M1 ($g_{\mathrm{ds},1}$) and M3 ($g_{\mathrm{ds},3}$) and source of transistor M2 ($g_{\mathrm{s},2}$). Hence,

$$\begin{aligned}
\Delta I_1 &= \frac{g_{\mathrm{ds},1}}{g_{\mathrm{ds},1} + g_{\mathrm{ds},3} + g_{\mathrm{s},2}} \delta I \\
&\approx 0, && (4.12) \\
\Delta I_2 &= -\frac{g_{\mathrm{s},2}}{g_{\mathrm{ds},1} + g_{\mathrm{ds},3} + g_{\mathrm{s},2}} \delta I \\
&\approx -\delta I, && (4.13) \\
\Delta I_3 &= \frac{g_{\mathrm{ds},3}}{g_{\mathrm{ds},1} + g_{\mathrm{ds},3} + g_{\mathrm{s},2}} \delta I && (4.14) \\
&\approx 0, && (4.15)
\end{aligned}$$

for $g_{\mathrm{ds},1} \ll g_{\mathrm{s},2}$, $g_{\mathrm{ds},3} \ll g_{\mathrm{s},2}$. The total change in the output current is given by $\Delta I_{\mathrm{out}} = \delta I + \Delta I_2 \approx 0$, so the effective transconductance of the nMOS cascode is approximately zero. This, in turn, implies that the device's contribution to the noise of the amplifier should be negligible. It is therefore possible to use very wide devices to get high source conductance without the concern that this will degrade the amplifier's noise performance.

Widening the gate has an additional effect that improves the frequency response of the amplifier. At the current branching nodes, the dominant capacitance is due to the bulk-to-junction capacitance of the nMOS cascode transistor. While

increasing the gate width will increase the overlap capacitance it also pushes the device towards weak inversion. In weak inversion, the source capacitance dramatically drops as changes in the channel charge shift from changes in the inversion charge to changes in the depletion charge. By operating the device in weak inversion, it is thus possible to use large devices, with reduced flicker noise and improved matching, without compromising the frequency response of the amplifier.

The final two nodes occur in the pMOS Wilson Current Mirror on the output branch. This circuit incorporates an internal feedback loop that cancels noise generated by the cascoding pMOS devices. Since the Wilson Mirror is a relatively common circuit element treated in many IC design books, *e.g.* [38], we will not provide a detailed analysis of it here and instead state the following results:

$$\begin{aligned}
\delta I_{\mathrm{out,M3}} &= -\left(\frac{a}{1+a}\right)\delta I_{\mathrm{N,M3}} \\
&\approx -\delta I_{\mathrm{N,M3}}, && (4.16)\\
\delta I_{\mathrm{out,M4}} &= \left(\frac{a}{1+a}\right)\delta I_{\mathrm{N,M4}} \\
&\approx \delta I_{\mathrm{N,M4}}, && (4.17)\\
\delta I_{\mathrm{out,M5}} &= \frac{g_{\mathrm{ds,M3}}}{g_{\mathrm{ds,M3}}+g_{\mathrm{s,M5}}}\left(\frac{1}{1+a}\right)\delta I_{\mathrm{N,M5}} \\
&\approx \frac{g_{\mathrm{ds,M3}}(g_{\mathrm{ds,M3}}+g_{\mathrm{ds,M7}})(g_{\mathrm{m,M4}}+g_{\mathrm{s,M6}})}{g_{\mathrm{m,M4}}g_{\mathrm{m,M6}}(g_{\mathrm{ds,M3}}+g_{\mathrm{s,M5}})}\delta I_{\mathrm{N,M5}}, && (4.18)\\
\delta I_{\mathrm{out,M6}} &= \frac{g_{\mathrm{m,M4}}}{g_{\mathrm{m,M4}}+g_{\mathrm{s,M6}}}\left(\frac{1}{1+a}\right)\delta I_{\mathrm{N,M6}} \\
&\approx \frac{g_{\mathrm{ds,M3}}+g_{\mathrm{ds,M7}}}{g_{\mathrm{m,M6}}}\delta I_{\mathrm{N,M6}}, && (4.19)
\end{aligned}$$

for $1 \ll a$, where a is the open loop gain of the Wilson current mirror,

$$a = \frac{g_{\mathrm{m,M4}}g_{\mathrm{m,M6}}}{(g_{\mathrm{ds,M3}}+g_{\mathrm{ds,M7}})(g_{\mathrm{m,M4}}+g_{\mathrm{s,M6}})}, \qquad (4.20)$$

$\delta I_{\mathrm{N,M}n}$ is the noise current generated by transistor Mn (n = 3, 4, 5, or 6) and $\delta I_{\mathrm{out,M}n}$ is the resulting change in the drain current of the respective transistor. As

long as the transistors remain saturated, causing the drain-to-source transconductance (g_{ds}) to be small relative to the source or gate transconductance (g_{s} and g_{m} resp.), equations 4.18 and 4.19 will both be ~ 0. Thus, it is possible to operate the cascode transistors well into weak inversion without a noise penalty, thereby reducing their gate capacitance and boosting their transconductance, raising the frequency at which this node begins to degrade the performance of the amplifier.

The final significant node is connected to the gates of the load transistors in the Wilson Current Mirror. Since this node occurs on only the non-inverting signal path in the amplifier, its degenerative effect on the frequency performance is mitigated by the effect of the inverting signal path. To illustrate this mitigation, recall that the signal at the output is the sum of the signals on the inverting and non-inverting branches. Thus, if one branch has no phase shift, but equal amplitude, the maximum phase shift a pole on the other branch can induce in the output to 45°. This effect can be seen directly in the frequency response of the Mixed–Mode PAD integrator amplifier, shown in figure 4.8. Near 1 MHz, a loss in phase margin begins to occur that plateaus between 20 MHz and 80 MHz. The initial phase margin loss is due to the node at the gate of the load transistors in the Wilson Current Mirror and the plateau occurs because this node affects the frequency performance of the circuit at much lower frequencies than nodes on the inverting signal path.

4.2.1.3 Integrator Amplifier – Performance Characteristics

The preceeding discussion of the amplifier's architecture offers a qualitative view of the integrator amplifier's design. This section aims to extend this view by offering a quantitative description of the performance limits of the amplifier.

To begin, it is necessary to specify the operating range for the integrator amplifier's external controls, *i.e.* bias currents and reference voltage. The amplifier is

designed to operate within specifications over a range of bias currents (I_{ioa}) from 2.5 μA to 7.5 μA. The upper limit of 7.5 μA is not due to the amplifier or performance requirements of the pixel but is rather a power constraint imposed by the power budget of the full ASIC. At $I_{\text{ioa}} = 7.5$ μA, 15 μA of current is consumed by each pixel. As the array contains 128×128 pixels, this per-pixel consumption results in a total current draw of ~ 0.25 A, the upper limit budgeted for this portion of the design. Comparatively, at the nominal operating bias of $I_{\text{ioa}} = 5$ μA[8] the integrator amplifiers draw nearly 165 mA, which is still a considerable amount of power at more than half of the average total Analog power draw of the ASIC.

The lower limit of the amplifier's operating range is set by the transconductance requirements of the front end as determined by the maximum flux specification. The Mixed–Mode PAD's maximum tolerable flux of 1×10^8 10 keV x-rays/pixel/s generates nearly 44 nA of photocurrent. For this current to be drawn onto the integration capacitor, there needs to be an equivalent current draw from the capacitor by the integrator amplifier, requiring a constant difference between V_{ref} and V_{pix}. To prevent this difference from causing a significant systematic error, we constrain this difference to be less than 0.2% of the total analog residual full range, *i.e.* 2 mV. Thus, the amplifier transconductance should be no less than 22 μA/V. We will show shortly that this is met when $I_{\text{ioa}} > 2.5$ μA.

As discussed earlier in this section, the transconductance of the amplifier will be roughly the transconductance of the input stage. Given the sizing of the input transistors, they will operate in moderate to weak inversion under most bias conditions so that a strong inversion approximation of the transconductance will be inaccurate. For this reason, we use operating point simulations to estimate the amplifier's transconductance. For a nominal operating bias of $I_{\text{ioa}} = 5$ μA TSpice small signal analysis reports that the g_m of the input pair is 47 μA/V at

[8] All simulation results reported in this section use this biasing condition.

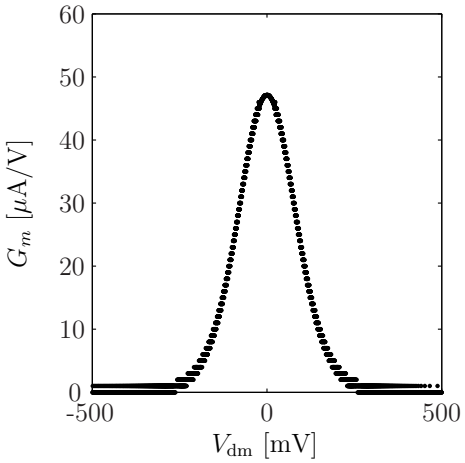

Figure 4.6: Mixed–Mode PAD integrator amplifier transconductance as a function of differential mode input voltage.

the quiescent operating point of the amplifier. To check this over a broader region, a simulation was conducted where the output was held at a fixed voltage while the differential mode input was swept. The transconductance was then computed by differentiating the resulting $I_{\rm out}$ vs $V_{\rm dm}$ curve. The results are shown in figure 4.6, where the maximum transconductance of 47 μA/V is in agreement with the operating point analysis expectations.

To relate this result to the operational range of the amplifier's bias current, we first note that the largest possible change in transconductance with bias current occurs when the the transistors are in weak inversion. In this limit, the transconductance is directly proportional to the bias current so that we can use our calculation of the nominal and minimal transconductance to set a conservative limit on the minimal bias current. To do this, we observe that the ratio of the nominal transconductance (47 μA/V) to the minimum transconductance (22 μA/V) is ~ 2. In the weak inversion limit, this ratio is equivalent to the ratio of the nominal bias current to the minimal bias current, telling us that the minimal bias current will be half of the nominal bias current, *i.e.* 2.5 μA.

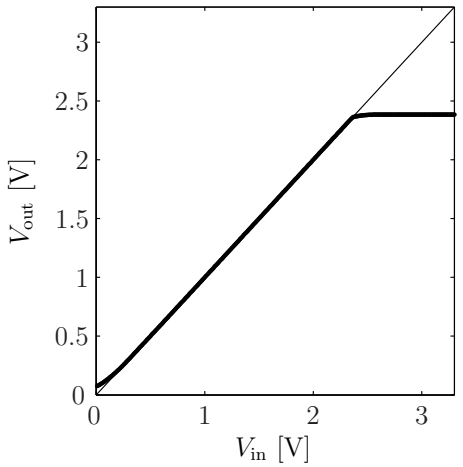

Figure 4.7: DC sweep simulation of the Mixed–Mode PAD integrator amplifier. A line of unity slope through the origin is included for reference.

Beyond the biasing conditions, the input range of the integrator amplifier is a consideration that impacts the setting of the integrator reference (V_{ref}). To determine this range, we first need the common mode input range of the amplifier. A DC sweep of the non–inverting input with the amplifier configured as a unity gain follower is shown figure 4.7. Defining the common mode input range as the range of voltages where the difference between the non–inverting input and the output is less than ± 1 mV then it extends from 0.55 V to 2.36 V. The integrator reference should be near the upper end of this range to allow for the specified 1 V output slew; however, an overshoot margin is desirable so it should not be set at the upper limit. Because of this, the operating range for the integrator reference is 2.1 V to 1.6 V.

Under the nominal operating conditions of $I_{\text{ioa}} = 5$ μA and $V_{\text{ref}} = 2.0$ V, the open loop frequency response of the amplifier was simulated—a Bode plot of the results is shown in figure 4.8. This data is taken from an amplifier within the pixel to as closely as possible mimic actual loading and feedback conditions. Accurate estimation of this design's frequency response is very important, because

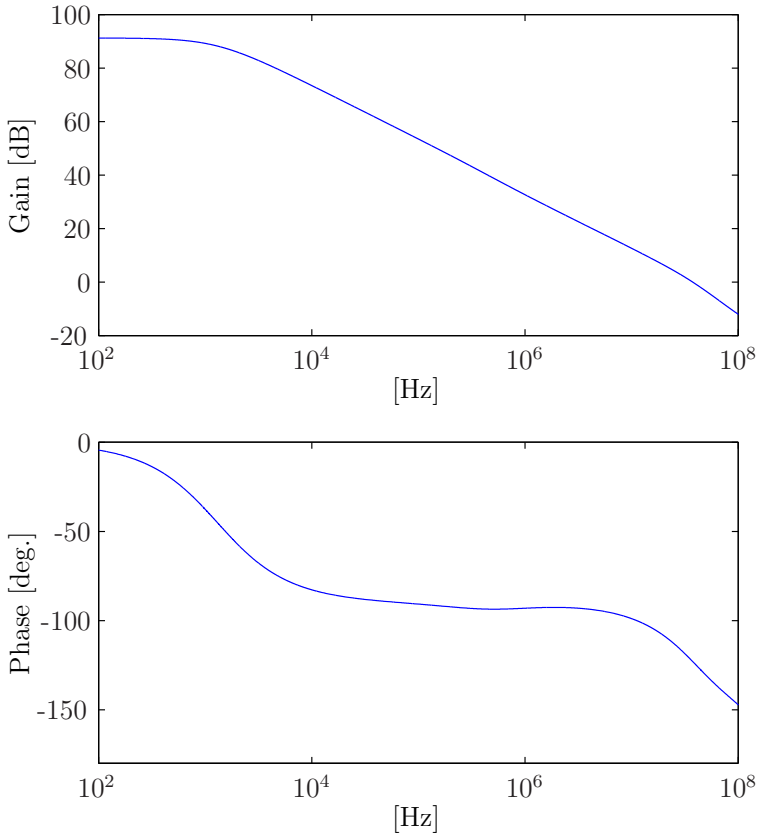

Figure 4.8: Bode plot depicting the frequency response of the Mixed–Mode PAD integrator amplifier under its nominal, 5 μA bias, operating conditions. This figure shows that the unity gain bandwidth of the amplifier is $\sim$30 MHz with a phase margin of nearly 45 deg..

the design strives for both wide bandwidth response and low power operation. As a result, the design operates in a regime where the first order transfer function of equation 4.9 is not truly applicable, because the amplifier's bandwidth extends to the point that parasitic poles contributed by nodes within the amplifier become significant. To further illustrate this point, during small area (16×16 pixels) prototyping fabrications, carried out before the full array submission, oscillations on V_{outp} were observed in some pixels when the reset switch (ϕ_{rst}) was closed, although test structures of the amplifier alone showed that it was unity gain stable. The cause of this instability proved to be a phase shift introduced by the RC feedback network created by the pixel parasitic capacitance, the integration capacitor, and the closed state resistance of the reset switch. The simulated gain and phase information is measured after the feedback network with the ϕ_{rst} switch closed to include any phase margin degradation this configuration produces. Also, the Correlated Double Sampling (CDS) circuit (which will be discussed in section 4.3.2) was active, *i.e.* pass gate and clamp switch open, for the minimal capacitive loading configuration. These simulations predict a DC differential gain of 91 dB with a unity gain frequency (gain bandwidth product) of 37.4 MHz and a corresponding phase margin of 57°.

4.2.1.4 Integrator Amplifier – Noise Performance

Because of the sensitivity required of the front-end, the noise performance of the integrator amplifier must be considered carefully. To present this analysis in a clear fashion, our discussion is divided between intrinsic sources of noise (noise generated by the amplifier alone) and the susceptibility of the design to extrinsic sources of noise (pick-up from fluctuations outside of the amplifier, *e.g.* power supply ripple). As a point of reference, one 10 keV x-ray produces a voltage shift of roughly 9.6 mV with an RMS uncertainty of 180 μV. Requiring the signal to noise ratio to be

at least three, at 10 keV, means that we can tolerate a noise level of up to 3.2 mV RMS.

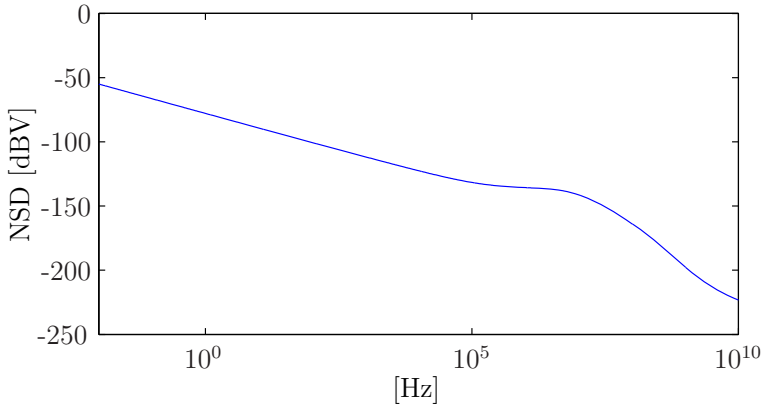

(a) Differential Noise Spectral Density

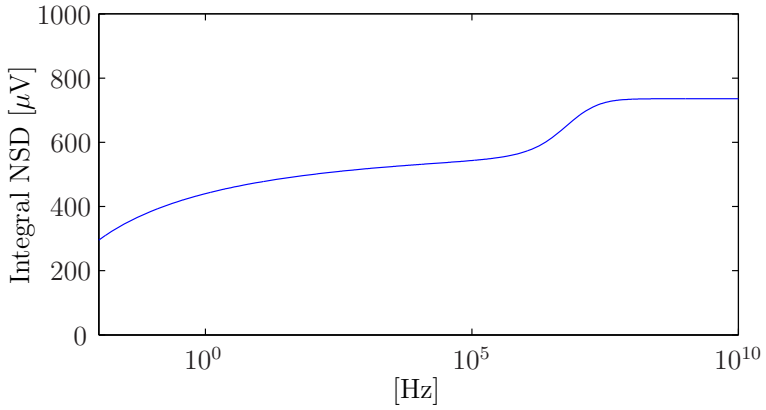

(b) Integrated Noise Spectral Density

Figure 4.9: Simulated noise power spectra for the Mixed–Mode PAD integrator amplifier. The first plot shows the differential NSD while the second shows the integrated NSD as a function of the sample & hold bandwidth, assuming a 100 second integration time.

The specifications presented in table 4.1 indicate that the total intrinsic noise from the amplifier should be less than 850 μV. This limit arises because the detector's thermal noise, *i.e.* kT/C noise, is expected to have a room temperature RMS value of 1.8 mV when observed at the output of the integrator (V_{outp}). An intrinsic noise contribution 850 μV increases the noise level of the front end to $\sim$2.0 mV, $\sim$10% above the detector's thermal noise baseline. This level was chosen to

Table 4.3: Thermal noise contributions from dominant amplifier noise sources. The *Integrated* column calculation assumes a 6 MHz bandwidth (to match the bandwidth of the sample and hold stage).

Transistor	Spectral [A^2/Hz]	Integrated [A^2]
M1, M2	5.19×10^{-25}	3.11×10^{-18}
M3, M4	4.10×10^{-25}	2.46×10^{-18}
M9, M10	4.27×10^{-25}	2.56×10^{-18}

Table 4.4: Flicker noise contributions from dominant amplifier noise sources. The *Spectral* column reports the value of equation 4.25 a 1 Hz.

Transistor	Spectral [A^2/Hz]	Integrated [A^2]
M1, M2	3.64×10^{-19}	4.54×10^{-18}
M3, M4	2.27×10^{-19}	2.83×10^{-18}
M9, M10	2.67×10^{-20}	7.28×10^{-19}

provide sufficient tolerance to achieve the desired sensitivity in spite of additional noise sources in the analog readout and digitization chain.

Typically, in amplifier design, one seeks to minimize the input referred noise to preserve the integrity of a transmitted voltage signal. However, for the Mixed–Mode PAD the situation differs in that the input signal is ideally a constant voltage. Because of the input voltage is static a more appropriate noise figure of merit for this application is the output referred noise, which is typically expressed as an RMS voltage. To calculate this level of fluctuation, one may express the noise contribution from each component in the amplifier as a current, propagate this

Table 4.5: Total noise contributions (RMS) from dominant amplifier noise sources. Combining these results yields a root mean square voltage noise from all amplifier sources of 600 μV.

Transistor	RMS Voltage Fluctuation [μV]
M1, M2	412
M3, M4	343
M9, M19	270

current to the output assuming an ideal circuit, and then multiply by the closed loop output impedance of the amplifier. Since the amplifier in the Mixed–Mode PAD integrator only uses linear current gain stages, the total output referred noise is the sum, in quadrature, of these component terms. This amplifier's closed loop output impedance ($Z_{\text{out,closed}}$) is given by

$$Z_{\text{out,closed}} = \frac{Z_{\text{out,open}}(\omega)}{1 + \beta A_{\text{open}}(\omega)}, \tag{4.21}$$

where, as discussed earlier in this section, $A_{\text{open}}(\omega)$ is the open loop gain of the amplifier, $Z_{\text{out,open}}(\omega)$ is the amplifier's open loop output impedance, and β is the gain of the feedback network,

$$\frac{1}{\beta} = 1 + \frac{C_{\text{pix}}}{C_{\text{int}}}. \tag{4.22}$$

In the limit $1 \ll \beta A_{\text{open}}(\omega)$ this reduces to

$$Z_{\text{out,closed}} \approx \frac{1}{\beta G_m} = \frac{C_{\text{int}} + C_{\text{pix}}}{G_m C_{\text{int}}}. \tag{4.23}$$

Within the amplifier two major noise sources were considered for each transistor: thermal, or white, noise whose *power spectral density* (PSD) is given by

$$\frac{\partial I^2_{n,,\text{ th}}}{\partial f} = \frac{8kT}{3} g_m, \tag{4.24}$$

and flicker, or $1/f$, noise whose PSD is given by

$$\begin{aligned} \frac{\partial I^2_{n,\text{ fl}}}{\partial f} &= \frac{K_f}{C_{\text{ox}} W L f^{A_f}} g_m \\ &\approx \frac{4\mu K_f}{L^2 f^{A_f}} I_{\text{DS}}, \end{aligned} \tag{4.25}$$

where the approximation is valid for transistors operating within the strong inversion regime.

Table 4.5 lists estimates of the noise contribution from each transistor that is expected to be a significant noise source. Based on these calculations, the

total noise contribution of the integrator amplifier is expected to be 600 μV RMS. For comparison, figure 4.9 shows the differential and integral output noise power simulated with TSpice. From these simulated results, the total noise contribution of the amplifier is expected to be 650 μV RMS.

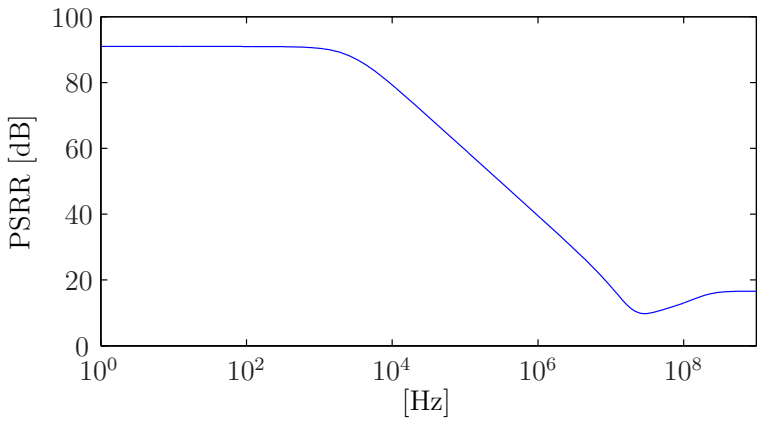

Figure 4.10: PSRR of Mixed–Mode PAD integrator amplifier in unity gain feedback configuration, *i.e.* reset switch closed.

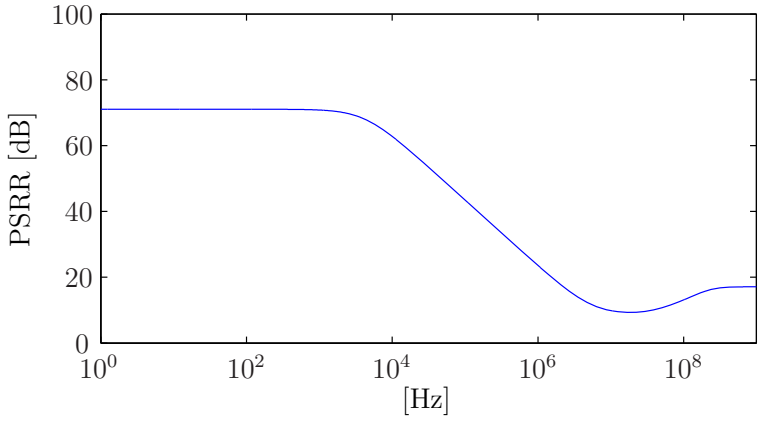

Figure 4.11: PSRR of Mixed–Mode PAD integrator amplifier in capacitive feedback configuration, *i.e.* reset switch open.

In addition to the amplifier's intrinsic noise, it is also necessary to consider the susceptibility of the amplifier to noise from extrinsic sources, most notably the power supply and voltage bias lines. Susceptibility to power supply pickup

is typically encapsulated in a design parameter called the *power supply rejection ratio* (PSRR), which describes the ratio of fluctuations on the power supply to the resulting fluctuations at a node of interest, *e.g.* the integrator output, in the circuit.

To evaluate the amplifier's susceptibility to power supply pickup, two feedback configurations are considered: the unity gain configuration that occurs when the ϕ_{rst} signal is low, representing the noise that is sampled when the reset switch is opened; and the capacitive feedback configuration that occurs when ϕ_{rst} is high, representing the noise sampled by the sample and hold at the end of the integration (see figure 4.3). The unity gain PSRR was 90.9 dB, well above the design specification. Early simulations of versions of the folded cascode architectures, which notably differed for the architecture shown in figure 4.4 by connection of the transistor bulk of the differential pair transistors (M1 and M2) to the analog supply voltage (V_{DDA}) rather than their source, in the capacitive feedback configuration yielded the very poor PSRR of $3.87\,[\mathrm{dB}] \approx 1.56$. Although this problem was fixed in the final design, the reason for the poor PSRR in this feedback configuration is worth discussing, because it is not a feature of the folded cascode architecture, in particular, but a problem endemic to any PAD using an integrator type front end.

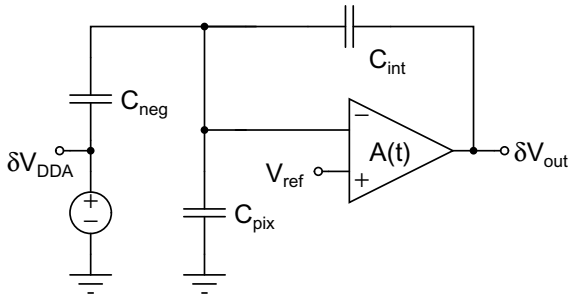

Figure 4.12: Model system for analyzing the effect of capacitive coupling between V_{DDA} and the inverting input of the amplifier.

The problem stems from fluctuations in the source, drain, and bulk voltages of

the transistor connected to the inverting input ($V_{\rm neg}$) of the amplifier differential pair,[9] caused by fluctuations of $V_{\rm DDA}$, capacitively coupling into the inverting input of the amplifier. To see how this parasitic coupling occurs, consider the model system shown in figure 4.12 where we assume an amplifier with an ideal PSRR and model the effect of fluctuation on the analog supply ($\delta V_{\rm DDA}$) coupling through the gate-to-bulk capacitance ($C_{\rm gb}$), gate-to-source capacitance ($C_{\rm gs}$), and gate-to-drain capacitance ($C_{\rm gd}$) of the transistor at the inverting input of the amplifier differential pair with a single explicit capacitor,

$$C_{\rm neg} = C_{\rm gb,M1}\frac{\partial V_{\rm b,M1}}{\partial V_{\rm DDA}} + C_{\rm gs,M1}\frac{\partial V_{\rm s,M1}}{\partial V_{\rm DDA}} + C_{\rm gd,M1}\frac{\partial V_{\rm d,M1}}{\partial V_{\rm DDA}}, \tag{4.26}$$

between the analog supply voltage ($V_{\rm DDA}$) and the inverting input ($V_{\rm neg}$) of the amplifier. The impedance seen through this capacitance, looking into the circuit, is

$$\begin{aligned} Z_{\rm in} &= \frac{1}{i\omega C_{\rm neg}} + \left(\frac{1}{i\omega C_{\rm int}(A+1)}\Big|\Big|\frac{1}{i\omega C_{\rm pix}}\right) \\ &= \frac{C_{\rm neg} + C_{\rm pix} + C_{\rm int}(A+1)}{i\omega C_{\rm neg}(C_{\rm pix} + C_{\rm int}(A+1))} \\ &\approx \frac{1}{i\omega C_{\rm neg}}, \end{aligned} \tag{4.27}$$

for $C_{\rm neg} \ll C_{\rm pix} + C_{\rm int}(A+1)$, where A is the gain of the amplifier at the frequency ω (and $i = \sqrt{-1}$). Thus, the power supply fluctuation $\delta V_{\rm DDA}$ will induce a current $\delta I \approx i\omega C_{\rm neg}\delta V_{\rm DDA}$ flowing onto the inverting input of the amplifier. Based on results derived in appendix B, this coupling will induce a fluctuation in the amplifier

[9] Transistor M1 in figure 4.4, except that we will consider here what happens when the bulk of this transistor ($V_{\rm b,M1}$) is connected to the analog supply ($V_{\rm DDA}$) rather than ($V_{\rm s,M1}$) the transistor source, as in this schematic.

output (V_{out}) of

$$\begin{aligned}\delta V_{\text{out}} &= -\frac{C_{\text{neg}}A}{C_{\text{pix}} + C_{\text{int}}(A+1)}\delta V_{\text{DDA}} \\ &\approx -\frac{C_{\text{neg}}}{C_{\text{int}}}\delta V_{\text{DDA}}, \qquad (4.28)\\ \Rightarrow \left|\frac{\partial V_{\text{DDA}}}{\partial V_{\text{outp}}}\right| &\approx \frac{C_{\text{int}}}{C_{\text{neg}}}, \qquad (4.29)\end{aligned}$$

for $1 \ll A$ and $C_{\text{pix}} \ll C_{\text{int}}A$.

To compare this analysis with TSpice simulations, we first must calculate C_{neg}. Since the bulk of the differential pair is connected to V_{DDA} its voltage coupling factor is simply

$$\frac{\partial V_{\text{b,M1}}}{\partial V_{\text{DDA}}} = 1. \qquad (4.30)$$

The voltage coupling factors between the analog supply source of transistor M1 and the analog supply drain of transistor M1 would be difficult to derive analytically because of their presence in the feedback loop. However, they are readily available from TSpice simulations and found (respectively) to be

$$\frac{\partial V_{\text{s,M1}}}{\partial V_{\text{DDA}}} = 2.66 \times 10^{-1}, \qquad (4.31)$$

$$\frac{\partial V_{\text{d,M1}}}{\partial V_{\text{DDA}}} = 3.20 \times 10^{-4}. \qquad (4.32)$$

The capacitances C_{gb}, C_{gs}, and C_{gd} of transistor M1 are given by TSpice small signal analysis to be 29.8 fF, 8.4 fF, and 8.4 fF, respectively. Combining these results gives $C_{\text{neg}} = 32$ fF. From our earlier analysis, the PSRR will be $\frac{\delta V_{\text{DDA}}}{\delta V_{\text{outp}}} \approx \frac{C_{\text{int}}}{C_{\text{neg}}} = \frac{50}{32} = 1.56 \approx 3.88\,[\text{dB}]$, which is in exceptionally good agreement with the directly simulated PSRR.

The value of C_{int} and the parameters comprising C_{neg} are restricted by other design considerations, so there is little to be gained trying to optimizing the design to minimize the ratio of $\frac{C_{\text{neg}}}{C_{\text{int}}}$. A more effective way to recover the PSRR is to remove the direct coupling between V_{DDA} and V_{neg} by connecting the bulk of the

differential pair to the source of the differential pair. This change results in a TSpice simulated PSRR of 69.08 [dB] = 2, 844. This change will require sacrificing some area, since the differential pair will need a separate N-well that is protected against the formation of parasitic field channels by a guard ring. However, given the dramatic improvement in PSRR it is a worth-while exchange.

Plots of the simulated PSRR for the amplifier under unity gain and capacitive feedback conditions are given in figures 4.10 and 4.11, respectively. Under both feedback configurations, the coupling between the power supply and integrator output is reduced to less than 0.1% for any fluctuations on time scales above 1 ms.

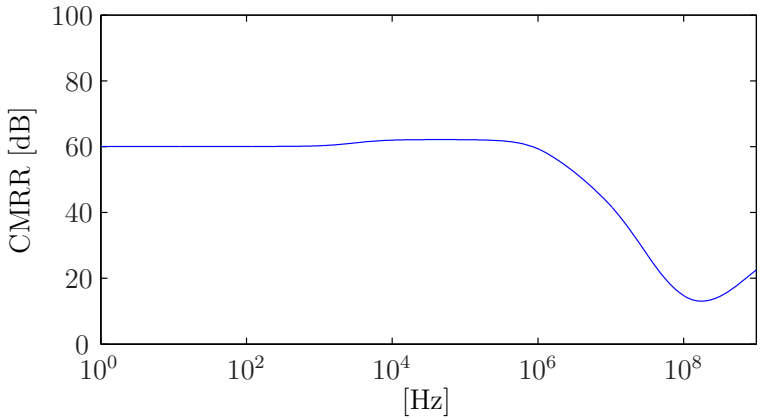

Figure 4.13: Simulated CMRR for Mixed–Mode PAD integrator amplifier.

In addition to the power supply, there are two other classes of control voltages that offer paths for pick-up in the integrator amplifier. These are the amplifier's bias control and reference voltages. Pick-up on the bias lines is not a great concern as it will only slightly alter the operating point of the amplifier but should not directly affect the output voltage. Fluctuations in V_{ref} are more of a concern.

Fluctuations on the bias control voltages do not strongly couple into the amplifier output voltage as these fluctuations either induce equal current shifts on the inverting and non-inverting branches of the circuit that cancel at the output, as is the case of the nMOS current source control voltage (V_{bn}) and pMOS current

source control voltage (V_{bp}), or are compensated for automatically by the amplifier, as is the case with the nMOS cascode bias (V_{cn}). The main impact of these fluctuations is to slightly shift the operating point of the amplifier. However, as long as the set point for the integrator leaves a reasonable tolerance for these shifts they should not create a problem.

Fluctuations on the integrator's reference voltage (V_{ref}) are more problematic. Depending on the feedback configuration, these either act as a common mode fluctuation or are boosted by the gain of the feedback network. The first condition arises when the amplifier is configured as a unity gain follower[10] which occurs when the integrator's reset is active. A plot of the modeled Common Mode Rejection Ratio (CMRR) as a function of frequency is shown in figure 4.13. The second condition arises when the reset is not active. In this case, the integration capacitor (C_{int}) forms a capacitive feedback network with the parasitic capacitance at the pixel front end resulting in a feedback gain of $\sim \frac{1}{\beta} = \frac{C_{\mathrm{int}}+C_{\mathrm{pix}}}{C_{\mathrm{int}}} \cong 5$. To minimize the effect for both of these feedback configurations, we low-pass filter the V_{ref} line, both on-chip and off.

4.2.1.5 Radiation Tolerance

As with noise performance, the tight coupling between the integrator amplifier and the overall performance of the analog front end warrants a detailed discussion of the effects of radiation damage on this circuit. Based on the discussion from section 2.3.2, there are two primary ways this amplifier architecture is susceptible to radiation damage:

- Bias current reduction due to radiation induced increase in pMOS transistor

[10] Technically, fluctuations on V_{ref} within the amplifier unity gain bandwidth directly propagate to V_{outp}. However our off-chip ADC system uses a differential front end with the reference voltage set to V_{ref} so that the actual quantity digitized is $V_{\mathrm{ref}} - V_{\mathrm{outp}}$.

thresholds.

- Increased leakage currents in the nMOS devices due to radiation induced reduction in nMOS thresholds and formation of parasitic channels.

The network that generates gate voltages for the amplifier bias transistors relies on a current input; therefore, any shifts in transistor threshold or leakage that are common to both pixels and bias network will be automatically accounted for. Radiation induced shifts in bias currents only occur if there is a difference in the total dose between the pixel and the bias network.

In many experiments, most notably High Energy Physics and space applications, it is safe to assume that chip will be uniformly irradiated so that this effect may be neglected. However, in experiments where the radiation dose is localized to specific portions of the chip, this assumption is not necessarily valid. Crystallography and synchrotron radiography, two likely applications of the Mixed-Mode PAD detector, are examples of this class of experiment. With regard to crystallography, one cannot assume uniform irradiation because crystal diffraction patterns produce localized spots or rings of x-ray intensity. While the location of these spots or rings may vary from experiment to experiment, experimenters will try to optimize their setup so that the majority of the signal is collected in the active area of the array. This bias will result in a disparity between the total dose accumulated at pixels in the array and the total dose accumulated by the bias network on the ASIC periphery. In synchrotron radiography, the problem is more pronounced due to the limited extent of the beam. As the beam footprint is typically smaller than a single chip with minimal divergence, the illuminated regions within the array will receive a total dose orders of magnitude higher than the bias network. As a final, more general, point, within the operating range of energies for the Mixed-Mode PAD the detector layer significantly attenuates the x-ray beam. The bias network,

though, is located at the edge of the ASIC near the wire bonding pads so that it does not benefit from this additional protection. Thus, even schemes to distribute the radiation more evenly across the pixel array and bias network would yield significantly mismatched total doses. For these reasons, the amplifier design must be robust against radiation induced shifts in its bias currents.

Reduction of the bias current is problematic because it reduces the transconductance of the input transistors as well as the current available to slew the output. The impact of these parameters on design performance was discussed previously in this section. The radiation induced shift is compounded by a potenital power supply droop of up to 30 mV from the wire bonded side of the pixel array to the far, opposite edge. Both radiation damage and power supply droop can be modeled as an increase in the threshold of the bias transistor.[11] To guard against this droop the sizing of the bias current generating transistor was chosen so that, over the range of operating currents the ratio g_m/I_{ds}, was as small as reasonably possible. This ratio allows one to estimate a percent change in I_{ds} for a given threshold shift. Based on published radiation studies of the TSMC 0.25 μm process [58], radiation studies performed at Cornell and reported in section 6.7, and studies of the power supply droop, we expect that the total threshold shift will be less than 50 mV at one megarad dose in the oxide. Assuming the bias transistor is operating in strong inversion (which it should be to minimize g_m/I_{ds}),

$$g_{m,\mathrm{M0}}/I_{ds,\mathrm{M0}} = 2\sqrt{\frac{WK'_p}{LI_{\mathrm{ioa}}}}.$$

Using K'_p from [94] and a 5 μA bias current one finds $g_{m,\mathrm{M0}}/I_{ds,\mathrm{M0}} \approx 6.7$ V^{-1} so that a 50 mV shift would reduce the bias current by roughly 34%, to 3.3 μA.

The effect of radiation induced leakage on the nMOS bias transistors has the potential to harm the amplifier's slewing capabilities. If the leakage current is

[11] In the case of supply droop one can assume an ideal (droop less) supply and increase the transistor threshold by the amount of the droop.

denoted I_{rad}, then the maximum output current will be limited to

$$\max(I_{\text{out}}) = I_{\text{ioa}} - I_{\text{rad}}.$$

The value of I_{rad} is, unfortunately, not a simple function of accumulated radiation dose. As reported in [58], it is also influenced by the time over which the dose is acquired and the environment of the chip during and between exposures.

In order to push the amplifier out of slew specifications, the maximum output current would have to be reduced to below 2.5 μA. This limit means that, starting from the nominal operating point of 5 μA and assuming a worst case radiation and power supply droop reduced bias current of 3.3 μA,[12] then 0.8 μA of nMOS leakage would be needed to pull the amplifier out of slew specification. Based on the measurements reported in [58], it is unlikely that this level of damage is attainable in the normal operating life of a hybrid.[13]

Arguably, if the nMOS bias devices used an enclosed layout structure, as discussed in section 2.3.2, it would mitigate most potential leakage problems. However, small W:L ratios, *e.g.* the 20:80 ratio used in the current design, are unattainable with these devices. Increasing the W:L ratio so that Enclosed Layout Transistor (ELT) devices could be used would degrade the frequency response, by lowering the high frequency poles and thus compromising the phase margin, as well as worsen the noise performance of the integrator amplifier. Because of these considerations, linear devices are used for the bias transistors.[14]

[12] As discussed previously in this section.

[13] The largest leakage level, at 1 MRad, reported by [58] was over an order of magnitude below this value for a minimum sized nMOS device. As noted in [58], increasing the time over which the dose is incurred will significantly reduce the effect of the radiation damage as will the use of longer device.

[14] More recently, our group has looked at a linear radiation hardened transistor. This device requires substantial extra perimeter area than a standard linear transistor, so it is not particularly good for digital circuitry, but it is a strong candidate building long nMOS devices. Details on this structure may be found in [26].

4.2.1.6 Integrator Linearity

We conclude our discussion of the amplifier by returning to its performance in the integrator, specifically looking at the linearity of the integrator's response to a constant current source. The design specification calls for a 1.0 V linear operating range for the analog residual voltage. This limit is based on the fact that we digitize with a 10–bit ADC and would like precision at the 1 mV level. Panel (a) of figure 4.14 shows a simulation of the integrator linearity in which a constant current source is applied to the integration node of the circuit. Panel (b) of this figure depicts the integrator's deviance from the ideal linear response as a function of the integrator's output voltage. Two linear regions exhibiting different slopes are evident in this plot, separated by a kink at ~ 0.8 V. The lower region ($V_{\text{outp}} < 0.8$ V) results from the cascoded pair of transistors, M8 and M10, being driven into their ohmic region. The resulting reduction in the output impedance leads to a proportional decrease in the amplifier's DC gain, in accordance with equation 4.7. The decrease in gain changes the effective capacitance of the integrator resulting in increased charge sharing with the pixel's parasitic capacitance.

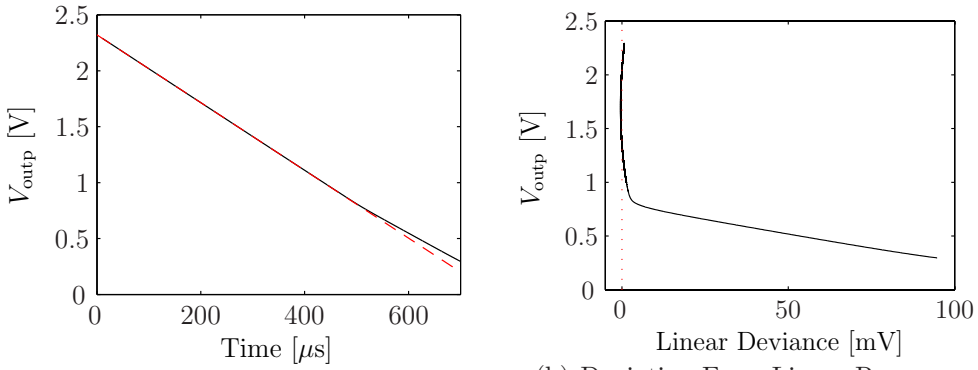

(a) Integrator Output vs Signal

(b) Deviation From Linear Response

Figure 4.14: Panel (a) shows a simulation of the change in the integrator output over time in response to a constant signal current allow with a dashed line showing the ideal response. Panel (b) shows the deviation of the simulated response from the ideal response.

As the upper linear range (0.8 V $< V_{\text{outp}}$) spans more than 1.5 V the design specification is easily met with this amplifier and operating conditions. Although not anticipated, if in the future a larger linear range were to be needed, it could be achieved by reducing the integrator amplifier bias current.

4.2.2 Quantized Charge Removal

Another critical component in the Mixed-Mode PAD analog front end is the quantized charge removal circuit used to accomplish the Δ-portion of the $\Sigma\Delta$-operation discussed in section 4.1. This circuit performs the task of removing a fixed quantity of charge from the integration node whenever sufficient charge has accumulated to trigger a removal operation. The digital logic details of the charge removal, specifically the trigger conditions initiating a removal operation and the details of the controller circuit, will be discussed in section 4.2.2.3. Here, we focus the discussion and analysis on the architecture of the analog components of the charge removal system.

The Mixed-Mode PAD uses a switched capacitor circuit to perform the quantized charge removal, as outlined in figure 4.15. Under normal circuit operation (*i.e.* when a removal operation is not occurring), the charge removal control clock (ϕ_{rem}) is high so that the switch SW1 is open while SW2 is closed. When a charge removal occurs SW1 closes and SW2 opens, effectively shorting the charge removal capacitor (C_{rem}) to the front end. This process results in a charge cancellation whereby a total charge (holes) of $\Delta Q = C_{\text{rem}}(V_{\text{ref}} - V_{\text{low}})$ that had accumulated on the integrator cancels with the charge (electrons) supplied by the charge removal capacitor.

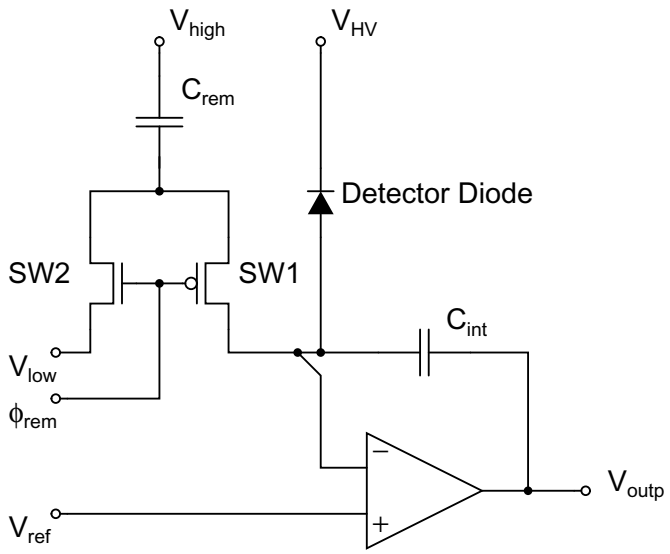

Figure 4.15: Schematic of the switched capacitor quantized charge removal circuit found in the analog front end of each pixel. This circuit performs the Δ-portion of the $\Sigma\Delta$-operation discussed in section 4.1.

4.2.2.1 Analog Components: The Gory Details

While *prima-facia* this circuit seems very straightforward, it directly connects to the most sensitive node in the pixel. Because of this connection, there are a number of subtle details that need to be attended to so that it is kept from degrading the performance of the analog front end. Foremost among these details is the choice of the passive (V_{high}) and active (V_{low}) reference voltages. These references need to be chosen with care to ensure the stability of the design, to minimize noise contribution and leakage onto or from the integration node, and to allow for accurate control over the quantity of the charge removed. Secondly, certain care needs to be taken in the design of the switches (SW1 and SW2) that remove the charge to avoid leakage current onto the integration node.

The passive reference node (V_{high}), so called because it does not draw current and should ideally remain at a constant voltage throughout the integration cycle, is capable of injecting charge into the analog front end whenever SW1 is closed.

Generally speaking, a charge of

$$Q = C_{\mathrm{rem}}(V_{\mathrm{high}} - V_{\mathrm{ref}}), \tag{4.33}$$

will remain on the removal capacitor when the charge removal operation ends. If δV_{ref} describes fluctuations of V_{ref} around its mean and δV_{high} describes the resp. fluctuation of V_{high} then, assuming these are independent, the resulting charge fluctuation in removed charge will be

$$\delta Q = C_{\mathrm{rem}}(\delta V_{\mathrm{high}}^2 + \delta V_{\mathrm{ref}}^2)^{\frac{1}{2}}. \tag{4.34}$$

To reduce this effect, we introduce a correlation between these noise sources by using a copy of V_{ref} to produce V_{high}. If V_{ref} and V_{high} are copies of the same signal then common fluctuations cancel so that $\delta Q \approx 0$.

The active reference node (V_{low}), so called because it draws current during portions of the integration cycle, likewise requires care in design. This node also contributes to fluctuation in the amount of charge removed in each charge removal cycle. When the switch SW2 opens any fluctuation in the difference of V_{low} and V_{high} will result in a fluctuation of the removed charge. As with SW1, this charge fluctuation may be reduced by correlating fluctuations on V_{low} and V_{high}. Correlations can be used to effectively remove external noise sources (*e.g.* pick-up from the power supply) if the reference voltages (V_{low} and V_{high}) are generated with identical, on chip digital-to-analog converters (DACs). In this case, the mirror pair of DACs, integrated onto the ASIC and driving similar loads, share both analog control signals and the transfer functions by which fluctuations on these control lines propagate to DAC outputs. Therefore, these common fluctuations should not induce any fluctuation in the removed charge.

These steps, though, do not address noise sources within these circuits or the effect of charge injected into the active node during every charge removal operation. With regard to these issues, the best recourse is to minimize the impedance

of these nodes through capacitive coupling, either with each other or ground. High frequency coupling is directly facilitated by the array of charge removal capacitors (C_{rem}, presenting $\sim$ 8 pF of capacitance across the array) while additional capacitance is available off-chip.

A final detail that needs to be considered with regard to this circuit is the leakage path presented by the switches SW1 and SW2. The TSMC 0.25 μm process used to manufacture our signal processing ASIC is primarily a digital process. As a result, this technology uses ion implantation to lower threshold voltage of its nMOS and pMOS devices to reduce switching times. This practice leads to substantially higher leakage than one would expect based on the reverse diode leakage of a minimal sized source or drain diffusion. More precisely, according to the manufacturer's specifications, these devices may exhibit leakage levels as high as 1 pA, a level that would be problematic for our front end, whereas leakage from the diffusions alone should be less than 1 fA. In addition, the long term effect of radiation on the nMOS switch (SW1) will be to further lower the effective threshold of the device and induce parasitic leakage paths around the edges of transistor gate.

Two techniques are used to minimize the leakage along this path. First, the nMOS switch uses an enclosed layout structure[15] to prevent the formation of parasitic leakage paths around the transistor gate. Second, the leakage induced by threshold lowering, either from ion implantation or radiation damage, may be stopped by raising the V_{low} above the bulk voltage (V_{GNDA}). Based on arguments outlined in [63], so long as the V_{low} is set no lower than $\sim$ 0.5 V, leakage through SW1 should be negligible. By extension, this same argument implies that leakage across SW2 should also be negligible.

[15]See section 2.3.2 for a discussion of enclosed layout transistors.

4.2.2.2 A Question of Fidelity: The Pixel Virtual Ground

As discussed in section 4.2.1, the integration node of the pixel (*i.e.* the point where the signal processing circuit connects to the detector diode) acts as a virtual ground due to the pixel's integrator. An important consideration that combines the topics of the previous two sections is the fidelity of this virtual ground. This voltage level needs to be maintained over a wide range of input conditions; this requirement becomes particularly important during the charge removal operation, because we rely on this node's fidelity to ensure the repeatability of the quantity of charge removed.

Traditionally, a stable virtual ground is obtained from an integrator by maximizing the integrator amplifier gain, often through the use of multistage amplifiers. In the Mixed–Mode PAD pixel, this approach alone is not sufficient due to the confluence of the limited settling time available in each charge removal operation and the limited frequency response available from high-gain two-stage amplifiers. The Mixed–Mode PAD design parameters require that each charge removal cycle occur in less than 1 μs, to meet the detector's maximum flux specification, which allows only a fraction of this time, 50% with typical settings, for the charge removal. Because of this timing requirement, the amplifier must have sufficient response at high frequencies to settle the virtual ground within a half a microsecond.

The remainder of this section begins by outlining a mathematical tool to estimate the charge removal performance of a given amplifier architecture using commonly available amplifier parameters, specifically the gain and frequency information contained in the amplifier's Bode plot. Using this tool, we then present estimates for the fraction of charge accounted for in each charge removal cycle as a function of the duration of the charge removal.

Very generally, to calculate the change in charge across the integration capacitor

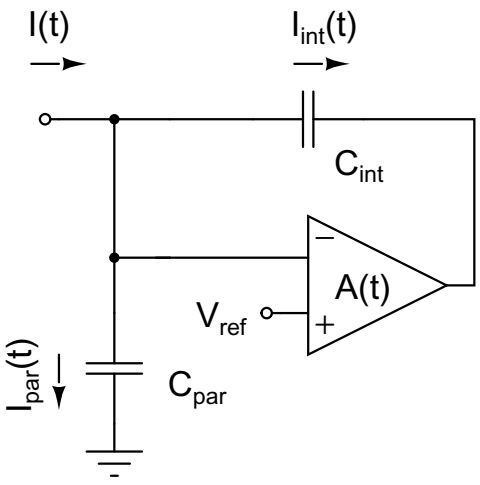

Figure 4.16: Analog input model used to derive the current transfer function, $\hat{H}(\omega)$. This model lumps the capacitance of the charge removal capacitor (C_{rem}) into the parasitic front end capacitance (C_{par}).

in response to current flowing onto or from the integration node we may begin with a time domain description of the current ($I(t)$) and a network response function ($H(t)$) describing the integrator's response to a unit current impulse. The time dependent flow of current onto or off of the integration capacitor is then given by the convolution of these functions, $H(t) * I(t)$.

A case of particular interest for our front-end architecture involves how the front end responds during a charge removal cycle. As was noted in equation 4.33, when a charge removal cycle ends some charge may remain on the charge removal capacitor. This equation, however, assumes that the charge removal operation allows enough time before the cycle ends for the integrator to re-establish its virtual ground level. A more general form of equation 4.33 would be

$$Q = C_{\text{rem}}(V_{\text{high}} - V_{\text{pix}}), \tag{4.35}$$

where V_{pix} is the voltage on the integrator node. This voltage will fluctuate during the charge removal process as the amplifier reacts to the sudden cancellation of charge caused by switching in the charge removal capacitor. Using the convolution

method outlined above, it is possible to assess this response by first integrating the area under the $H(t) * I(t)$ curve over the region of time when the charge removal operation is active to get the total charge drawn from the integration capacitor. Then, the charge not drawn from the integration capacitor ($Q_{\rm res}$) will cause a deviation in $V_{\rm pix}$ of

$$V_{\rm pix} - V_{\rm ref} = \frac{Q_{\rm res}}{C_{\rm int} + C_{\rm par}}, \tag{4.36}$$

where $C_{\rm par}$ is defined in figure 4.16, so that the resulting error in charge removed ($\Delta Q_{\rm err}$) will be

$$\Delta Q_{\rm err} = \left(\frac{C_{\rm rem}}{C_{\rm int} + C_{\rm par}} \right) Q_{\rm res}. \tag{4.37}$$

To determine the time dependence of the charge removal response, we will first present an analysis assuming that the circuit is operating in its linear range, subsequently treating what happens when it is not. To carry out an analysis of the linear range response a particular integrator amplifier architecture's virtual fidelity, we will need a few mathematical tools. To derive these, begin by defining the boxcar windowing function,

$$\Theta_{\tau_{\rm rem}}(t) = \begin{cases} 1, & t \in [-\tau_{\rm rem}, \tau_{\rm rem}], \\ 0, & \text{otherwise}, \end{cases} \tag{4.38}$$

where $\tau_{\rm rem}$ represents half the charge removal time. With this function, we can isolate the response of the integrator during time when a charge removal is occurring in the pixel; *i.e.,*

$$\begin{aligned} \Delta Q - Q_{\rm res} &= \int_{-\infty}^{\infty} dt\, \Theta_{\tau_{\rm rem}}(t) \cdot \{H(t) * I(t)\} \\ &= \int_{-\infty}^{\infty} dt\, \Theta_{\tau_{\rm rem}}(t) \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega\, \hat{H}(\omega) \cdot \hat{I}(\omega) e^{i\omega t}, \end{aligned} \tag{4.39}$$

where $\hat{H}(\omega)$ and $\hat{I}(\omega)$ are respectively the Fourier transforms of the current transfer function and the input current, and we have cast the problem in terms of how much

charge is accumulated onto the charge removal capacitor during the time the charge removal circuit is active. We model the effect of switching in the charge removal capacitor as an impulse current source of magnitude $\Delta Q = C_{\text{rem}}(V_{\text{ref}} - V_{\text{low}})$, so that,

$$I(t) = \Delta Q \delta(t). \tag{4.40}$$

Fourier transforming this equation gives

$$\hat{I}(\omega) = \Delta Q. \tag{4.41}$$

A precise form of $H(t)$ will be presented shortly. For our current purposes, though, it is sufficient and more convenient to express $H(t)$ as its Fourier transform, written in its polar form as,

$$\hat{H}(\omega) = \hat{a}(\omega) e^{i\hat{p}(\omega)}. \tag{4.42}$$

Substituting equations 4.41 and 4.42 into equation 4.39, we find that

$$\begin{aligned}
\Delta Q - Q_{\text{res}} &= \int_{-\infty}^{\infty} dt\, \Theta_{\tau_{\text{rem}}}(t) \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega) e^{i\hat{p}(\omega)} \Delta Q e^{i\omega t} \\
&= \frac{\Delta Q}{2\pi} \int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega) e^{i\hat{p}(\omega)} \int_{-\infty}^{\infty} dt\, \Theta_{\tau_{\text{rem}}}(t) e^{i\omega t} \\
&= \frac{\Delta Q}{2\pi} \int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega) e^{i\hat{p}(\omega)} \int_{-\tau_{\text{rem}}}^{\tau_{\text{rem}}} dt\, e^{i\omega t}.
\end{aligned} \tag{4.43}$$

Integrating the temporal integral gives us that

$$\begin{aligned}
\Delta Q - Q_{\text{res}} &= \frac{\Delta Q}{2\pi} \int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega) e^{i\hat{p}(\omega)} \cdot \left[\frac{e^{i\omega\tau_{\text{rem}}}}{i\omega} - \frac{e^{-i\omega\tau_{\text{rem}}}}{i\omega} \right] \\
&= \frac{\Delta Q}{\pi} \int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega) e^{i\hat{p}(\omega)} \frac{\sin(\omega\tau_{\text{rem}})}{\omega} \\
&= \frac{\Delta Q}{\pi} \int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega) \left[\cos(\hat{p}(\omega)) + i \sin(\hat{p}(\omega)) \right] \frac{\sin(\omega\tau_{\text{rem}})}{\omega} \\
&= \frac{\Delta Q}{\pi} \int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega) \cos(\hat{p}(\omega)) \frac{\sin(\omega\tau_{\text{rem}})}{\omega} \\
&\qquad + i \frac{\Delta Q}{\pi} \int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega) \sin(\hat{p}(\omega)) \frac{\sin(\omega\tau_{\text{rem}})}{\omega}.
\end{aligned} \tag{4.44}$$

Now, recalling that the transform's magnitude ($\hat{a}(\omega)$) will be a even function of ω while its phase ($\hat{p}(\omega)$) will be odd, we see that the second (complex) integral is an odd function integrated over a symmetric domain. Therefore, we may eliminate it to get that

$$\begin{aligned}\Delta Q - Q_{\text{res}} &= \frac{\Delta Q}{\pi}\int_{-\infty}^{\infty} d\omega\, \hat{a}(\omega)\cos(\hat{p}(\omega))\frac{\sin(\omega\tau_{\text{rem}})}{\omega} \\ &= \frac{2\cdot\Delta Q}{\pi}\int_{0}^{\infty} d\omega\, \hat{a}(\omega)\cos(\hat{p}(\omega))\frac{\sin(\omega\tau_{\text{rem}})}{\omega},\end{aligned} \tag{4.45}$$

where the final result was reached by noting that the argument of the remaining integral is an even function over a symmetric domain.

As for the exact form of $H(t)$, it can be derived by observing that when a capacitor is placed across an amplifier, as C_{int} is in the integrator shown in figure 4.16, the amplifier has the effect of boosting the capacitor capacitance by a factor of $1+\hat{A}(\omega)$. At a given frequency, the fraction of the source current accumulated onto the charge removal capacitor is given by the divider ratio

$$\hat{H}(\omega) = \frac{C_{\text{int}}(\hat{A}(\omega)+1)}{C_{\text{par}} + C_{\text{int}}(\hat{A}(\omega)+1)}. \tag{4.46}$$

To take this analysis further we need performance parameters, specifically frequency dependent gain and phase response, for the particular amplifier we wish to analyze. Using the Mixed–Mode PAD integrator amplifier's open loop gain ($\alpha(\omega)$) and phase ($\phi(\omega)$) information, a portion of which is shown if figure 4.8, we can compute the network response function's gain and phase response,

$$\hat{a}(\omega) = \left|\frac{C_{\text{int}}(\alpha(\omega)e^{i\phi(\omega)}+1)}{C_{\text{par}} + C_{\text{int}}(\alpha(\omega)e^{i\phi(\omega)}+1)}\right| \tag{4.47}$$

$$\hat{p}(\omega) = \angle\left(\frac{C_{\text{int}}(\alpha(\omega)e^{i\phi(\omega)}+1)}{C_{\text{par}} + C_{\text{int}}(\alpha(\omega)e^{i\phi(\omega)}+1)}\right), \tag{4.48}$$

with $C_{\text{int}} = 50\,[\text{fF}]$ and $C_{\text{par}} \approx 200\,[\text{fF}]$ in the Mixed–Mode PAD pixel design. These results were then plugged into the integrand of equation 4.45,

$$\hat{a}(\omega)\cos(\hat{p}(\omega))\frac{\sin(\omega\tau_{\text{rem}})}{\omega}, \tag{4.49}$$

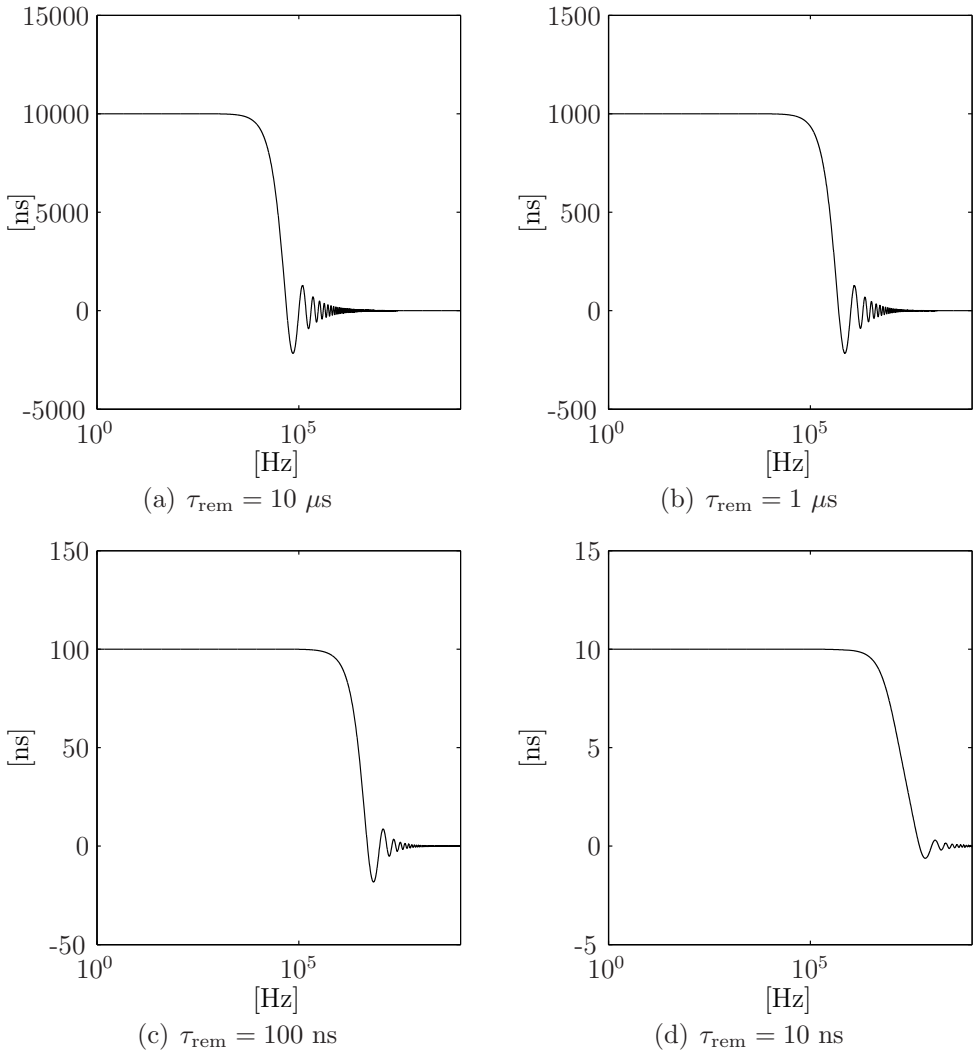

(a) $\tau_{\rm rem} = 10$ μs (b) $\tau_{\rm rem} = 1$ μs

(c) $\tau_{\rm rem} = 100$ ns (d) $\tau_{\rm rem} = 10$ ns

Figure 4.17: Examples of the current transfer function (the integrand of equation 4.45) for the Mixed–Mode PAD integrator amplifier at four different $\tau_{\rm rem}$ values.

to compute the current transfer function. Figure 4.17 shows examples of the current transfer function for the Mixed–Mode PAD integrator amplifier at four different τ_{rem} values.

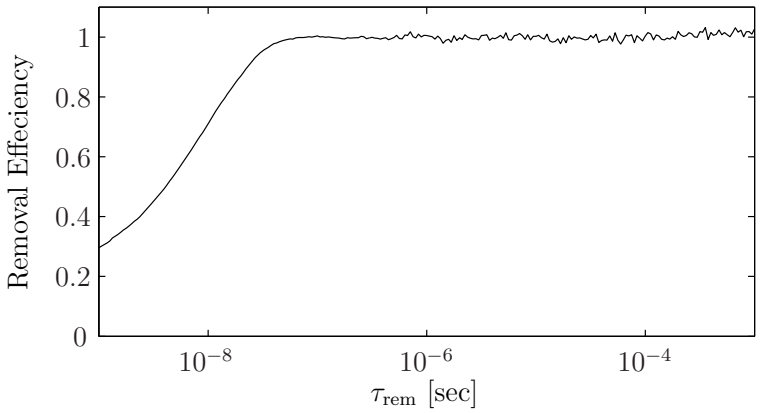

Figure 4.18: Fraction of charge accumulated onto the charge removal capacitor, for a pixel operating in the linear range, during a charge removal operation of duration $t_{\text{rem}} = 2\tau_{\text{rem}}$ using the Mixed–Mode PAD integrator amplifier. In most cases a few additional considerations are required because the quantity of charge removed will, temporarily, take the pixel out of the range of linear approximation. These considerations are outlined at the end of section 4.2.2.2.

Numerically integrating these current transfer functions and multiplying by $\frac{2}{\pi}$ gives the fraction of charge drawn onto the integrator during the charge removal operation. Figure 4.18 shows these results for the Mixed–Mode PAD over a range of τ_{rem}. From it, we can see that within less than 0.2 μs (recall that τ_{rem} is defined to be half of the charge removal time) over 99% of the impulse charge has been accumulated and that the quantity of charge removed is very stable above this limit.

One caveat of the proceeding analysis is that it assumes that the integrator amplifier is operating in its linear range. However, under the expected operating conditions, this will not be the case immediately after the charge removal cycle is initiated. The electron charge injected by the C_{rem} will be roughly 50 fC ($1\,[\text{V}] \cdot 50\,[\text{fF}]$) onto a total pixel capacitance of roughly 250 fF. This should result in

an immediate shift in the voltage of the pixel integration node relative to the integrator reference voltage ($V_{\rm pix} - V_{\rm ref}$) of 0.2 V. Although the integrator amplifier transconductance ($G_{\rm m}$) is nearly 50 μA/V, the slew rate is limited by the integrator amplifier bias current ($I_{\rm ioa}$) of 5 μA. Consequently the pixel will be slew rate limited until it reenters the linear range of the integrator amplifier; *i.e.* until roughly

$$G_m \cdot (V_{\rm pix} - V_{\rm ref}) = I_{\rm ioa}. \tag{4.50}$$

Conservatively, it will take the amplifier less than 50 ns[16] to slew ($V_{\rm pix} - V_{\rm ref}$) from 0.2 V to 0.1 V, after which point the fractional impulse response curve of figure 4.18 becomes applicable, now representing the time it takes to acquire a given fraction of the charge remaining after the initial slew response. To illustrate this for the operating conditions considered above, ∼50% of the charge will be accumulated in the first 50 ns while the integrator slews the shift in the pixel integration node from 0.2 V to 0.1 V. It will then take another ∼80 ns to accumulate better than 98% of the remaining charge so that over 99% of the impulse charge has been accumulated. This is notably longer than if the integrator were to remain entirely within its linear range, but still more than a factor of five shorter than our design specification of 0.5 μs.

4.2.2.3 Charge Removal Controller

The primary purpose of the charge removal controller is to ensure that a consistent quantity of charge is removed from the front end with each removal operation. From the discussion from section 4.2.2.2, stable charge removal is possible as long as the duration of the charge removal ($\tau_{\rm rem}$) is long enough to sufficiently settle the front end. Based on the discussion at the end of section 4.2.2.2, this constraint

[16] If all of the amplifier current of went to the integration capacitor it would only require 5 ns, however there is other capacitive loading of the integrator output which must be considered. The 50 ns limit assumes that the total capacitive loading of the integrator is less than $10 \cdot C_{\rm int}$.

implies that the charge removal controller needs to produce a control clock ($\phi_{\rm rem}$) with a duration of ~ 0.1 μs, or longer, whenever sufficient charge has accumulated on the integrator.

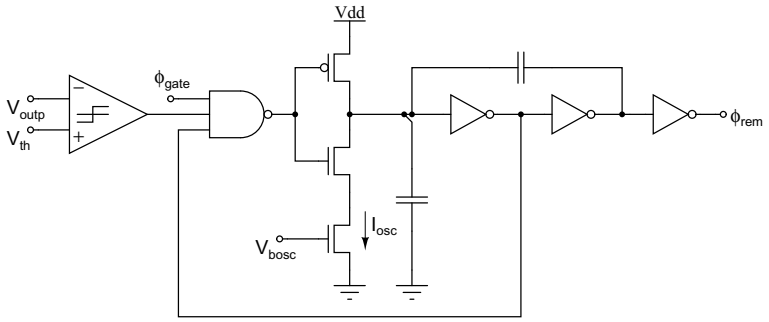

Figure 4.19: Schematic of the charge removal control circuit.

The circuit that produces this control clock is the comparator triggered gated oscillator shown in figure 4.19. Before explaining the details of this circuit, however, it is worth discussing the rationale behind using an oscillator as opposed to some form of single-pulse generator. Initial prototypes of the Mixed–Mode PAD pixel incorporated a circuit that generated a single pulse with each comparator crossing; yet, this circuit exhibits a potential lock-up condition if insufficient charge is removed during any charge removal operation to ensure that the integrator output is above the comparator threshold ($V_{\rm outp} > V_{\rm th}$) when the removal ends. Under normal operation, this should never pose a problem. However, there is a novel way of using the pixel, with the oscillator architecture, to acquire a sequence of frames with no associated readout dead time. Specifically, the architecture of the analog residual measurement, which will be discussed in section 4.3.2, allows one to take a non-destructive snapshot of the state of the pixel. If charge removal is then inhibited during digital readout, it is possible to continue the integration into the next frame without reseting the pixel. The problem comes in when the amount of charge accumulated during the read out exceeds twice what can be removed in

a single charge removal. With a single-pulse generator, this will cause the circuit to lock-up. However, with the oscillator, one may acquire, during the readout, nearly two full charge removals, with standard $V_{\rm ref}$ and $V_{\rm low}$ settings, or roughly 200 10 keV x-rays without degrading the pixel response—for a 1 ms readout, this is equivalent to a flux of 2×10^5 10 keV x-ray/s. This feature is quite attractive for high frame rate experiments that meet the flux requirements, where the time to read out the detector is a substantial fraction exposure duration.

To explain the operation of the charge removal control circuit: during the low phase of $\phi_{\rm rem}$ in each cycle, charge is removed from the integrator by connection of the charge removal capacitor to the integration node of the pixel. During the high phase, this connection is broken and the charge removal capacitor is reconnected to the $V_{\rm low}$ in preparation of the next removal cycle. As mentioned earlier, the trigger for the oscillator is provided by the comparator. Due to the delay built in to the oscillator digital feedback loop, once the trigger initiates a charge removal cycle the comparator has no influence on the oscillator until the full cycle is completed. If the comparator's input returns low or is still low when a cycle completes, then a subsequent cycle will be initiated, continuing in this fashion until sufficient charge has been removed from the integrator to raise $V_{\rm outp}$ above $V_{\rm th}$. Under normal circumstances, though, only a single cycle is needed to accomplish return the pixel to its normal operating state.

The duration of the charge removal cycle is controlled by the oscillator bias current ($I_{\rm osc}$). The capacitors shown in figure 4.19 are laid out to provide a capacitance of 0.1 pF for an effective capacitance of 0.2 pF. The duration of the charge removal operation is then

$$\tau_{\rm rem} = \frac{V_{\rm DDD}}{2I_{\rm osc}} \cdot 0.2\,[\mathrm{pF}]\,, \tag{4.51}$$

where $V_{\rm DDD}$ denotes the digital supply voltage, so that the bias currents ranging

from $\sim$ 200 nA to $\sim$ 2 μA yield performance within the operating specifications of the pixel front end. This broad range of acceptable settings gives the controller robustness against the pixel-to-pixel variations that occur across a large array—a result of device mismatch and power supply droop. In addition, this range helps to protect against additional variation induced by radiation damage.

As an nMOS device is used to generate I_{osc} there is a potential radiation damage concern. However, the sizing of the device ($W = 0.9$ μm, $L = 3$ μm) serves to minimize the radiation damage effects, as discussed in section 2.3, so that along with regular annealing, as will be discussed in section 6.7, this circuit is acceptably robust against radiation damage. Still, in future revisions of this design, radiation hardening through the use of the radiation hard linear transistors reported in the Ph.D. thesis of Dr. Alper Ercan [26], formerly of the Cornell PAD research group,[17] should be implemented to extend the long term reliability of the imager.

4.2.2.4 Charge Removal Conclusions

From the forgoing analysis of the Mixed–Mode PAD charge removal circuit, a number of important conclusions may be drawn regarding its impact on the performance of the Mixed–Mode PAD as an imager, particulalry in terms of its effect on the noise in the pixel signal measurement and acquisition of signal when the charge removal is active.

Although not presented here, a number of other charge removal architectures (*e.g.* reset-to-zero charge removal and constant-current/constant-time charge removal [34]) were investigated for the Mixed–Mode PAD front end. These, ultimately, were rejected because of the strong coupling they introduce between the stability of the timing circuit and the uncertainty in the charge removed. In the

[17] Unfortunately this layout technique is some what risky as it involves violation of MOSIS and TSMC design guidelines. Consequently there was not time in the Mixed–Mode PAD project schedule to vet a layout with this change in time to incorporate it into the large area ASIC.

presented switched capacitor charge removal architecture the coupling between uncertainty in the timing diminishes as the duration of the charge removal operation is lengthened until it ultimately becomes negligible.

In addition to robustness against timing uncertainty, the Mixed–Mode PAD switched capacitor charge removal architecture offers the benefit that it automatically compensates for uncertainty in the comparator threshold. To see why, suppose that a noiseless comparator triggers a charge removal operation when the integrator output crosses the comparator threshold voltage ($V_{\mathrm{outp}} < V_{\mathrm{th}}$). Neglecting the signal accumulated during the charge removal, we have that $V_{\mathrm{outp}} = V_{\mathrm{th}} + \frac{C_{\mathrm{rem}}}{C_{\mathrm{int}}}(V_{\mathrm{ref}} - V_{\mathrm{low}})$ following the removal. If we then allow a fluctuation in the comparator threshold (δV_{th}) such that the charge removal triggers at $V_{\mathrm{outp}} = V_{\mathrm{th}} + \delta V_{\mathrm{th}}$, then following the removal $V_{\mathrm{outp}} = V_{\mathrm{th}} + \delta V_{\mathrm{th}} + \frac{C_{\mathrm{rem}}}{C_{\mathrm{int}}}(V_{\mathrm{ref}} - V_{\mathrm{low}})$. In this way, the effect of the threshold variation is retained in the analog residual voltage and may, therefore, be removed when the analog and digital data are recombined in post processing.

Finally, because of the switched capacitor architecture used to accomplish the charge removal, this circuit is capable of accurate photocurrent collection while the charge removal is in process. Consequently, this circuit operates, essentially, without dead time during an exposure.

Due to the impact of its accuracy on the ultimate accuracy of the detector, the charge removal circuit is one of the most important and subtly complicated elements of the Mixed–Mode PAD pixel. The architecture we have just finished presenting is very well matched to the Mixed–Mode PAD application because it has a very low susceptibility to noise and minimally impacts the operation of the rest of the front end circuity.

4.2.3 In-pixel Counter

The counter is one of the largest area commitments in the pixel, taking up nearly one quarter of the total transistor area in the AE207 (fall 2006) layout. Because of the substantial fraction of the pixel area required, a compact counter architecture is very important. In addition, it must be possible to read out and reset the counter rapidly to minimize the interframe deadtime. Different architectures were investigated for this circuit, ultimately settling upon the canonical pseudorandom counter solution. At the end of this section we will offer a brief discussion regarding this choice.

4.2.3.1 Pseudorandom Counter

One popular architecture that accomplishes these tasks utilizes a two tap Fibonacci mode Linear Feedback Shift Register (LFSR).[18] Practically speaking this is a shift register, as shown in figure 4.20, where two selected elements are tapped off, exclusively OR'ed, and then used to generate the first data element of the shift register for the subsequent clock pulse. When n and m (as defined in figure 4.20) are chosen appropriately, the sequence produced by clocking this circuit will have $2^m - 1$ elements with a predictable sequence of states.[19] This allows the number of counter clock cycles the system has undergone to be determined from knowledge of the final state, as long as the initial state is also known.

This circuit is an elegant application of the theory of Finite Fields and, as such, the mathematics underlying its operation are quite interesting and complicated. A detailed discussion of the mathematical underpinnings of this circuit is offered

[18] So called because the sequence of elements generated is due to a *linear recurrence*; that is, each element may be expressed as a linear combination elements that preceeded it in the sequence. As the famous Fibonacci Sequence $(0, 1, 1, 2, 3, 5, 8, 13, 21, \dots)$ is generated by a similar recurrence over the integers, this register structure also bears his name.

[19] While the sequence of states is predictable, the resulting distribution numerical values is sufficiently flat that this circuit is often used as a random number generator [47].

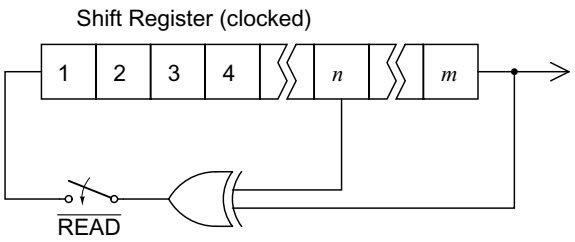

Figure 4.20: General architecture of the linear feedback shift register based pseudorandom counter. Figure adapted from [47].

in appendix A. For reasons discussed in this appendix, the 18-bit pseudo random counter used in the Mixed–Mode PAD (*i.e.* $m = 18$) is able to attain the $2^{18} - 1$ possible pseudorandom counter states if a feedback tap is placed at the 11-th bit (*i.e.* $n = 11$).

4.2.3.2 Linear Alternatives

As mentioned, this was not the only circuit investigated for the for the in-pixel counter. The pseudorandom counter offers the advantage that it is structurally very simple, being simply a linear shift register with minimal feedback, and, therefore, requires minimal space to implement. However, this structural simplicity and compactness comes at the price of data complexity and limited functionality. The data complexity is a consequence of the pseudorandom series of counter states the circuit enters into and the resulting, effectively random, pattern of bits making up the digital data stream. Because of this high level of complexity, debugging the digital data stream a painful exercise, not recommended for persons with a short temper. To illustrate what is meant by limited functionality of this circuit, consider replacing the pseudorandom counter with a linear counter (*i.e.* a structure where the bits correspond to significance in the numerical value of the counter). Then, if a least significant bit readout architecture is used, the time to read out the detector could be reduced in low flux/short exposure experiments by minimizing

the number of pixels needing to be addressed. Alternatively, one can imagine designing a counter with fewer bits than the current counter, but where the highest order bits are regularly read off the detector as the exposure is acquired, yielding an imager whose well depth is effectively unbounded.

Because of these disadvantages, an effort was made to develop a compact linear counter architecture to replace this canonical design, that culminated in the shift register element detailed in section 4.3.1.1. The unfortunate drawback of this circuit is that, because of its very compact layout, its radiation tolerance is not sufficient to act as a counter in a synchrotron application, as the long term effects of radiation damage steadily activate the elements nMOS components, slowing down transitions for the devices LOW to HIGH state and ultimately pinning every state node in the LOW state—until the damage is removed through annealing as will be discussed in section 6.7. The limited radiation tolerance of this circuit is, however, not a concern in its application in the pixel diagnostic circuit, which we will discuss in further detail in section 4.3.1.1, as LOW states in this register simply disable the diagnostic structures.

4.2.3.3 Counter Conclusions

For better radiation tolerance, the decision was ultimately made to adopt the canonical pseudorandom counter using a common register element.[20] Although this circuit does not use a radiation hardened layout, the larger spacing of its transistors combined with the fact that, unlike a linear ripple counter, it is externally clocked, and therefore does not couple state changes in one element with driving state changes in another, yields higher radiation tolerance, as evidenced by the radiation

[20]Each register element is a master/slave pair of D-latches operating with opposite logic. The D-latches are simply a cross coupled inverter pair with a CMOS switch in the feedback path and a inverting buffer as an input, that is active when the feedback switch is open and inactive when it is closed. Schematics of this circuit may be found in [7].

damage assessment that will be presented in section 6.7.

4.3 Periphery Pixel Components

In addition to the basic elements of the pixel already discussed, there are a number of peripheral components whose operation supports the pixel and allows it to interface with the external world. These include:

Pixel Diagnostic Circuit: An externally programmable circuit responsible for routing signals nodes from within the pixel onto a diagnostic bus for external monitoring as well as providing test inputs to the pixel.

Pixel Correlated Double Sampling: An analog noise reduction circuit, primarily intended to reduce the low frequency noise of the pixel.

Pixel Sample and Hold: The final analog output and buffering stage within the pixel.

The following sections will provide a discussion and analysis of these circuits.

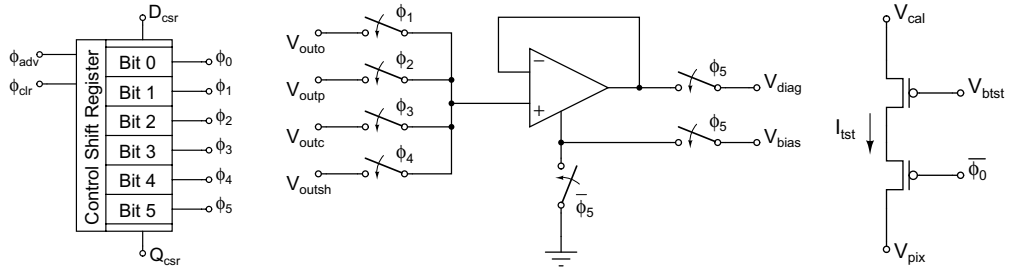

(a) CSR (b) Analog MUX and Output Buffer (c) I_{tst} Source

Figure 4.21: Elements of the pixel diagnostic circuit. Panel (a) shows the shift register used to control the diagnostic circuit (*i.e.* the Control Shift Register or CSR). Panel (b) shows the analog MUX and output buffer used to drive waveforms within the pixel to test point on the ASIC periphery. Panel (c) shows the test current source connected to the integration node of each pixel.

4.3.1 Pixel Diagnostic Circuit

In a complicated, large-area arrayed device, such as the Mixed-Mode PAD, it is very important to have diagnostic capabilities within each pixel. This allows characterization of how the pixel performance varies with location in the array. It also provides an invaluable tool for isolating array-level problems to specific portions of the design. For this reason, every pixel within the Mixed–Mode PAD contains the externally programmable diagnostic circuit shown in figure 4.21. The diagnostic structure may be broken down into three components: a control shift register used to configure the circuit's behavior; a MUX and buffer combination to transmit waveforms from critical internal nodes within the pixel to monitoring points, probe and wire bonding pads, on the ASIC periphery; and a test current source that allows a variable level of current to be injected onto the pixel's integration node.

4.3.1.1 Control Shift Register

Because of the space restrictions within the pixel and array, the space available for diagnostic elements is limited as is the number of output lines for diagnostic monitoring. To monitor multiple critical nodes within multiple pixels, it is necessary to have a programmable multiplexing system that allows individual signals of interest to be selected and driven to probe or bonding pads on the ASIC periphery. Control of this multiplexer, as well as other aspects of the diagnostic circuit as detailed in table 4.6, comes from a six-bit shift register included in every pixel, the Control Shift Register (CSR). These registers are chained together in a column-wise fashion to facilitate programming. From a functional standpoint this design seems unremarkable. However, at the heart of this register is a unique, small-area single-phase shift register element, modeled after a circuit developed at Caltech

[69].

Table 4.6: Summary of pixel diagnostic bits. Offsets are given in a big-endian format.

CSR Bit Name	Offset	Function
COUNT	0	Charge Removal Control Pulse Select
CALEN	1	Diagnostic Test Source Enable
OUTP	2	Integrator Output Select
OUTC	3	Comparator Output Select
OUTSH	4	Sample & Hold Output Select
AMPEN	5	Diagnostic Buffer Amplifier Enable

The fundamental register element, shown in figure 4.22, is made up of two complementary pairs of cross–coupled inverters. Each branch of this circuit is activated by a choke transistor (M7, M8, M11, and M12 in figure 4.22) that either supplies its branch with or starves it of current, depending on the state of ϕ_{adv}. In addition, there are two complementary sets of pass transistors (M5, M6, M9, and M10) that provide a gate between the supplied branches and the current starved branches, allowing the supplied branches to write their current state into the starved branches. Because complementary sets of transistors are used it is possible to drive this architecture using a single clock signal—whereby alternating ϕ_{adv} between low and high states sequentially supplies nMOS and pMOS cross coupled inverter stages, thus shifting data states through the register.

The risk this architecture presents is that a driven stage may overwrite the state of the driving stage during a ϕ_{adv} transition. To prevent this, the pass transistors must be made sufficiently weak relative to the choke transistors to ensure that write back is not possible. To find conditions for reliable operation we can consider the worst case scenario example of the transistor stack M10, M11, and M13, drawn from panel (a) figure 4.22, under the steady state assumption that:

$$V_{g,\text{M10}} = V_{g,\text{M11}} = V_{g,\text{M13}} = V_{d,\text{M10}} = V_{\text{dd}}. \tag{4.52}$$

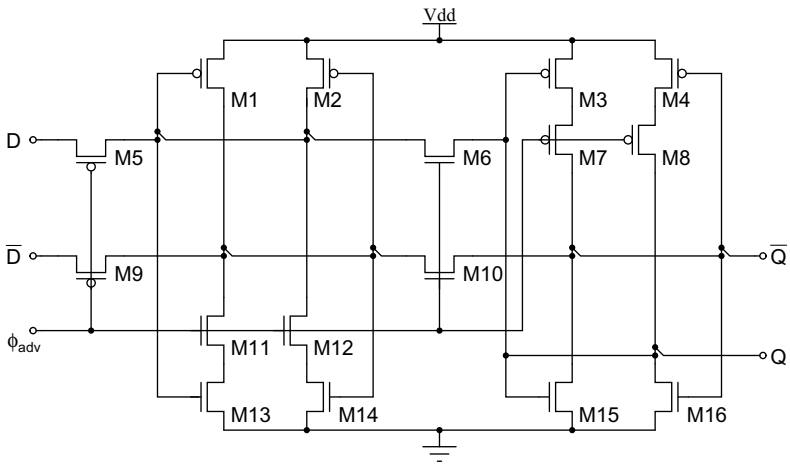

(a) Small-Area Single-Phase Register Element: Standard Schematic

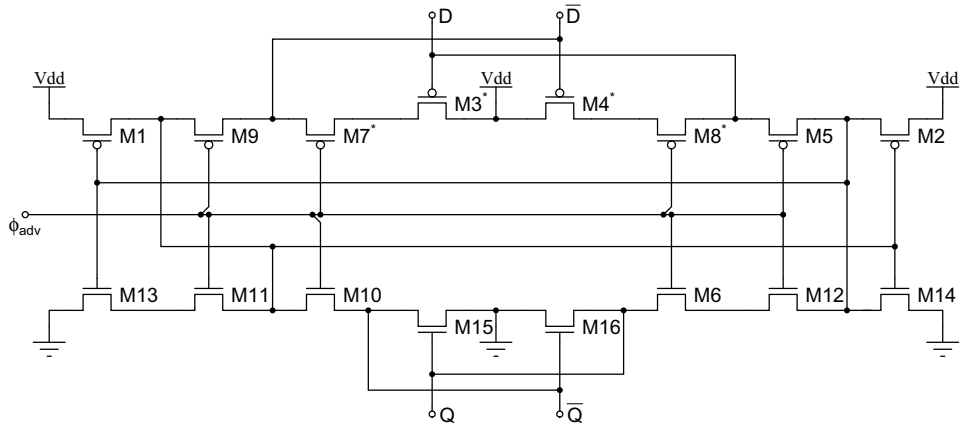

(b) Small-Area Single-Phase Register Element: Topological Schematic

Figure 4.22: Small-area single-phase shift register element.

Under these circumstances, assuming that the width of all transistors is the same, then the drain voltage of the choke transistor M11 ($V_{d,\mathrm{M11}}$) will be roughly:

$$V_{d,\mathrm{M11}} \approx \left(\frac{L_{\mathrm{M11}} + L_{\mathrm{M13}}}{L_{\mathrm{M10}} + L_{\mathrm{M11}} + L_{\mathrm{M13}}} \right) V_{\mathrm{dd}}. \tag{4.53}$$

Prima-facia, this result is independent of process. However, since this analysis is intended to represent the worst case input voltage of one inverter in the cross coupled pair, one has to consider the transition voltage of this inverter. Assuming we have sized our nMOS and pMOS widths to balance the inverters ($W_{\mathrm{PMOS}} \approx \frac{\mu_{\mathrm{e}}}{\mu_{\mathrm{h}}} W_{\mathrm{NMOS}}$) then $V_{d,\mathrm{M11}} = 0.20 \cdot V_{\mathrm{dd}}$ should, conservatively, yield reliable operation. This constraint implies that $L_{\mathrm{M10}} = 4(L_{\mathrm{M11}} + L_{\mathrm{M13}})$ or $L_{\mathrm{M10}} = 8L_{\mathrm{min}}$, where L_{min} is the minimal gate length allowed in the technology.

The complexity of this circuit's description belies its true elegance. Topologically, the circuit shown in panel (a) of figure 4.22 is equivalent to that shown in panel (b) of the same figure. In the latter figure, the transistor network has been unraveled to show how the large number of common source/drain nodes along with its complementary pair structure may be exploited to create a very compact layout, with an element of this register requiring roughly 50% of the area needed for the more traditional register architectures used elsewhere in the Mixed–Mode PAD design.

4.3.1.2 Analog MUX and Output Buffer

The MUX and output buffer used in the diagnostic circuit both utilize relatively standard architectures. The MUX is a series of CMOS pass gates that connect in common to the input of the output buffer. The output buffer, in turn, is a basic, five transistor, amplifier with a pMOS input stage, configured as a unity gain follower.

The simplicity of this architecture has consequences that need to be mentioned.

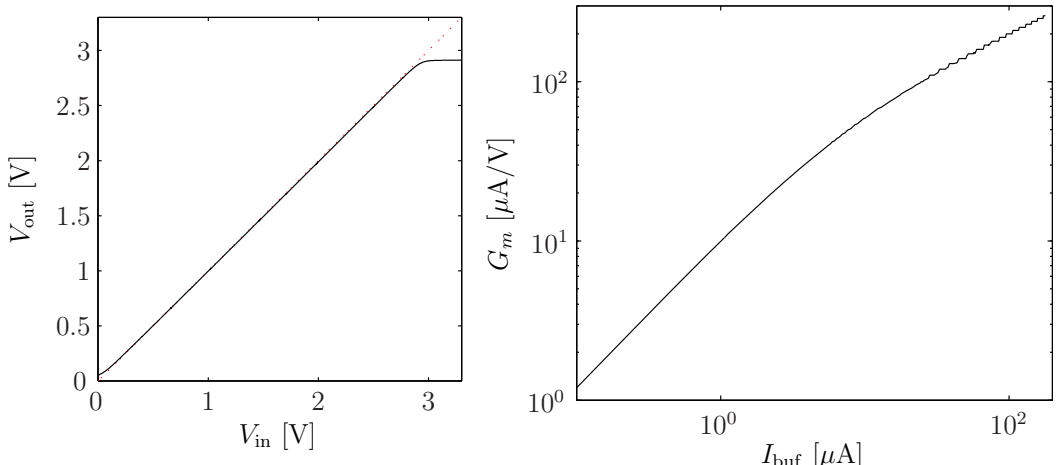

(a) Follower Range

(b) Transconductance

Figure 4.23: Performance characteristics of the diagnostic buffer amplifier.

With regard to the Analog MUX, when power is initially applied to the ASIC the CSR will enter into a random state, arbitrarily opening and closing the channels of the MUX and thereby shorting together active nodes within the pixel. If the pixel is powered at this point, then this random CSR configuration can result in an unacceptably high system power draw. To prevent this potentially damaging power-up situation, a clear signal (ϕ_{clr}) is included in the CSR which, when active, causes cycling ϕ_{adv} to clear all register elements, opening all the switches in the analog MUX. Consequently, the Mixed–Mode PAD requires a special start-up sequence wherein the CSR of all pixels is cleared, the DACs are programmed to set minimal bias currents and reference voltages,[21] and then the CSR is programmed before the operating bias currents and reference voltages are loaded into the DACs.

Finally, there are a few points that are worth mentioning regarding the limits of the output buffer in accurately reproducing waveforms from within the pixel. These limits affect the slew rate, bandwidth, and range of the output buffer. The

[21]Each ASIC contains a bank of DACs on its periphery that control bias currents and voltages. Both the reference voltages and the current mirror voltages used to set bias currents within the array are buffered to prevent loading issues. However, as a consequence, if these DACs are zeroed, the inputs of these buffers may float leading to unpredictable behavior in the ASIC.

slew rate limitation comes about from the substantially larger capacitance the output buffer has to drive (C_{load}), relative to capacitances within the pixel. This load is the result of the parasitic capacitance of the diagnostic bus, estimated to be on the order of 10 pF, in parallel with the load of the off chip probing tool. The setup typically used for work in this thesis was a model 12C Picoprobe from GHB Industries that presents a 0.1 pF capacitive load with 1 MΩ of shunt resistance. With this measurement configuration and the output buffer bias current (I_{buf}) at its nominal level of 10 μA, the slew rate is limited to 1 V/μs, which is at least an order of magnitude below slew rates on nodes within the pixel.

The effect of output loading also extends to the small-signal performance of the amplifier. In this architecture, the frequency response of this amplifier has a low-pass characteristic with a unity-gain bandwidth of $G_{\text{m}}/C_{\text{load}}$, where G_{m} is the amplifier transconductance which, as with the folded cascode architecture discussed in section 4.2.1, is determined by the transconductance of the input transistors. Panel (b) of figure 4.23 depicts G_m simulations for the output buffer as a function of the buffer's bias current. With the nominal I_{buf} setting of 10 μA, this analysis then predicts a unity gain bandwidth of $\sim$ 10 MHz, which will result in noticeable shaping effects on all monitored nodes except for the track and hold voltage of the pixel sample and hold circuit (V_{outsh}).

A final consideration is the output range over which the buffer is capable of following an input signal. As figure panel (a) of 4.23 depicts, the output range of the amplifier is not rail-to-rail. This occurs because the output branch of the circuit is also one of the differential input branches. As a result, it is possible for the output to rise high enough to drive the pMOS bias transistor out of saturation and into its ohmic region, thereby inhibiting the output's ability to track the input.

4.3.1.3 Test Current Source

The final element of this circuit is a current source that supplies a test current (I_{tst}) with which the basic functionality of the pixel may be checked.[22] The structure is a simple current mirror with one small, but significant, modification. Because it connects to the integration node of the circuit, it will introduce some leakage onto this node. Therefore, care must be taken to keep this leakage at a minimum. As was discussed in section 4.2.2, modern integrated circuit technologies often use ion-implantation to effectively lower device thresholds. This processing leaves the device in a partially on state with current levels that could be as high as 1 pA. To avoid this problem, a negative gate to source voltage (relative to the transistor type) is applied, to eliminate leakage from the channel, by setting the source voltage of the current mirror (V_{cal}) a few tenths of a volt below the off logic state of the source ($\bar{\phi}_0 = V_{\text{DDA}}$).

4.3.2 Mixed-Mode PAD CDS

Correlated Double Sampling (CDS) is a technique that has been used in CCDs and other precision measurement systems to remove low frequency noise through a time-correlated difference measurement. The Mixed-Mode PAD pixel was designed with an analog CDS system integrated into each pixel to allow correlated double sampling in parallel among pixels. While at first this method seems to yield the same behavior as the serial CDS used in CCDs and most other devices utilizing this technique, there are important differences that can have a degenerative impact on the performance of the detector. In fact, initial tests with the Mixed–Mode PAD yielded surprisingly poorer performance results when CDS was used in comparison

[22] The accuracy of this source, particularly its temporal stability, is generally not sufficient for quantitative evaluations of the pixel performance.

to tests where CDS was not active. The analysis of the correlated double sampling method presented in this section is intended as an explanation for this difference.

The section begins with a brief discussion of the Mixed-Mode PAD CDS implementation and proceeds to a general analysis and discussion of the CDS transfer function. Then, we compare CDS with the effect of performing only a single sample at the end of the integration. This analysis is followed by a discussion of non-ideal behavior; that is, behavior not included in the typical first-order analysis, that can affect the CDS measurement. Finally, conclusions are drawn, based on this analytical work, as to the limitations of analog CDS for this sort of measurement.

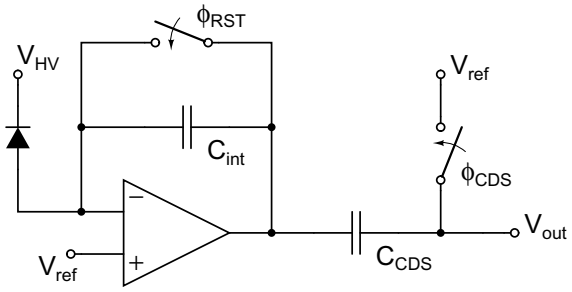

Figure 4.24: Schematic of the Mixed-Mode PAD CDS implementation.

Figure 4.24 shows a schematic of the Mixed-Mode PAD CDS circuit. The operation of this circuit is controlled by two clock signals: ϕ_{rst} which gates the pixel reset; and ϕ_{CDS} which clamps the tracking node of the CDS capacitor to a reference voltage. The circuit works in four stages: 1) both switches are closed to reset the pixel and CDS; 2) the pixel reset is opened, sampling the front end noise and injecting charge through clock feedthrough; 3) the CDS clamp switch is opened, allowing the tracking node to follow the pixel output; and 4) at the end of integration the voltage on the tracking node is sampled and recorded.

4.3.2.1 General CDS

From a signal processing standpoint, the principle underlying CDS is that of a high-pass filter, removing low-frequency noise through a cancellation operation. Before deriving this transfer function, it will be useful to review the particular noise sources that CDS attempts to remove by looking at what happens during the opening of the reset switch.

Initially, the noise at the integration node within the pixel can be approximated as a bandwidth limited white noise source of RMS intensity:

$$\delta V_{\text{pix}} = \sqrt{\frac{kT}{C}} = \left[4kTR \cdot \frac{1}{2\pi} \int_0^\infty \frac{d\omega}{1 + (\omega RC)^2} \right]^{\frac{1}{2}}, \tag{4.54}$$

where the familiar kT/C result has been expanded to emphasize the dependence on the effective resistance seen from the pixel integration node even though there is a cancellation effect of R in the source and bandwidth terms. When the pixel reset switch is closed, the effective resistance at the integration node is small, $R_{\text{int,on}} < 10^4$ ohms, whereas when the reset switch is opened, the effective resistance of this node increases by many orders of magnitude, $R_{\text{int,off}} > 10^{12}$ ohms. Because of this change, opening the reset switch effectively amplifies the per-unit-bandwidth noise at the integration node but limits the noise bandwidth to lower frequencies. Additionally, the act of opening the switch samples some of the closed-state noise onto the pixel, the exact quantity of which depends on the waveform used to control the reset switch, as well as injecting charge onto the integration node via clock feedthrough. This sampled noise and injected charge can be thought of as very low frequency contributions to the open-reset-state noise spectrum.

If these contributions are not removed and an unfiltered measurement of the pixel output is made, then the resulting noise contribution of the integration node,

both thermal and clock feedthrough, will be

$$\delta V_{\text{int}} = \sqrt{2\frac{kT}{C} + \frac{\delta Q_{\text{CLK}}}{C}}. \tag{4.55}$$

Seen from this perspective, it should be clear why a high-pass filter, like CDS, is desirable to suppress the low-frequency noise on the integrator front end when the reset switch is in the open state.

4.3.2.2 CDS Transfer Function

We begin with a simplified description of the CDS transfer function,

$$V_{\text{out}}(t) = V_{\text{in}}(t)\delta(t) - V_{\text{in}}(t)\delta(t - \Delta t_s), \tag{4.56}$$

where $V_{\text{out}}(t)$ is the voltage from the pixel and Δt_s is the time between sampling of the reset noise and sampling of the signal. The Fourier transform of this transfer function is

$$\hat{V}_{\text{out}}(\omega) = \hat{V}_{\text{in}}(\omega)(1 - e^{-i\omega\cdot\Delta t_s}), \tag{4.57}$$

whose norm is

$$\left|\hat{V}_{\text{out}}(\omega)\right|^2 = \left|\hat{V}_{\text{in}}(\omega)\right|^2 (2 - 2\cos(\omega \cdot \Delta t_s)). \tag{4.58}$$

For a given input noise spectra ($\hat{N}_{\text{in}}(\omega)$), the spectrum after the CDS ($\hat{N}_{\text{out}}(\omega)$) then becomes

$$\left|\hat{N}_{\text{out}}(\omega)\right|^2 = \left|\hat{N}_{\text{in}}(\omega)\right|^2 (2 - 2\cos(\omega \cdot \Delta t_s)), \tag{4.59}$$

for a total output noise (N_{CDS}, rms volts) of

$$\begin{aligned} N_{\text{CDS}} &= \left[\frac{1}{2\pi}\int_0^\infty \left|\hat{N}_{\text{out}}(\omega)\right|^2 d\omega\right]^{\frac{1}{2}} \\ &= \left[\frac{1}{2\pi}\int_0^\infty \left|\hat{N}_{\text{in}}(\omega)\right|^2 (2 - 2\cos(\omega \cdot \Delta t_s))\right]^{\frac{1}{2}}. \end{aligned} \tag{4.60}$$

4.3.2.3 Effect of CDS on Low-Pass-Filtered White Noise Source

To illustrate the effect of CDS, let us consider a low-pass-filtered white noise of spectral density A_n and filter time τ_n, so that,

$$\left|\hat{N}_{\text{in}}(\omega)\right|^2 = \frac{A_n}{1+(\tau_n\omega)^2}. \tag{4.61}$$

After CDS this spectrum becomes

$$\left|\hat{N}_{\text{out}}(\omega)\right|^2 = \frac{A_n}{1+(\tau_n\omega)^2}\left(2-2\cos(\omega\cdot\Delta t_s)\right). \tag{4.62}$$

Figure 4.25 shows this result for four different combinations of Δt_s and τ_n, illustrating the dramatic changes in the post CDS spectra that relative changes in these two time constants produce.

To further extend this result, we can integrate the post CDS spectrum to find the total output noise,

$$N_{\text{CDS}} = \left[\frac{1}{2\pi}\int_0^\infty \frac{A_n}{1+(\tau_n\omega)^2}\left(2-2\cos(\omega\cdot\Delta t_s)\right)\,d\omega\right]^{\frac{1}{2}}. \tag{4.63}$$

This integral may be solved analytically using residue calculus [5]. After a little work one finds

$$N_{\text{CDS}} = \left[2A_n\cdot\left(1-e^{-\frac{\Delta t_s}{\tau_n}}\right)\right]^{\frac{1}{2}}, \tag{4.64}$$

the result of which is plotted in figure 4.26.

An important corollary that we can draw from equation 4.64 is that the break-even point for low-pass-filtered white noise occurs when $\Delta t_s = -\tau_n\cdot\ln\left(\frac{1}{2}\right)$. Therefore, as long as $\Delta t_s < -\tau_n\cdot\ln\left(\frac{1}{2}\right)$, CDS will reduce the total output noise; however, when $\Delta t_s > -\tau_n\cdot\ln\left(\frac{1}{2}\right)$, CDS will actually amplify it.

4.3.2.4 Noise Comparison without CDS

To provide a metric with which one may gauge the effectiveness of CDS, it is useful to consider the noise of a measurement without CDS. Here, there are two

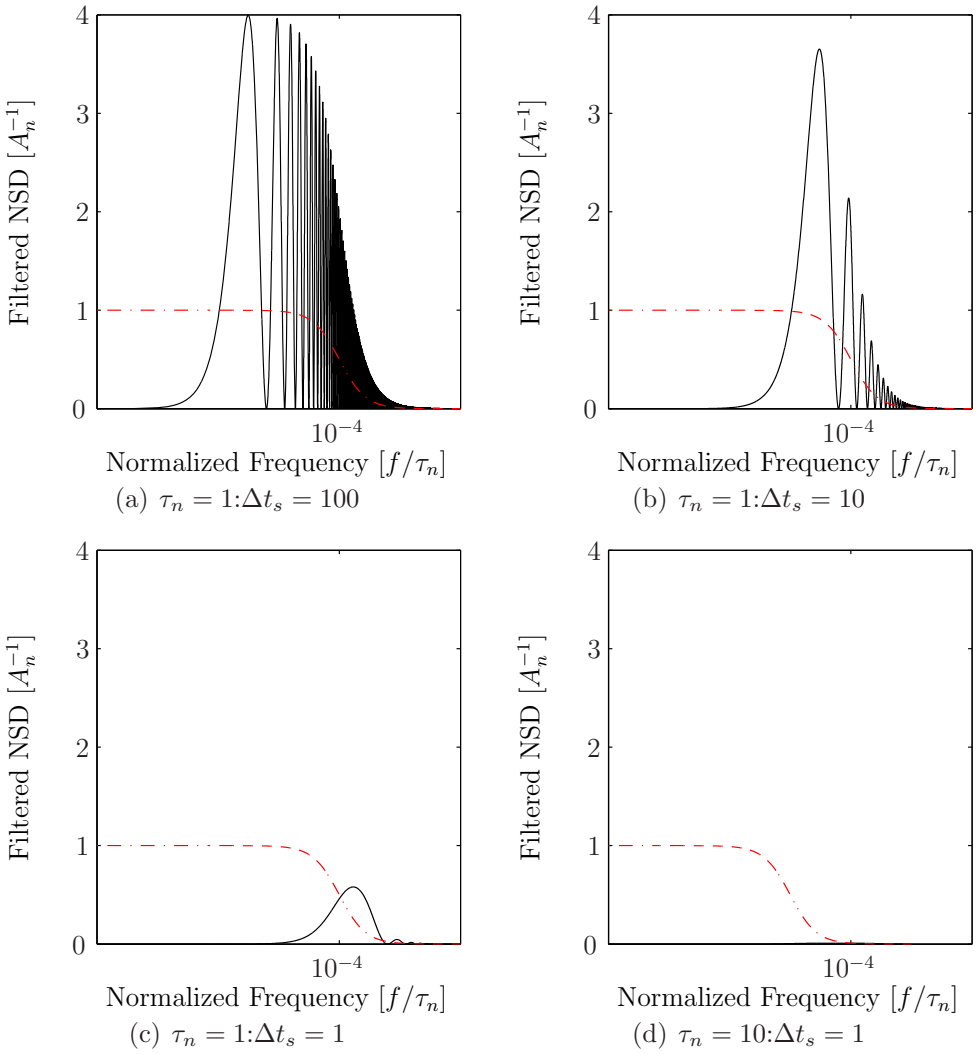

(a) $\tau_n = 1$:$\Delta t_s = 100$ (b) $\tau_n = 1$:$\Delta t_s = 10$

(c) $\tau_n = 1$:$\Delta t_s = 1$ (d) $\tau_n = 10$:$\Delta t_s = 1$

Figure 4.25: Post CDS-filtering of low-pass-filtered white noise spectra for different combinations of τ_n and Δt_s. These figures illustrate how strongly the effectiveness of CDS is influenced by the ratio of the these two time constants.

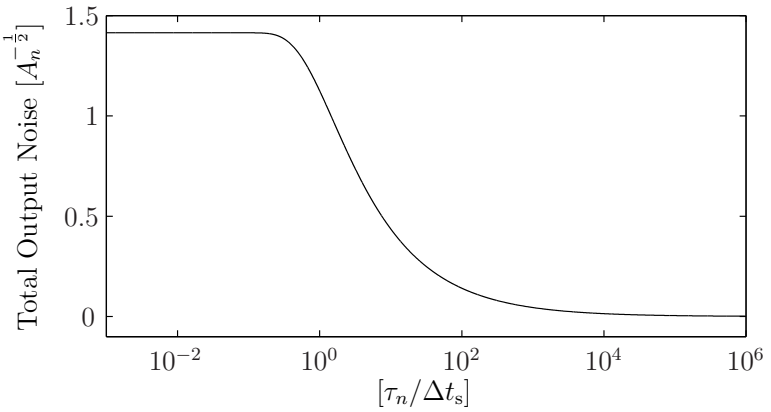

Figure 4.26: Normalized total output noise (N_{CDS}) of a low pass filtered white noise spectrum with noise power A_n and filter constant τ_n after CDS with sampling time Δt_s.

dominant noise sources to look at: the noise contributed by thermal fluctuations on the integration node; and noise contributed by the active devices within the integrator amplifier.

The example of a low-pass-filtered white noise spectrum illustrates how CDS can, in the presence of some noise spectra, increase the sampled noise relative to a single sampling. However, a single sample does not accurately represent what happens in a pixel without CDS. When CDS is not used, the noise at the pixel front end is still sampled twice, once when the pixel is reset and once when the signal is read, with a white noise contribution to these measurements of $\sqrt{\frac{kT}{C_{\mathrm{pix}}}}$ in both cases.[23] This fundamental noise sets the minimal single sample total output noise, neglecting clock feed through, at

$$N_{\mathrm{sing}} \geq \sqrt{2 \cdot \frac{kT}{C_{\mathrm{pix}}}}, \tag{4.65}$$

[23]In a strict sense, this statement is not true for the sample taken at the end of the integration. At this time, the integrator effectively has a capacitive feedback network formed by C_{int} and C_{pix}. The gain of this network, relative to fluctuations of V_{pix} will be $\frac{C_{\mathrm{pix}}+(1+A(\omega))C_{\mathrm{int}}}{(1+A(\omega))C_{\mathrm{int}}}$, where $A(\omega)$ is the frequency dependent gain of the integrator amplifier, which is expected to vary between 1 and ~ 5 based on the Mixed–Mode PAD design specifications for C_{int} and estimates of C_{pix}. However, the bandwidth of thermal fluctuations on the integration node when the reset switch is open should be low enough that the approximation $\frac{C_{\mathrm{pix}}+(1+A(\omega))C_{\mathrm{int}}}{(1+A(\omega))C_{\mathrm{int}}} \approx 1$ is valid.

which is the uppermost limiting value of the CDS filtered total output noise as $\tau_n/\Delta t_s \rightarrow 0$. Therefore, although CDS may amplify the noise it does not make it worse than a measurement without CDS.

The noise from the integrator amplifier will be an additive contribution to the integration node thermal noise, both with and without CDS, and, therefore, may be treated independently. In the case of a single sampling measurement, the front-end electronic noise is sampled twice, once when the reset switch is opened and again when the signal is sampled. With the first sampling, though, a charge of $(C_{\rm pix}+C_{\rm int})\delta V_{\rm outp}$, where $\delta V_{\rm outp}$ is the sampled fluctuation in the integrator output caused by the active circuitry in the amplifier, is sampled onto the integration node leading to a fluctuation in the integrator's output voltage of $-\frac{C_{\rm pix}+C_{\rm int}}{C_{\rm int}}\delta V_{\rm outp}$. The noise spectrum will typically not change significantly with the opening of the switch and, like CDS, the time between samples will destructively correlate the results.

Assuming the amplifier has a single-pole low-pass characteristic and following a similar analysis to the preceeding CDS work, we find that

$$\left|\hat{V}_{\rm ioa,\overline{CDS}}(\omega)\right|^2 = \left|\hat{V}_{\rm ioa}(\omega)\right|^2 \left(1+\frac{C_{\rm pix}+C_{\rm int}}{C_{\rm int}}-2\cos(\omega\cdot\Delta t_{\rm int})\right), \tag{4.66}$$

where $\hat{V}_{\rm ioa}(\omega)$ is the noise spectrum of the integrator amplifier and $\hat{V}_{\rm ioa,\overline{CDS}}(\omega)$ is the spectrum of the integrator amplifier's contribution to a measurement without CDS. If, as with the CDS derivation, we assume a low-pass-filtered white noise spectrum for the front-end amplifier of low frequency amplitude A_n and shaping time constant τ_n then the total noise contribution is

$$\begin{aligned} N_{\rm OA,sing} &= \left[\frac{1}{2\pi}\int_0^\infty \frac{A_n}{1+(\tau_n\omega)^2}\left(1+\frac{C_{\rm pix}+C_{\rm int}}{C_{\rm int}}-2\cos(\omega\cdot\Delta t_{\rm int})\right)\,d\omega\right]^{\frac{1}{2}} \\ &= \left[\frac{C_{\rm pix}}{4C_{\rm int}}A_n + 2A_n\cdot(1-e^{-\frac{\Delta t_{\rm int}}{\tau_n}})\right]^{\frac{1}{2}}. \end{aligned} \tag{4.67}$$

CDS sampling of the amplifier's output noise will yield the same result as equation 4.64, which is smaller than the above result by the term $\frac{C_{\rm pix}}{4C_{\rm int}}A_n$ in the square root.

Thus, the electronics noise from the amplifier should be increased in a measurement without CDS relative to one using CDS.

For both of the noise sources, this analysis argues that CDS should improve the noise performance over measurements without CDS. It also shows us that this advantage is marginalized when $\tau_n \ll \Delta t_{\text{int}}$. In this limit, the noise of the CDS response will approach that of a measurement without CDS, for the low-pass-filtered white noise spectra considered.

4.3.2.5 Analog CDS Fidelity

One item that has not been considered in the preceeding analysis is the long term fidelity of the CDS storage capacitor. In CMOS electronics, the transistors and integrated passive components tend to deviate from their idealized models; consequently, an effect like leakage onto or from the CDS capacitor is a concern. While there are a number of potential sources of capacitor leakage, the fact that the mean leakage may be treated as analogous to dark current means that only sources prone to significant variation need to be considered. For the CDS architecture used in the Mixed–Mode PAD, the most variable leakage current sources are photo and thermal electrons generated within the ASIC bulk.

Here, the leakage current results from minority carriers generated in the bulk silicon, which diffuse into the reverse biased diode region surrounding the transistor source/drain diffusions. The diffusion length for electrons in p-type silicon is typically hundreds of microns to a few centimeters for doping concentrations from 10^{15} to 10^{18} acceptors per cm^{-3} [49]. This level of diffusion is sufficient that a substantial fraction of these minority carriers will be drawn into the transistor source/drain border diodes with the assistance of the weak field generated between the substrate grounding connection and the diode edge.

In the case of thermal generation, there is a strong dependence between the

density of minority carriers and temperature,

$$p \approx \frac{n_i^2}{N_{\mathrm{D}}}; \qquad n_i^2 \propto (kT)^3 e^{-\frac{E_g}{kT}}, \tag{4.68}$$

$$\Rightarrow \qquad p \propto (kT)^3 e^{-\frac{E_g}{kT}}. \tag{4.69}$$

For this reason, the temperature stability of the ASIC is critical to the hold ability of the CDS.

While thermal effects can be significant if there is not sufficient temperature control, photo current generated in the ASIC bulk can have an even greater effect. To understand this claim, consider the process. Photons are generated by sources, such as room lighting, which have both statistical and systematic (*e.g.* 60 Hz modulation) fluctuations.[24] Additionally, the probability of conversion in the silicon obeys an exponential decay governed by the mean free path (λ_ω) of the photon. The depth of conversion influences the fractional charge yield, as conversion deeper into the silicon will increase blooming (*i.e.* lateral diffusion of the resulting charge cloud). The variation in the source and absorption processes combine multiplicatively to create variablity in the CDS leakage current.

4.3.2.6 Conclusions on the Mixed–Mode PAD Analog CDS

A key difference between the serial implementation of CDS in a typical CCD and the parallel implementation in the Mixed–Mode PAD is that the CCD has much better control over the sampling time and the time constant of the dominant low-pass filter. The extent of the control is described in detail in [51] wherein it is explained how CCD designers are able to tune their CDS sampling time and low-pass-filter time constants so as to maximize the signal to noise ratio. In a parallel CDS device like the Mixed–Mode PAD, there is little control over the CDS

[24] Here, we are mainly considering the effect of ambient lighting; however, the claim holds for source lighting, optical or x-ray, as well.

sampling time because it is fundamentally set by the integration time. In addition, the sampling time is necessarily much longer than the bandpass of the front-end electronics, to permit signal acquisition. Therefore, the benefits of Analog CDS in the Mixed–Mode PAD are primarily restricted to low-frequency noise reduction.

This conclusion, however, presumes that one has an ideal CDS system. As discussed earlier in this section, there are environmental factors that diminish the fidelity of the CDS circuit. These factors can be reduced by operating the chip at colder temperatures, improving the thermal stability, eliminating excess light on the detector, and using wells to isolate transistors connected to charge sensitive node from current generation in the ASIC bulk. Ultimately, though, they impose a limit on the duration of exposures beyond which using CDS will increase the noise of the detector. For longer integrations, a different approach to correlated double sampling is appropriate; this method, called digital CDS, will be discussed in section 4.3.3.

4.3.3 Pixel Sample and Hold

The pixel Sample & Hold (S&H) is the final analog stage within the pixel. It records the state of integrator and retains this state until read out and reset. This action is necessary in order to define an exposure window, because there is no means to decouple the integrator from the detector diode. While sampling the state of the integrator, the sample and hold acts as a bandwidth limiting element, reducing the noise introduced by the integrator amplifier. Finally, this circuit acts as a buffer to drive the analog residual voltage to high power buffers at the edge of the chip.

The circuit that accomplishes these tasks is shown in figure 4.27. It comprises two unity-gain follower stages, one to drive the integrator's output voltage onto the

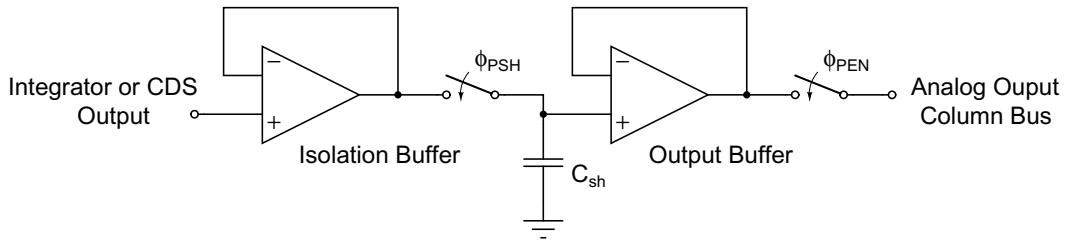

Figure 4.27: Schematic description of the pixel sample and hold circuit.

sample and hold capacitor ($C_{\rm sh}$) while isolating the integrator from this capacitance and the other used to drive the sample and hold voltage to buffers on the edge of the chip. The first stage also acts as the bandwidth limiting element of the circuit, through a combination of capacitive loading a transconductance control. This effect is identical to the situation discussed in section 4.3.1, wherein the unity gain follower's bandwidth is the ratio of the buffer's transconductance to the capacitance of the sample and hold capacitor ($G_{\rm m}/C_{\rm sh}$). The sample and hold capacitor is designed for 1 pF of capacitance and the transconductance of the sampling buffer is dependent on the buffer's bias current ($I_{\rm cds}$), as depicted in figure 4.28. With this follower's nominal bias current setting of 2 μA the sampling buffer should have a transconductance of $\sim$6 μA/V yielding a bandwidth of 6 MHz.

The second portion of the sample and hold is the output buffer responsible for driving the sampled residual voltage to the edge of the ASIC. This amplifier and its compatriot buffer at the chip edge determine the maximum frequency at which the ASIC's analog data may be read out. The architecture of the output buffer is a simple, five-transistor, amplifier with a pMOS input stage, identical to that of the diagnostic output buffer, discussed in section 4.3.1, whose characteristics are depicted in figure 4.23. Because the loading conditions for this circuit are nearly identical to that of the diagnostic output buffer, the anticipated bandwidth of the sample and hold output buffer is also $\sim$ 10 MHz at the nominal bias current of

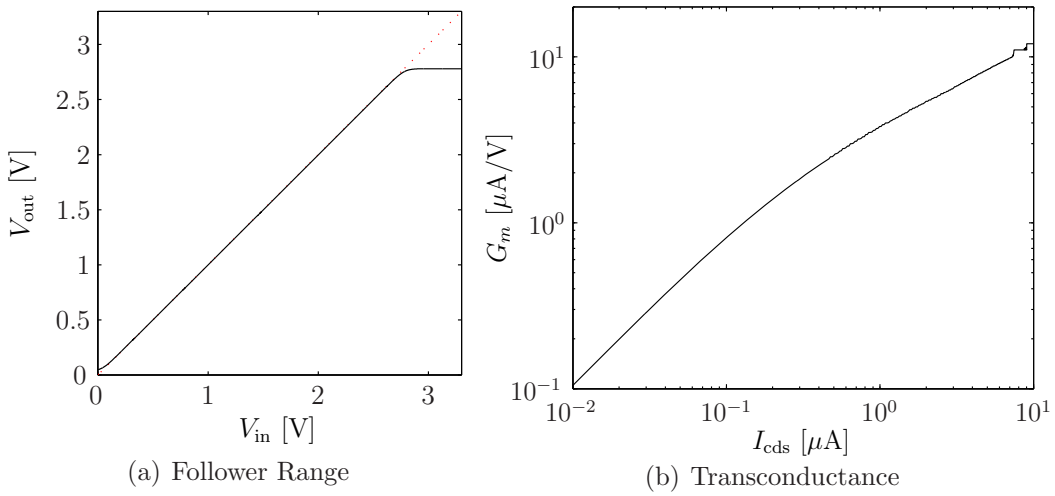

(a) Follower Range (b) Transconductance

Figure 4.28: Performance characteristics of the sample and hold isolation buffer.

10 μA.

4.3.3.1 Digital CDS

An interesting use of a sample and hold circuit is as a replacement for the analog CDS for long (> 0.1 s) low-to-moderate intensity integrations. This operation is accomplished by using the CDS to record the integrator output immediately after the reset switch is opened. The recorded voltage is then read out while the exposure is being acquired. At the end of the exposure, the detector is read out as usual; however, offline, the initial analog reading is subtracted from the final reading.

Mathematically, this operation is the same as the analog CDS described in section 4.3.2 with two important differences. First, the measurement is free of the degenerative effects associated with using CDS over long integration times. However, as the full analog readout chain must be used to make each measurement, the noise spectral density before the CDS is expected to be larger. At the time of this writing, digital CDS has not yet been implemented in the Mixed–Mode PAD.

However the method has been well studied and reported in [50].

4.4 Design Reflections

The Mixed–Mode PAD is a device with demanding performance specifications designed for the harsh radiation environment of the modern synchrotron light source. The flux available from modern synchrotrons enables experiments where signal intensities of interest can vary from billions of x-rays per mm^2 per second to a fraction of an x-ray per mm^2 per second. For reasons discussed at length in chapters 1 and 3, no x-ray imager currently in use or in development is capable of measuring this range of signal intensities, apart from the Mixed–Mode PAD. This range is made possible by combining the the high flux tolerance of an Analog PAD front end and the geometric well depth to circuit area relationship found in Digital PADs. In addition, by performing the most significant portion of the analog-to-digital conversion in-pixel and as the exposure is taken, the Mixed–Mode PAD is capable of breaking the interdependence of well depth, precision, and frame rate that normally limits analog imagers, to obtain a broad dynamic range while operating at frame rates beyond the capabilities of more conventional x-ray imagers.

Thus, within the limits of the project's original goals, the pixel design fabricated in the final 128×128 pixel hybrid imager rises to the design challenges this project has presented. Hindsight, however, offers a remarkable perspective for evaluating a design, allowing us to see that there are ways that these goals could be extended and a more capable detector developed in the future.

Generally, when one thinks of an imager the model that comes to mind[25] is

[25]At least to those of the author's generation and others who preceeded the digital-camera revolution.

the film camera, where light sensitive material is exposed for a period of time during which the illumination pattern is passively recorded. Today, nearly all x-ray imagers still follow the film-camera paradigm of passively recording an exposure then reporting the results after the exposure has completed. However, given modern levels of circuit integration and the availability of accessible and inexpensive reconfigurable logic devices (*i.e.* a Field-Programmable Gate Arrays (FPGAs)), it is possible to go beyond this concept. The step the Mixed–Mode PAD made from passive pixels to active pixels (*i.e.* introducing pixels which respond during an exposure to the signal they detect) was a significant advance in this direction, resulting in a dramatic increase in well depth and frame rate without significant sacrifices in flux tolerance. Yet, the Mixed–Mode PAD was still conceived as a device that would collect all its signal during an exposure in the pixel during and report the results only afterwards.

What custom signal processing and accessible reconfigurable logic offer is the opportunity to extend the concept of an active pixel to that of an active imager. By an active imager we mean a device where an active control system operates during the exposure to extend and improve the capabilities of the imager. To better explain this, consider the following two simple examples.

For the first example, consider a device like the Mixed–Mode PAD, but one where the 18-bit pseudorandom counter is replaced by a 10-bit linear counter and a latching 11^{th} bit that transitions high on the change of state of the 10th counter bit. To turn this into an active imager suppose that the 11^{th} bit of each pixel were read out and cleared, once every millisecond. This architecture would meet or exceed all the design specifications of the current Mixed–Mode PAD while offering new advantages. Notable among these, three advantages of this should be evident. First, by reducing the number of bits in the in-pixel counter, there is a savings

in terms of area within the pixel that could be used, for example, to implement a more radiation hard layout. Second, the space devoted to accumulating the 11^{th} bit overflows in the FPGA can be extremely large so that the well depths attainable with this design would far exceed anything that could possibly be implemented within a pixel alone. Finally, it is possible to process the 11^{th} bit overflow data within the FPGA, as the exposure is being taken. This, in turn, would make it possible to impose other end conditions on the exposure, for example requiring acquisition of a minimum signal in some part of the image.

As a second example, consider the problem faced by the Mixed–Mode PAD, and any other integrating device, for long exposures where a portion of the interesting x-ray signal is very weak. Because of the combined effects of accumulating dark current and low-frequency noise, it is very difficult to achieve single x-ray sensitivity in this case. Now suppose that, instead of only measuring the analog voltage at the end of the integration the sample and hold circuit was used to track it at regular intervals during the integration, akin to its operation in the digital CDS technique discussed in section 4.3.3.1. So long as the sampling is more frequent than the arrival of x-rays, digital signal processing within the FPGA should make it possible effectively count the x-rays as they arrive. This could be implemented in the current Mixed–Mode PAD hybrid, although it was not part of the original design intention.

Despite the room for future development, the Mixed–Mode PAD is an imager capable of performing experiments presently unaccessible to any other x-ray imager. In the remaining chapters we will demonstrate these capabilities, first presenting characterization measurements of the imager performance, then concluding with results from the first experiments with the Mixed–Mode PAD.

CHAPTER 5

SINGLE HYBRID CAMERA

As a final step before presenting measurements taken by the Mixed–Mode PAD, we present a discussion of the system built to exercise it. Just as it is said, "no man is an island,"[1] no hybrid is an imager. For all its complexity, the Mixed–Mode PAD hybrid is merely the part, although a very essential one, of a larger system whose purpose is to quantify patterns of x-rays, translating them into meaningful data values in the form of digital images. This chapter addresses the topic of the imager, beyond the detector hybrid, in two parts. The initial portion offers a presentation of the systems that support the detector hybrid and allow users to interact with it. The later portion looks in more detail at how these systems interact with the detector hybrid by focusing our discussion on the control signals generated by this system and the response of the hybrid.

5.1 System Breakdown and Decomposition

Our discussion up to this point has focused on the fixed portion on the Mixed–Mode PAD, the detector hybrid. While a substantial amount of time and effort went into designing this device, in the end it is still only a component of the imager. A complete set of support electronics, appropriate housing, as well as control, data acquisition, and data analysis software are also needed for a functional imager. Beyond this, it is the quality of these support systems that ultimately determines the performance potential the imager is capable of achieving.

Ultimately, the goal of the Mixed–Mode PAD project is to produce a large-area detector, 2048 pixels × 2048 pixels for roughly 310 cm × 310 cm active area constructed from 64 single detector hybrids. The task of building a support system

[1] John Donne (1572–1631), *Devotions Upon Emergent Occasions, Meditation XVII.*

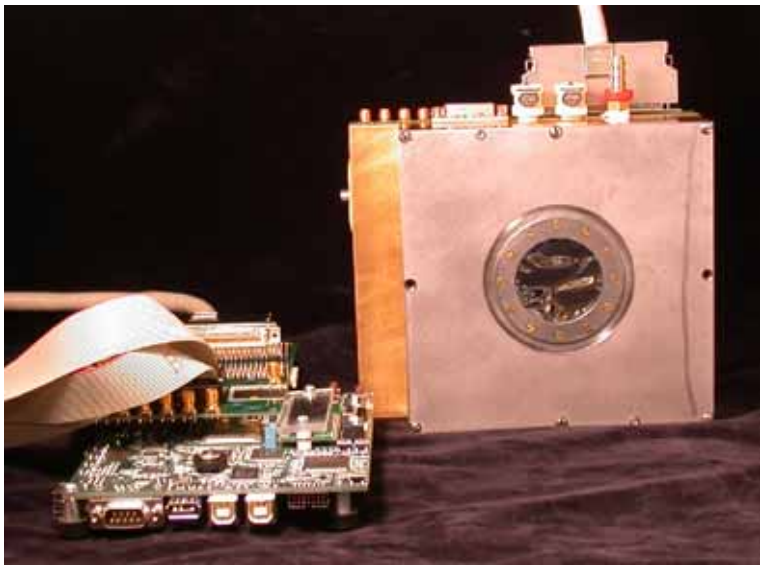

Figure 5.1: Photograph of the cryostat housing the Mixed–Mode PAD single hybrid camera along with the FPGA control and frame buffer used in the camera. Not shown is the electronics rack containing the data acquisition control computer.

for this detector extends far beyond the scope of this thesis and is primarily the responsibility of our commercial collaborators at Area Detector Systems Corporation (ADSC). In contrast, the support system, here after referred to as the camera, built at Cornell and used to generate the bulk of the material presented in this thesis was designed for a single detector hybrid. The purpose of its construction was explicitly to validate and characterize the performance of the these imagers; although the experiments presented in chapter 7 prove it is capable of doing much more.

To simplify our discussion of the single hybrid camera, we distinguish three sub-systems:

- Camera Housing and Detector Cryostat
- Low-Noise Support Electronics
- Data Acquisition and Control

The remainder of this section will address each of these sub-systems in turn.

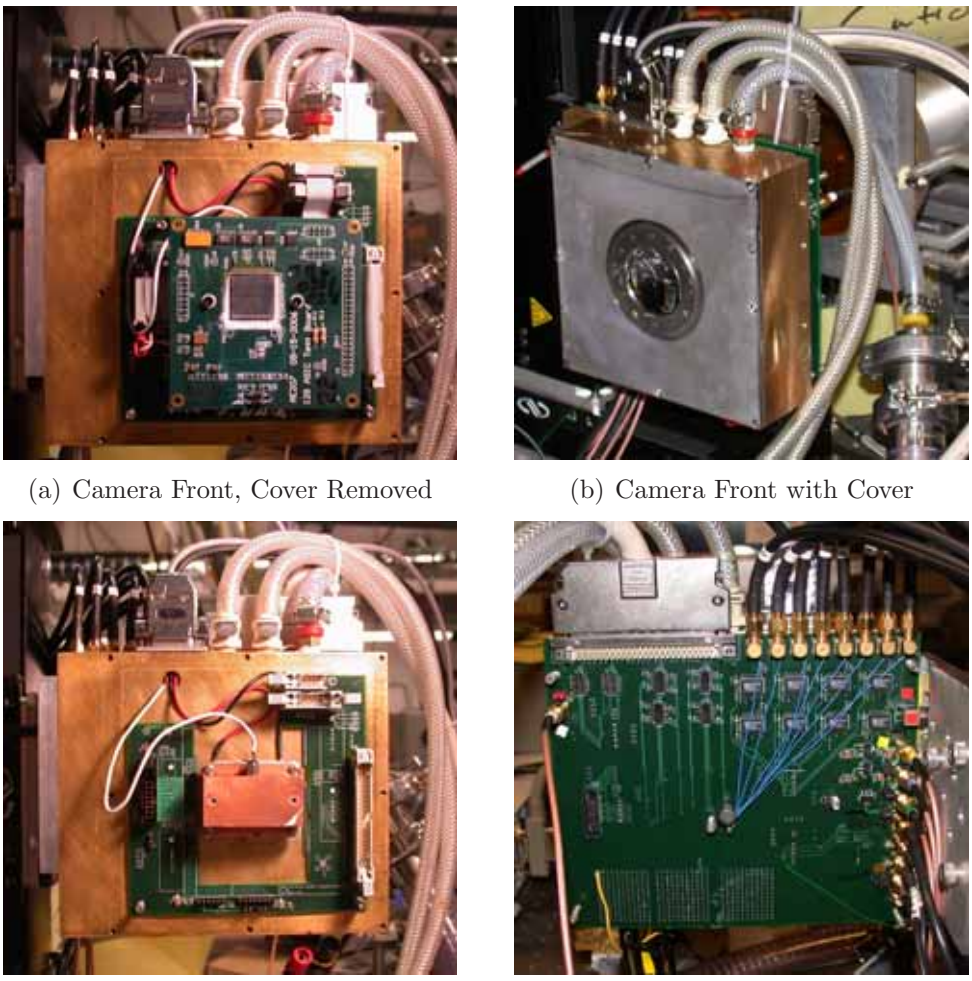

(a) Camera Front, Cover Removed

(b) Camera Front with Cover

(c) Camera Front, Hybrid Removed

(d) Camera Back

Figure 5.2: Photographs of the Mixed–Mode PAD from different perspectives. The plastic tubing snaking from the top of the back plate carry chilled water, left and center tube in panel (a); and supply the vaccuum connection, right tube in panel (a). Panel (c) exposes the thermelectric-cooled cold finger.

5.1.1 Camera Housing and Detector Cryostat

The detector hybrid is housed in a custom vacuum cryostat, as shown in figure 5.2. Built of a brass housing with an aluminium face plate and a 50 μm thick aluminized mylar x-ray window, it maintains a roughing pump vacuum of better

than 20 mTorr. In this environment, the detector is typically operated at $-35°$ C; achieved by a two-stage thermoelectric that pulls heat from a copper brick (cold finger), that is thermally connected to the detector hybrid, into the back plate of the camera housing, which, in turn, is water cooled. A small surface heating element may be attached to the front of the detector to prevent condensation on the x-ray window, when needed. With the exception of the thermoelectric control and monitoring lines, all electrical signals are passed through the camera back plate, via custom high-speed feed throughs,[2] directly to a supporting printed circuit board attached to the back of the detector.

5.1.2 High-Speed, Low-Noise Support Electronics

One of the design goals of the Mixed–Mode PAD is a total dead time of a millisecond or less. To comprehend this task, consider the quantity of data produced by the Mixed–Mode PAD within a single frame. A single detector hybrid has 16,384 pixels, each with 18 bits of digital data and 1 analog value that needs to be digitized to ± 1 mV on a 1 V range. Each hybrid is divided into eight 128×16 pixel banks, each bank with one digital and one analog output, so that, at the bank level, 36,864 bits and 2,048 analog values must be read out within this millisecond.

A primary challenge faced in the design of the Mixed–Mode PAD support electronics is the dilemma of how to maintain the fidelity of the analog residual voltage in the presence of the high-speed signals used during detector readout. For reasons that will be explained later in section 5.2.4, the analog readout clock and the digital readout clock are interleaved so that both data streams are drawn off the hybrid simultaneously. This presents a potential for crosstalk between the

[2]These feed throughs were built from back-to-back Mictor (Matched Impedance Connector from Tyco Electronics/AMP) interconnected by PCB epoxied into the back of the camera housing.

digital and analog data if care is not taken to protect the analog voltages from the high-speed (specified to be up to 100 MHz) transitions on the digital clock and data lines as well as, more significantly, the resultant transients on the power supply lines. To accomplish this a combination of minimization, isolation, and rejection techniques are employed.

To minimize the size of power supply transients, Low-Voltage Differential Signaling (LVDS) [53] is employed on all digital lines to and from the camera. While these circuits require a slightly higher quiescent current than modern single ended logic, as each line must constantly drive a fixed impedance, this drawback is made up for by reduced noise during switching, as current consumed by these drivers changes in direction rather than magnitude.

Isolation of the analog electronics is accomplished in two ways. First, the digital control and data lines, as well as the power planes, are kept physically separated from the analog circuits to minimize parasitic coupling. Second, as an electrical connection between the analog and digital electronics must exist to establish a relative potential and this potential must, in turn, be referenced to an absolute ground, particular care is taken with the grounding network. The absolute ground of the design is connected to the digital ground plane (V_{GNDD}). This is done because the digital circuit is far more robust against pick-up from this ground line. It also requires a lower impedance ground connection to moderate the ground bounce caused by its own switching transients. The analog electronics, on the other hand, are designed to draw a constant level of current, a result of the reliance on differential elements, and, thus, transients are much smaller. More important to the analog electronics than absolute potentials is that the potential difference between the high and low power supply lines is maintained at a constant level. Therefore, the analog ground plane (V_{GNDA}) is connected to the digital ground

plane, within the camera, by a large (10 mH) inductor. This low-pass connection maintains the two supplies near enough to prevent accidental forward biasing of the transistor diffusions on the hybrid while effectively isolating the analog supply from high-frequency noise on the digital supply.

Finally, to make the analog electronics robust against transients remaining after the aforementioned minimization and isolation steps, all measurements are preformed differentially with a reference that should exhibit similar noise coupling. Jumper settings within the camera electronics make it possible to select either $V_{\rm low}$, $V_{\rm ref}$, or $V_{\rm GNDA}$ (as defined in the pixel architecture discussion presented in chapter 4) as an ADC reference—though the last option is only offered as a failsafe measure. In the event of a transient on one of the analog supply voltages, the similar structures generating the analog residual voltage ($V_{\rm outsh}$, as also defined in chapter 4), $V_{\rm ref}$, and $V_{\rm low}$ should yield closely matching induced fluctuations. Because the digitized value is either $(V_{\rm outsh} - V_{\rm ref})$ or $(V_{\rm outsh} - V_{\rm low})$, this operation effectively eliminates the induced noise. As a final measure, judicious care is taken in setting the filtering constants before the ADC so as to suppress the high-frequency pick-up while passing as much of the low-frequency signal as possible.

5.1.3 Data Acquisition and Control

The data acquisition and control system for the Mixed-Mode PAD camera was the result of a evolutionary process beginning with the first 16×16 prototype ASIC and carried on to the large-area chip. The premise underlying this evolution has been to provide a rapid testing platform for the Mixed–Mode PAD prototypes. Because of this, we have avoided the complicated customized electronics adapted by our collaborators at ADSC,[3] instead relying on off-the-shelf electronics wherever

[3]The division of labor between our collaborators at ADSC and the group at Cornell has been that ADSC would design the compact, high-performance support electronics for the final multi-

possible.

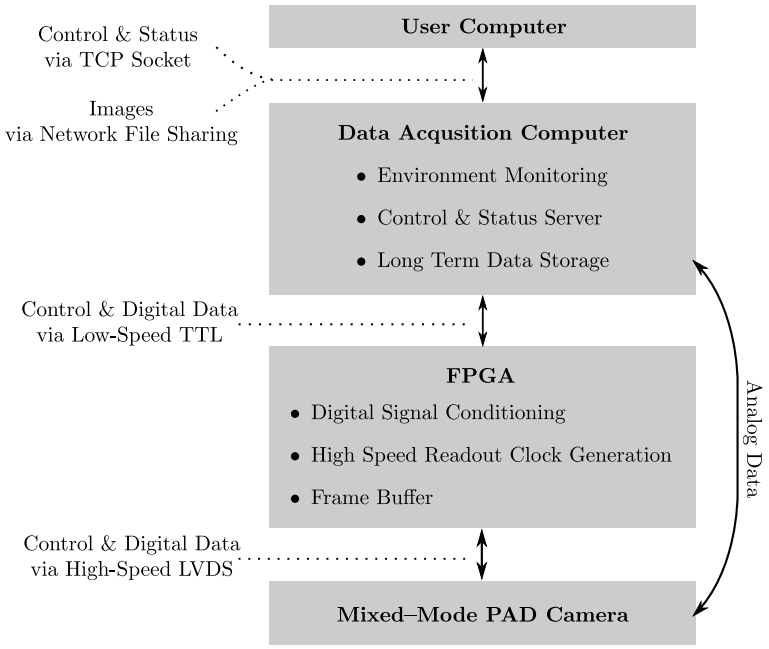

Figure 5.3: Control and data flow within the Mixed–Mode PAD Single Hybrid Prototype data acquisition & control system.

The data acquisition and control system for the Mixed–Mode PAD is comprises four elements:

- A client application running on a user computer.
- A server application running on the data acquisition and control computer that is responsible for controlling the GPIB[4] hardware, pattern generation, pattern capture, and analog-to-digital conversion hardware.
- An FPGA[5] that conditions the signals generated by the pattern generator,

hybrid detector while the Cornell group would focus on rapid testing and verification of detector prototypes.

[4]General Purpose Interface Bus, also commonly known as IEEE-488 and HP-IB (Hewlett-Packard Instrument Bus). A relatively slow, but simple, data and control bus predominantly used to automate test equipment.

[5]Field-Programmable Gate Array: an integrated circuit composed of programmable logic, called "logic blocks," and programmable interconnects in which custom, reconfigurable digital logic may be implemented.

generates the high speed clock sequences that read out the detector, and acts as a frame buffer for the digital data.

- The Mixed–Mode PAD Single Hybrid Prototype Camera.

This system is illustrated in the flow diagram shown in figure 5.3.

The use of a simple TCP/IP socket based protocol to exchange commands and queries between the Mixed-Mode PAD data acquisition controller and the user applications has proved very beneficial, as it allows a wide variety of software tools to control the imager. In particular, standard data acquisition and control tools, such as LabView (National Instruments–Austin, TX) and MATLAB (MathWorks–Natick, MA) as well as more synchrotron-specific software tools, such as Spec (Certified Scientific Software–Cambridge, MA) and ADX (Area Detector Systems Corporation–San Diego, CA), are easily extended to control the Mixed–Mode PAD. Most of these tools offer command line access to socket based communications allowing the interface with the Mixed–Mode PAD acquisition controller to be accomplished through runtime scripts. Where this is not an option, a C library is available along with an example application (*camcli*) to facilitate development of the control and status connection.

The Mixed–Mode PAD acquisition control application, known as *camserv*, is a multi-threaded server coded in a mixture of C and C++. When active, *i.e.* having an established connection with a user client, it operates two threads: a listener, which waits on the socket connection for requests from use client; and a monitor, which checks and logs selected parameters within the Mixed–Mode PAD environment, *e.g.* power supply voltages and current, supplying alarms when these values step too far out of range. Interactions with the Mixed–Mode PAD camera are carried out using control clock pattern generation and frame-buffering/frame-capture systems described in the subsequent sections.

5.1.3.1 Control Clock Pattern Generation

At the lowest level, each detector hybrid is controlled by a sequence of clock patterns. Generating these clock patterns is the combined task of a pattern generation module[6] and a Field-Programmable Gate Array (FPGA).[7] The purpose of the pattern generator is to provide an easy and fast means to alter the clock pattern sequence, a task that is made possible because the test pattern generator generates patterns based on a list of test vectors loaded into the generator's memory before every execution cycle. This pattern is created by software tools built into the Mixed–Mode PAD acquisition controller and, thus, very easily altered.

There are, however, practical limits to the pattern generator. The most notable of these is the maximum rate at which it can reliably output a sequence of test vectors. Our system shows good performance up to roughly 10 MHz. However, at higher rates the reliability is compromised. In addition, there exists a problem with maintaining a given output state over the periods between sequence generation because, during the reprogramming that occurs when a new series of sequences is loaded into the memory of the pattern generator, the pattern generator will return to the default, null, state. Our solution to these problems is to place an FPGA between the pattern generator and cryostat electronics to condition the control digital signals and generate any fast clock signals. The FPGA processes every pattern, taking a snapshot of it and then sending it on to the camera. Based on control signals from the pattern generator the FPGA may also latch certain states, holding them until the latch is released, or generate specific patterns such as the high-speed clocks that read out the detector, which the pattern generator is incapable of generating.

[6]Model UC.7221 (Strategic Test–Woburn, MA).
[7]Model Virtex-4 LX25 (Xilinx–San Jose, CA).

5.1.3.2 Readout and Frame Buffering

The readout of the detector is divided between the analog and digital portions of the detector data stream. The details of the readout clock sequencing will be discussed in section 5.2.4; here, though, we will be concerned with the paths these two data streams take from the hybrid to the control computer.

The analog data stream is comprised of a sequence of voltages derived from the residual voltages at the integrator output (V_{outp}) of each pixel. From the analog output on the detector hybrid, this voltage is initially buffered by a high-bandwidth, unity-gain buffer located very close, physically, to the hybrid. This buffer drives the voltage to a high-bandwidth instrumentation amplifier[8] located just before the outputs of the camera. As discussed in section 5.1.2, this device is a key element in the noise reduction system of the off-chip analog electronics, as it isolates the camera analog electronics from the external analog-to-digital conversion system, bandwidth limits the analog output signal, and makes the signal differential to eliminate common noise. The resulting voltage is driven to a high-speed analog-to-digital converter located in a cPCI[9] bus module[10] on the data acquisition computer.

This analog readout chain provides high fidelity data but suffers from rate limitations imposed, primarily, by the slew limits of the instrumentation amplifier. Between the internal limits of these amplifiers and the line capacitance each must drive ($\sim$ 3 m of coax cable) the analog readout rate is limited to $\sim$ 500 kHz, setting the deadtime of the single hybrid camera at just under 5 ms. While this is longer than the $<$ 1 ms targeted for the final Mixed–Mode PAD camera, tests on hybrids alone indicate that the $<$ 1 ms readout time is achievable if the bandwidth

[8]Model AD524 (Analog Devices–Norwood, MA).

[9]CompactPCI–a 3U or 6U industrial computer architecture, where modules are connected via a PCI backplane.

[10]Model UC.3021 (Strategic–Woburn, MA).

Table 5.1: Mixed–Mode PAD digital control signals. A line above a signal name indicates that the signal is active low.

Signal	Description
READ	Read enable.
CKA	Analog readout clock.
$\overline{\text{MRST}}$	Master reset.
CKD	Digital readout clock.
CKEN	Clock enable.
AMPEN	In-pixel output buffer enable.
PRST	In-pixel integrator reset.
PCL	CDS clamp.
PSH	In-pixel sample & hold.
CSRIN	Pixel control register data.
$\overline{\text{CSRCL}}$	Pixel control register clear.
CSRCK	Pixel control register clock.
DACIN	On-chip, refrence generating DAC data.
$\overline{\text{DACCL}}$	On-chip, referencegenerating DAC clear.
$\overline{\text{DACLD}}$	On-chip, reference generating DAC load.
DACCK	On-chip, reference generating DAC clock.

of the off-chip transmission and analog-to-digital conversion circuity is sufficiently large. To accomplish this, however, requires high-rate analog-to-digital converters integrated into the support electronics of the camera, a task that is underway as part of the custom control and data acquisition electronics being developed by our commercial collaborators at ADSC.

The digital data stream is somewhat more complicated, as the pattern capture electronics require periodic data while, for reasons that will be discussed in section 5.2.4, the digital data from the detector hybrid has interspersed pauses to allow the interwoven analog data to be sampled and recorded. To overcome this, the FPGA, which generates the readout clock, is configured to act as a frame buffer, temporarily holding one frame (exposure) worth of digital data until it may be read into the pattern capture cPCI module[11] in the data acquisition computer.

[11]Model UC.7021 (Strategic Test–Woburn, MA).

5.2 Selected Control Clock Patterns

There are a total of sixteen digital control lines used to operate each Mixed–Mode PAD hybrid, as detailed in table 5.1. In this section, we lay out, primarily for the benefit of those who may someday need to modify this system, how these clocks work together to control systems in the detector hybrid.

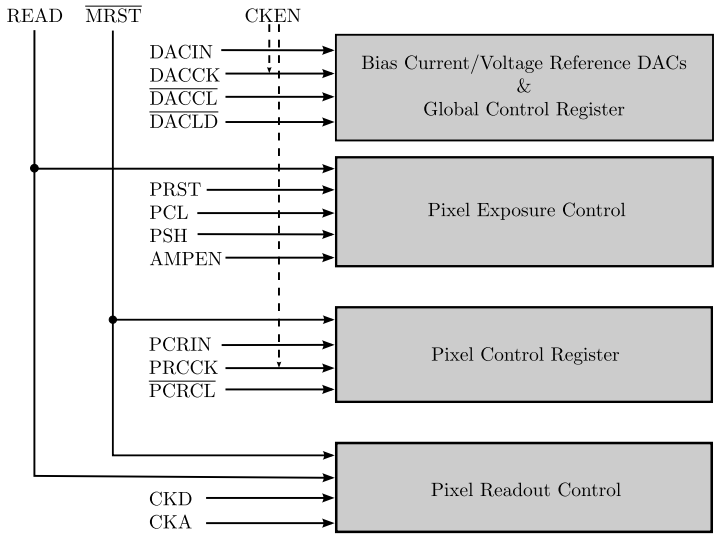

Figure 5.4: Relation between the Mixed–Mode PAD digital control signals, as defined in table 5.1, and systems on the detector hybrid. The CKEN signal does not directly affect any system on the chip, but is intended to act as a gate for the various system clocks to prevent errant cycles. On the AE207 submission, however, there is an error in the implementation of this line, and, thus, its use is not advised.

At a very high level each hybrid may be divided into four, nominally independent, logical systems, as illustrated by figure 5.4. The tasks carried out by these systems involve: configuring the global environment of the hybrid; controlling the pixel array during an exposure; configuring the in-pixel diagnostic register; and reading out the detector. For each of these tasks, a subset of the Mixed–Mode PAD control signals are used in a manner that will be explained.

Table 5.2: Summary of the elements of the Mixed–Mode PAD global environment register. This register contains the settings for the 6-bit DACs that control the reference voltages and bias currents used throughout the pixel array as well as additional bits that control aspects of the detector's behavior. More detailed information on these register elements may be found in [7].

Element	Type	Description
IMAST	DAC	Master Bias Current
ISS1	DAC	Integrator Amplifier Bias Current
ISS2	DAC	Comparator Bias Current
ISS3	DAC	Gated Oscillator Bias Current
ISS4	DAC	Sample & Hold, Sample Stage, Bias Current
ISS5	DAC	Pixel Output Bias Current
ISS6	DAC	Utility Buffer Bias Current
ISS7	DAC	Test Source Current
VREF	DAC	V_{ref} Voltage (typically set from 1.6 to 2.0 V)
VCAL	DAC	V_{cal} Voltage (typically set from 2.3 to 2.6 V)
VLOW	DAC	V_{low} Voltage (typically set from 0.8 to 1.2 V)
VTH	DAC	V_{th} Voltage (typically set from 0.6 to 1.1 V)
TMAST	BIT	DAC Master Current Test
TSLAV	BIT	DAC Slave Current Test
CLMODE	BIT	CDS Select
TCRNGE	BIT	Test Source Current Range

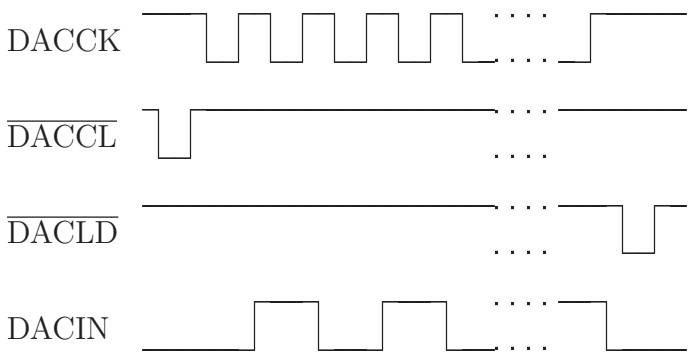

Figure 5.5: Timing for programming the on-chip bias and reference generating 6-bit DACs as well as the global control register. Data is latched on the falling edge of the DACCK signal so that the waveform shown here would load a hypothetical sequence of 01010 . . . 1.

5.2.1 Bias/Reference DACs & Global Control Register

The first logical system involves the global environment in which the Mixed–Mode PAD pixels operate. This includes programming the bias currents and reference voltages that control the circuit elements within each pixel as well as setting global logic bits that affect their operation. To accomplish this, in each detector hybrid there are a series of 6-bit Digital-to-Analog Converters (DACs) that generate the reference voltages and bias currents used within the pixel array. Two sets of registers program these DACs: the first is a daisy-chained shift register (with six successive bits for each DAC element and four most significant control bits, as organized in table 5.2) into which the DAC settings are loaded; the second is a latching register, with one for each element of the shift register, into which the shift register bits are loaded when the $\overline{\text{DACLD}}$ line (active low) is asserted. At the end of the shift register, there are four additional bits, with associated latch register elements, that provide diagnostics for the DACs and control various global aspects of the Mixed–Mode PAD behavior, as detailed in table 5.2. Figure 5.5 offers a timing diagram that illustrates the programming of this register.

5.2.2 Pixel Exposure Control

The second logical system within the detector hybrid controls the device during an exposure; the detailed operation of which depends on whether or not the analog CDS circuit is being used, as selected by a bit in the global control register. As figure 5.6 illustrates, the clock timing which controls exposures is relatively straightforward, with only a few details worth mentioning. First is the slight timing difference induced by the use of the analog CDS circuit. In both panel (a) and panel (b) of figure 5.6, the exposure duration is given by $\Delta t_{\text{exp}} = t_{\text{fin}} - t_{\text{bgn}}$; however, the transition demarking t_{bgn} changes depending on the use of CDS. The

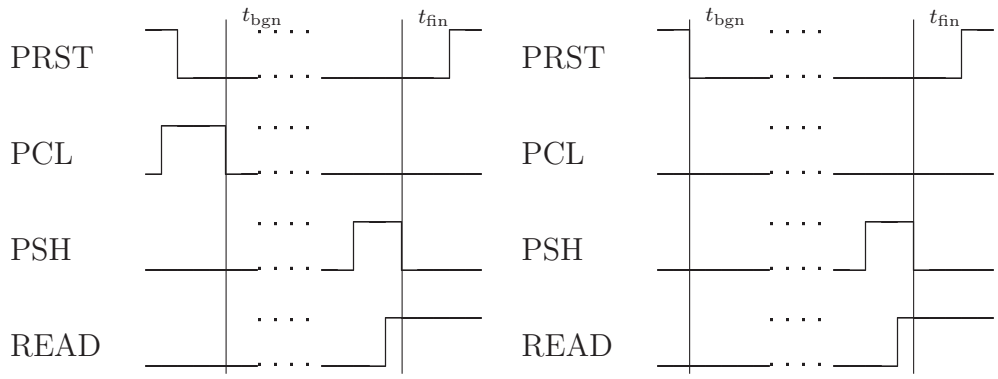

(a) Exposure with CDS (b) Exposure w/o CDS

Figure 5.6: Timing diagram for the control of an exposure in the cases where analog CDS is used, panel (a), and where it is not, panel (b). Note that the location of t_{bgn} changes between these two cases.

second point involves the timing of the PSH and READ signals at the end of the exposure. A consequence of asserting the READ signal is that the $\Sigma\Delta$ operation of the pixel is inhibited, allowing the integrator output (V_{outp}) to slew below the charge removal threshold (V_{th}) without initiating a removal. As such, it must be asserted after the PSH causes the pixel sample and hold to start to track the integrator or (if active) CDS output, but before the tracking ends in a sample.

5.2.3 Pixel Control Shift Register

The structure of the pixel control shift register (CSR) was already described in section 4.3.1; here, therefore, we focus on its programming. From the standpoint of the CSR, the array may be seen as one (row-wise) 128 element shift register coupled to 128 (column-wise) $128 \times 6 + 1$ bit shift registers where the first 128×6 bits represent the CSR elements of a complete column of pixels and the 1 extra bit is the CSR column select bit. To program the CSR, the last bit of the last register of each column is shifted into the array, with the right-most column shifted in

first. After 128 cycles of the CSRCK, clock an internal clock divider on the ASIC generates an internal load clock, causing the contents of the row-wise register to shift into the column-wise register. This sequence repeats, bit-by-bit, until the column-wise register is full.

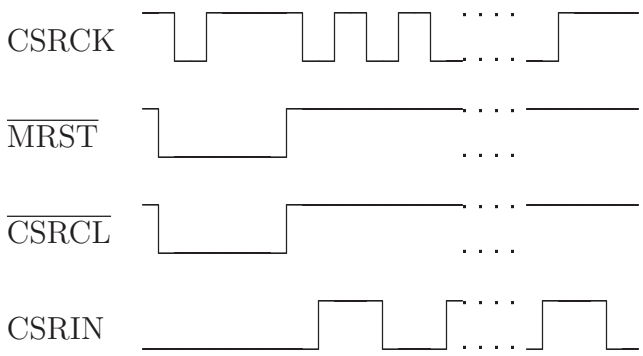

Figure 5.7: Timing for programming the pixel control shift register. The data shown represents a hypothetical sequence of $101\ldots01$.

5.2.4 Pixel Readout Control

The readout clock timing is by far the most delicate in the Mixed–Mode PAD, as it involves the careful interweaving of the CKD and CKA signals to bring the analog and digital data off the chip at high rates. To see the need for this, one must understand a few points of the bussing and multiplexing architecture used to convey analog and digital data from within a pixel to output buffers at the edge of the detector hybrid. Within the array, each column contains two data bus lines, one which carries analog and one which carries digital data. The analog data bus feeds into a multiplexer that connects to analog line drivers at the edge of the ASIC. The digital data bus, in turn, leads to a latching shift register that samples one bit of data from each column whenever latched, shifting this data

out at high rates in the interval between latches. This architecture is designed to provide maximal slewing and settling time for signals on the bus lines within the array by allowing this operation to occur as other data is being read off the chip.

Despite the tight interweaving of the analog and digital data streams, the control clocks that sequence this data (CKA and CKD, resp.) are largely autonomous. This is because, historically, the readout of the analog and digital data was completely independent. In this primordial design, the readout scheme for the digital data utilized separate shift registers that spanned each column, funneling the data into the, still existent, latching shift register. This architecture, however, was deemed to be too susceptible to catastrophic failures, caused by the loss of single shift register elements, to be acceptable—resulting in the move to the current, bussed, architecture. As the transition was relatively straightforward to carry out within the framework of the existing design, no effort was made to integrate the two clock sequences, consequently leaving a mildly maddening phase space of potential clock pattern sequences with hidden pitfalls that would corrupt the data of the unwary.

In the remainder of this section, we will try to unravel these digital, and then analog readout clocks in an intelligible manner. This process begins by describing each independently, in terms of the signals they generate within the hybrid, then merging our discussion to show how they can be interwoven to read off the data reliably.

5.2.4.1 Digital Readout Clock Timing

The digital readout clock generates two derived signals within each detector hybrid, as illustrated in figure 5.8. The first, labeled DLATCH, causes data on the digital bus to be recorded into the latching shift register, whence it is subsequently clocked off the hybrid with each CKD cycle. The second, labeled DIGADV causes new

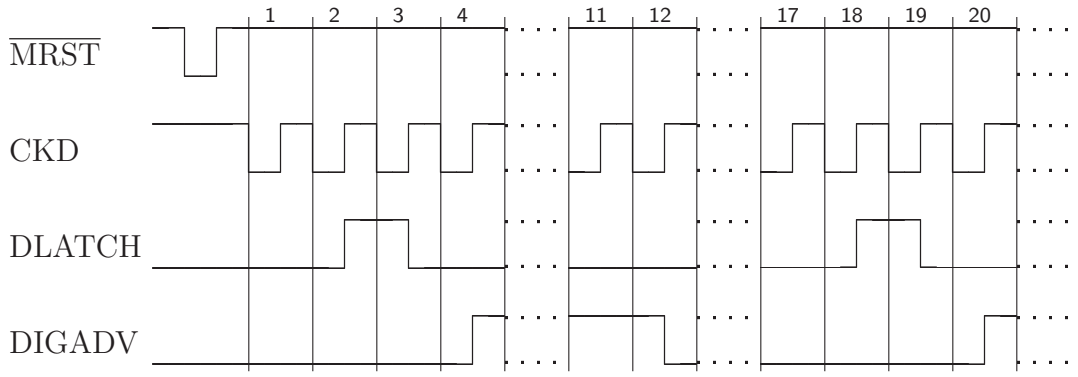

Figure 5.8: Timing controlling the readout of the Mixed–Mode PAD digital data. $\overline{\text{MRST}}$ and CKD are external signals defined in table 5.1 while the DLATCH and DIGADV signals are derived signals generated internally on each hybrid. The DLATCH signal causes the digital data on the array bus to be latched into the output shift register while DIGADV shifts data from the in pixel data register onto digital data bus of the array.

data to be placed on the digital bus by advancing the data register of each pixel in the active row—a process that, incidentally, resets the register as a new initial state is clocked in to replace the data that is clocked off.

5.2.4.2 Analog Readout Clock Timing

Slightly simpler than the system driven by CKD, CKA drives a single bit through a 17-bit shift register. When present in one of 16 bits of the register, one of the voltages on the analog bus is multiplexed to the analog line drivers and sent off-chip to be digitized. When present in the 1 remaining bit, the row-select register (responsible for gating the analog and digital data of a given row onto the readout buses) is advanced.

5.2.4.3 CKD & CKA Interweaving

The first CKA falling edge following a $\overline{\text{MRST}}$ cycle advances the row-select register, placing the analog and digital data from the first row of pixels onto the analog and

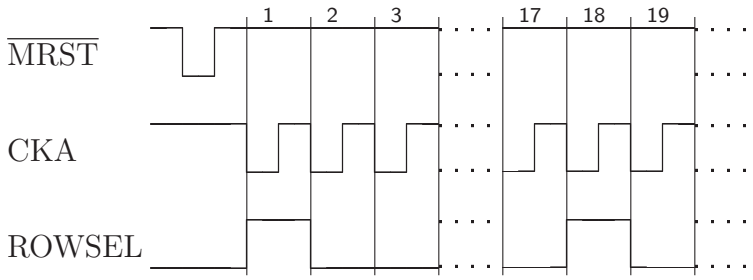

Figure 5.9: Timing controlling the readout of the Mixed–Mode PAD analog data as well as the row select logic for the digital data. $\overline{\text{MRST}}$ and CKD are external signals defined in table 5.1 the ROWSEL signal is a derived signal generated internally on each hybrid. The ROWSEL signal is responsible for advancing the row select shift register (a 128 element single-shot shift register, reset when $\overline{\text{MRST}}$ is asserted).

digital output buses of the array. Before this occurs, however, care must be taken to prepare the state of the data registers within the pixel to prevent write back from occurring when the digital data is first connected to the bus. This situation is the unfortunate consequence of the omission of a buffer between these data registers and the digital output bus. The flip-flops used in these registers are a sequence of two identical stages, each of which may either be in a follow or drive mode. When in the follow mode the stage takes on the state of its data inputs. With a transition from the follow to the drive mode, the last state seen in the follow mode is recorded by the stage and driven on its outputs. The clocks that drive these stages are 180 deg out of phase so that when one stage is in the follow state the other is in the drive, and vice-versa. The problem of write back occurs when the second stage of the last element in the pixel data registers is connected to the bus while in drive mode, as it is possible for the state of the bus to overwrite the state recorded in this register due to its much larger capacitance. To avoid this, the stage must be in the record mode so that the data is isolated from the bus, and, hence, the DIGADV-derived clock must be low.

Figure 5.10 illustrates how the data readout clocks are sequenced in the single

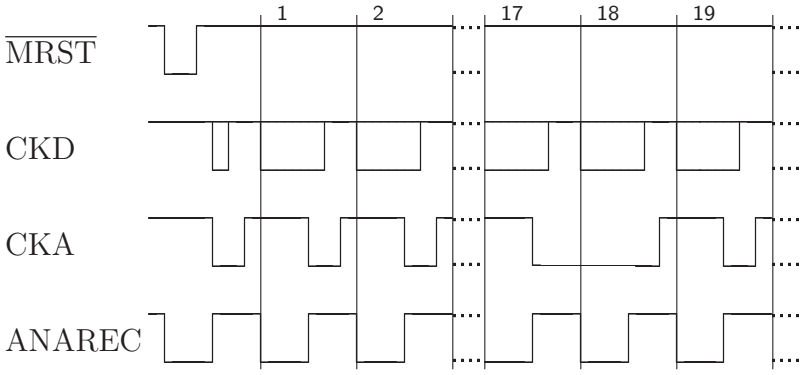

Figure 5.10: Combined readout clock sequencing used in the single hybrid camera. In this diagram, the CKD signal is not shown at full resolution due to space limitations. Instead, in its first active region, there are two cycles, and, in each subsequent region, there are sixteen cycles (denoted by the high/low logic region). The ANAREC signal denotes the sampling clock used by the ADCs to time recording of the analog data. Because of the internal structure of the ADCs it must be periodic, resulting in two redundant samples for every 16 pixels worth of analog data. The first valid digital data comes off the chip in period 2.

hybrid camera. An additional signal (ANAREC) is included on this diagram to indicate the timing of the sampling clock for the ADC.[12] This sequencing is contrived to offer as much settling time as possible on the various data lines while still being relatively straightforward to implement within the FPGA.

[12]Although it is not evident from figure 5.10, the ANAREC falling (sampling) edge preceeds the activity on the CKD lines by a brief period, smaller than the resolution of the timing diagram.

CHAPTER 6

DETECTOR CHARACTERIZATION

A prerequisite to conducting experiments with the Mixed–Mode PAD is sufficient understanding of the response of the imager to allow confident interpretation of the images it produces. In general, an imaging device will introduce artifacts into the raw images it produces. In a pixelated device like the Mixed–Mode PAD, these result from a combination pixel-level effects, such as nonlinearities in the response and electronic noise, pixelation effects, which strongly influence the imager's spatial response, and pixel-to-pixel variations. An imager needs to be characterized with regards to these artifacts and, where necessary, calibration and correction procedures developed to remove these effects via post acquisition processing.

This chapter deals with these characterizations, identifying primary factors contributing to uncertainty in measurements made with the instruments and presenting baseline measurements of imager performance. Following this characterization, the degree to which this performance meets what would be expected from an ideal imager is evaluated through a metric called the detector quantum efficiency, which will be defined in section 6.5. Based on this assessment, areas in which calibration and correction could improve image quality are identified, followed by a discussion of procedures to achieve this. As we will see, for a complicated device like the Mixed–Mode PAD, the intricacies of characterization and calibration is one of the chief obstacles to achieving the performance limits of the detector. Finally, as this device is intended for use in an intense radiation environment, its tolerance to radiation damage is assessed.

6.1 Linearity

Though linearity is often assumed by detector users, it should not be by instrumentation developers. Demonstrating the linearity of an imager is therefore an essential first step in assessing the performance of an imager. For the Mixed–Mode PAD, this assessment in nontrivial due, on the one hand, to the fundamental nonlinearity of the device,[1] and, on the other, to the large well depth of the each pixel, which spans from 1 to more than 25 million 10 keV x-rays, for which there is no practical means of performing an end-to-end linearity assessment. Fortunately, these issues may be addressed simultaneously by demonstrating the limited linearity of the digital and analog components of the Mixed–Mode PAD data stream separately, and, then showing that there exists a scaling factor which allows these separate portions of the data stream to be smoothly combined into a response that is linear over its full domain. In this way, we may satisfy ourselves that the system is indeed linear over its entire range.

The digital portion of the linearity assessment was conducted by our collaborators at ADSC. It involved injecting a known test current onto the pixel integration node via a needle probe connection to the bump bonding pad of a pixel on an un-hybridized ASIC. While this was occurring, a portion of the pixel diagnostic monitoring system was used to track this pixel's analog residual voltage (V_{outp}), recording the resulting rate of charge removal operations. The test current was sourced and measured by an HP4145 Semiconductor Parameter Analyzer and spanned a range from 1 pA to 100 nA. The charge removal frequency was measured by a Fluke 45 Multimeter. A large resistance ($\sim$10 MΩ) was placed in series with the needle probe to decouple the capacitive portion of the load pre-

[1] Only taken together and with appropriate calibration factors do the analog and digital data streams from the Mixed–Mode PAD form a linear data set. Independently, both data streams are decidedly nonlinear due to the effect of the charge removal operation.

sented by the current source from the input to the front-end amplifier. Figure 6.1 shows the results of this measurement. At current levels below 20 pA the data appear to notably depart from the fitted linear response. The actual magnitude of this deviation is much smaller than the figure suggests, with its appearance enhanced as a result of the log-log scale, and is at a level that is consistent with the variation present in other portions of the data. Generally, a low-signal non-linearity is unlikely to be due to the Mixed–Mode PAD as the $\Sigma\Delta$ front-end within each pixel is a charge sensitive, rather than current sensitive, architecture. As a result, any deviance from a linear response would be much more likely under high-signal condition. As a verification of these arguments, later in this section we will present another, independent, measurement that tests the low-signal linearity of the Mixed-Mode PAD response.

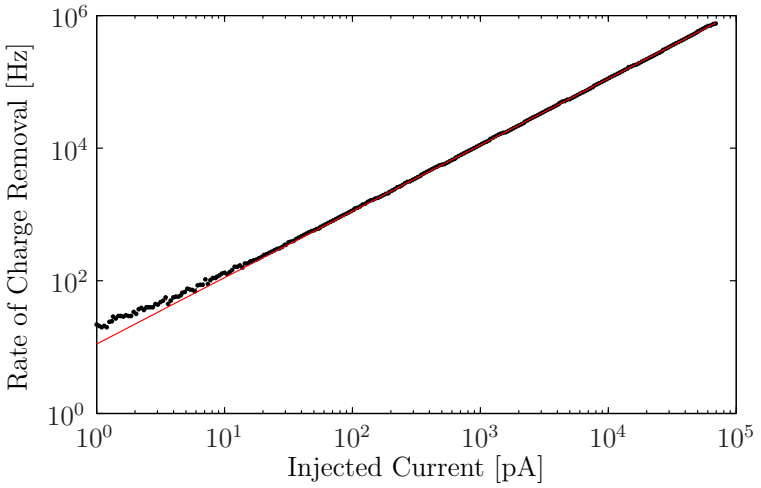

Figure 6.1: Linearity of the digital portion of the Mixed–Mode PAD data stream, shown as the rate of charge removals as a function of stimulating current. This measurement was made by sourcing a known current onto the pixel integration node via a needle probe connection to an unbonded pixel, in the manner discussed within the text.

To evaluate the linearity of the analog portion of the data stream and to verify that it is possible to combine this data with the digital response in a smooth

fashion, one would like a stable current source capable of generating the equivalent signal of a few thousand 10 keV x-rays in a reasonable amount of time. Fortunately, as we will see, such a source is readily available in the form of the leakage from the Mixed–Mode PAD detector diode. To understand how the detector's dark current can become a good weak signal calibration source, we need to look at how this current is generated.

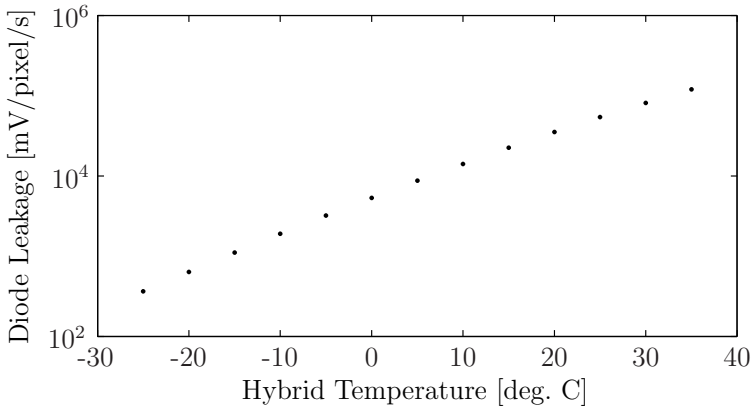

Figure 6.2: Average per-pixel leakage current from observed by interior pixels of a Mixed–Mode PAD hybrid as a function of temperature. The dependent axis is plotted in terms of mV/pixel/second, because this is what is directly measured from the integrator. The direct conversion to charge depends on the absolute conversion factor of the integrator in each pixel, which, in turn, depends on the size of the integration capacitor. The integration capacitor was laid out have a capacitance of 50 fF, however, measurements indicate that its actual capacitance is 20% to 30% larger than expected.

The Mixed–Mode PAD detector diode leakage current consists of a combination of generation current from within the detector depletion region and diffusion current from the quasi-neutral areas at the depletion region boundary. As such, the total leakage current (J_{lkg}) is given by

$$J_{\mathrm{lkg}} = \int_{\partial \mathrm{dep}_n} dx^2 \, q \sqrt{\frac{D_{\mathrm{h}}}{\tau_{\mathrm{h}}}} \frac{n_i^2}{N_{\mathrm{D}}} + \int_{\mathrm{dep}} dx^3 \, q n_i G_{\mathrm{h}}, \tag{6.1}$$

where the first integral is over the area of the quasi-neutral region within the n-type silicon and the second integral is over the complete volume of the depletion

region. The first term in this expression is associated with the diffusion current and the second is the thermal generation current, as described by Shockley-Read-Hall recombination [64]. Of the parameters within this expression, the dominant external influnce comes in through the intrinsic carrier concentration (n_i) along with the depletion region volume and bounding. The intrinsic carrier concentration is dependent on the detector temperature via

$$n_i \propto \left(\frac{T}{300}\right)^3 \exp\left\{-\frac{E_G}{2kT}\right\}, \tag{6.2}$$

while the dimensions of the depletion region are determined by the diode reverse bias voltage [93]. If these are kept stable, then fluctuations in the detector leakage current will be governed by the detector shot noise, which, as a Poisson process, will exhibit integral fluctuation of

$$\delta J_{\mathrm{lkg}} = \sqrt{q J_{\mathrm{lkg}} t_{\mathrm{exp}}}, \tag{6.3}$$

where J_{lkg} and t_{exp} denote the detector diode leakage current and exposure duration respectively. The average per-pixel leakage at 20 deg. C was measured to be ~ 300 fA; for pixels within the array interior, this source should exhibit fluctuations of less than 0.5 fA$\cdot\sqrt{\mathrm{s}}$. Thus, so long as the high voltage and temperature are sufficiently stable, it is possible to use the diode leakage to test the analog linearity of the imager.

Achieving thermal and bias voltage stability with this detector is relatively straightforward, as studying a hybrid within its camera enclosure provides for these. Care, though, needs to be taken to ensure that no light may enter the camera while the measurement is made, so the window needs to be covered with an opaque material. The first test measures the linearity of the analog response by investigating a single pixel under the influence of the detector leakage. The temperature of the detector was set to 20.0 ± 0.2 deg. C and held there by a

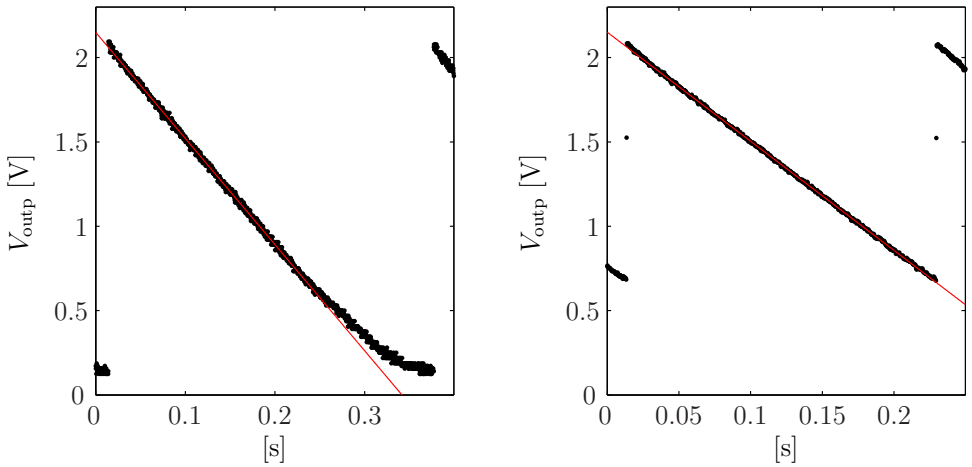

(a) V_{outp} Extended Range (b) V_{outp} Linear Range

Figure 6.3: Traces from a pixel integrator output under the stimulus of the diode leakage current. Panel (a) was taken with V_{ref}, V_{low}, and V_{th} set to extend the range of V_{outp} into the nonlinear region of the integrator. In panel (b), V_{ref}, V_{low}, and V_{th} were set to show the linear range of the integrator. Under normal operating conditions V_{ref}, V_{low}, and V_{th} are set so that the integrator output will remain within a ~ 1 V subset of this range. In both panels a linear fit is shown to illustrate the integrator linearity or deviance therefrom.

thermoelectric within the cryostat. As discussed in section 4.2.1.3, there is a portion of the range of the integrator output where its response is not linear, as illustrated in panel (a) of figure 6.3. However, as panel (b) of figure 6.3 shows, there is a sufficiently large linear range of the integrator response to meet the 1 V of linear slew specification. For the measurements reported here care was taken to ensure that the integrator output was set within the linear operating range.

The final step of the linearity test involves showing that it is possible to merge the analog and digital data sets in smooth fashion. From the discussion in section 4.1, the equivalent voltage (V_{eqv}), that is the voltage that the analog output would ideally slew if no charge removals occurred, is given by

$$V_{\text{eqv}} = \frac{dV_{\text{eqv}}}{dN_{\Delta Q}} N_{\Delta Q} + V_{\text{outp}}. \tag{6.4}$$

When the stimulus is constant, with the exception of the brief discontinuity when

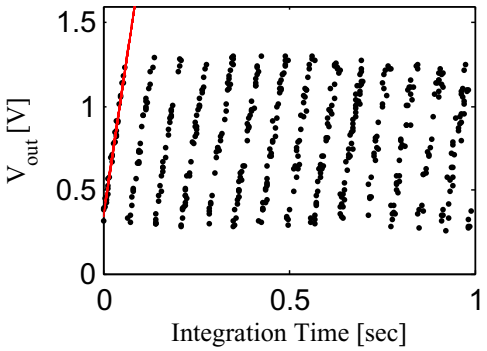

(a) Raw Analog Data

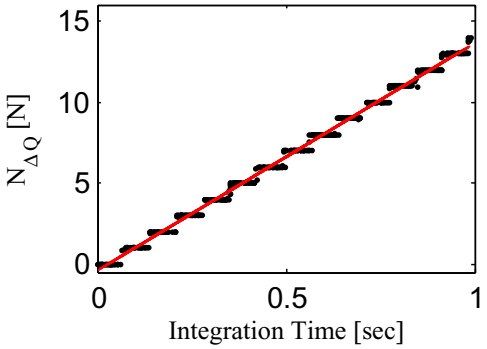

(b) Raw Digital Data

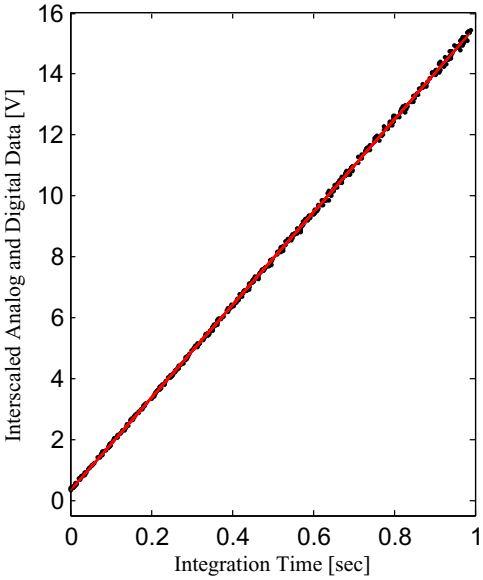

(c) Combined Analog and Digital Data

Figure 6.4: Typical analog (V_{outp}), digital ($N_{\Delta Q}$), and merged (V_{eqv}) data for one pixel from leakage current integration series. Detector was held at 20 deg. C, isolated from ambient light during these exposures.

charge is removed, the rate of change in the analog residual voltage and the rate of change in the equivalent voltage are related by

$$\frac{dV_{\rm outp}}{dt} = \frac{dV_{\rm eqv}}{dt}. \tag{6.5}$$

In addition, when the stimulus is constant

$$\frac{dV_{\rm eqv}}{dt} = \frac{V_{\rm eqv}(t_1) - V_{\rm eqv}(t_0)}{N_{\Delta Q}(t_1) - N_{\Delta Q}(t_0)} \cdot \frac{N_{\Delta Q}(t_1) - N_{\Delta Q}(t_0)}{t_1 - t_0} = \frac{dV_{\rm eqv}}{dN_{\Delta Q}} \cdot \frac{dN_{\Delta Q}}{dt}, \tag{6.6}$$

where $t_1 \neq t_0$ and these times are chosen to be the midpoint of a $N_{\Delta Q}$ step, *i.e.* the average integration time yielding a given $N_{\Delta Q}$ value. Then, combining these results we find that

$$\begin{aligned} & \frac{dV_{\rm outp}}{dt} = \frac{dV_{\rm eqv}}{dN_{\Delta Q}} \cdot \frac{dN_{\Delta Q}}{dt} \\ \Rightarrow \qquad & \frac{dV_{\rm eqv}}{dN_{\Delta Q}} = \frac{dV_{\rm outp}}{dt} \left(\frac{dN_{\Delta Q}}{dt} \right)^{-1}. \end{aligned} \tag{6.7}$$

While a single measurement will not yield $\frac{dV_{\rm outp}}{dt}$ or $\frac{dN_{\Delta Q}}{dt}$, they may be easily extracted from an ensemble of measurements taken under identical constant stimulus conditions with varying integration times. This method is illustrated in figure 6.4. Here, panels (a) and (b) contain, respectively, the analog and digital data from one pixel for a variety of integration times with the stimulus held constant. The digital count scaling constant is derived from the ratio of the slopes of the two linear fits, and the merged data set is shown in panel (c). These data sets were taken using the leakage from the detector diode at 20 deg. C as a stimulus where, as before, the camera cryostat was used to maintain the thermal stability of the hybrid and care was taken to isolate the detector from ambient lighting. The series spans exposure durations from 500 μs to 1 s with individual integration times randomly ordered to remove any temporal systematics. Each data point shown represents the averaged response of 25 consecutive frames taken with identical settings. The full data set contains 2,000 different integration times taken over the period of a day.

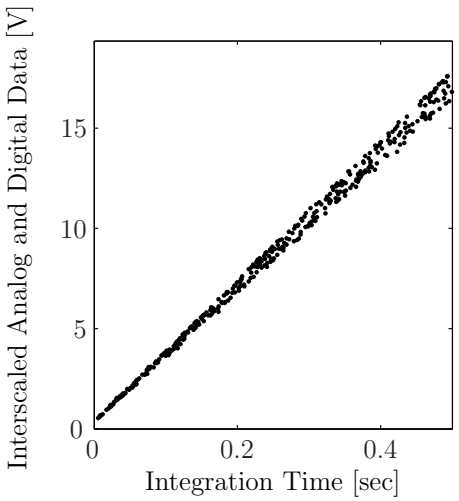

Figure 6.5: Typical merged analog and digital data (V_{eqv}) from one pixel in a Cu x-ray tube exposure series. For this series, the tube was operated at 25 kV with a current of 0.4 mA, hybrid temperature was set at 20 deg. C. Data scaling factors were calculated from dark current integration series as discussed in the text.

As a check, both of this interscaling method and as a validation of the front-end linearity at low signal levels, we utilized our calculations of $\frac{dV_{\mathrm{eqv}}}{dN_{\Delta Q}}$ to merge a similarly recorded data set in which the detector, operated at -35 deg. C, was illuminated with a flat flood field produced by a Cu x-ray tube operated with a 25 kV bias and 0.4 mA tube current. The results are shown for a representative pixel from the array in figure 6.5. The signal produced by this flux resulted in roughly 33 charge removals per second as opposed to 18 with the dark current alone. If the Mixed–Mode PAD was noticeably nonlinear at low signal rates, as discussed earlier, then one would expect to see a saw tooth pattern superimposed upon this ramp. However, as figure 6.5 shows, the interscaling is quite smooth, with no apparent nonlinearities.

6.2 Pixel Electronic Noise

Along with linearity, the electronic noise of individual pixels needs to be evaluated to understand the pixel-level performance. The purpose of this is twofold: first, it provides a good assessment of the front end circuit design, and, second, it helps to decouple effects of individual pixels from pixel-to-pixel variations in assessing aggregate performance metrics like the detector quantum efficiency.

Fortunately, this measurement is straightforward, as the data from the leakage-current-based tests of the detector linearity may also be used to measure the detector electronic noise. As mentioned, at each integration time, a set of 25 images were taken. Within each of these sets and for each pixel, one may calculate a RMS of the merged analog and digital data. Collecting the RMS calculations for pixels across the array into a histogram yields a distribution that is statistically Gaussian,[2] as figure 6.6 illustrates. This is not surprising, as it merely indicates that the distribution represents quantities derived from the same statistical ensemble, which one expects so long as the total signal acquired is not sufficient to bring out systematic differences in the pixel response or in the leakage currents that the pixels are measuring.

What is perhaps surprising is the effect that a global correction of the mean response of the array has on the data. In figure 6.6, the distribution labeled A shows the uncorrected results for this frame set. Distribution B describes the same data; however, in it, each image was globally corrected by subtracting off the image mean before calculating pixel-by-pixel RMS. Figure 6.7 collects this data for a series of integration times, showing that the system noise response is consistently

[2]There may be outliers to this distribution (*e.g.* hot pixels due to surface damage to the detector or edge leakage, or silent pixels due to unconnected bump bonds) that do not represent *normal* pixel behavior, but can significantly skew statistical calculations, such as the mean and variance. By restricting our consideration to the Guassian set of pixels, we automatically cut out these aberrant elements.

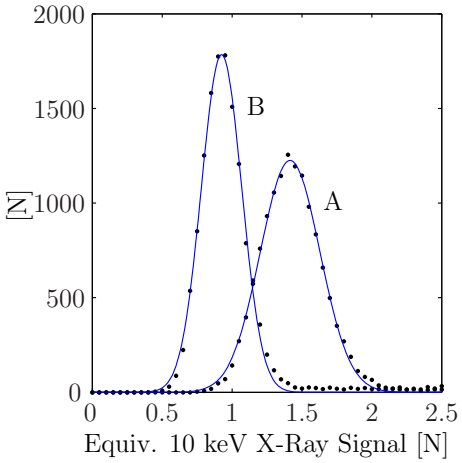

Figure 6.6: Two pixel RMS distributions derived from the same series of 25 frames taken where all frames had the same integration time. The Gaussian peaked at a right, larger average RMS, was derived from uncorrected data while the Gaussian at the left, smaller average RMS, was corrected for global shifts in the array via a mean subtraction. This data assumes a 1 mV = 1 keV conversion gain.

improved by this global correction.

This improvement indicates that there is a source of global noise within the prototype camera that is significant at low signal levels. There are numerous areas within the external electronics that could be the source of this noise, including: the bias on the detector diode, the power supply lines, the current reference from which the on-chip digital-to-analog outputs are derived, etc. Within the prototype, a full understanding of this noise source is not essential, as it is easily accounted for in post processing—so long as it is possible to identify a reference region within each image with which the global shift may be calculated. This, however, is something that needs to be evaluated and, if possible, corrected, in the process of building custom support electronics for the multi-hybrid, large active area imager.

This point aside, figure 6.7 reveals some very interesting points regarding the weak-signal performance of the Mixed–Mode PAD. First, although the data spans nearly 20 charge removal operations there are no apparent discontinuities in the

noise curve. This indicates that the noise contributed by each charge removal operation is negligible relative to the noise in the source, the detector diode leakage current at 20 deg. C. Secondly, the level of the electronic noise is far below the fluctuation one would expect in the Poisson Statistics of a similar quantity of x-rays signal. Thus, for flux levels of at least a few x-rays/pixel/s, fluctuations in the x-ray signal should dominate any measurement.

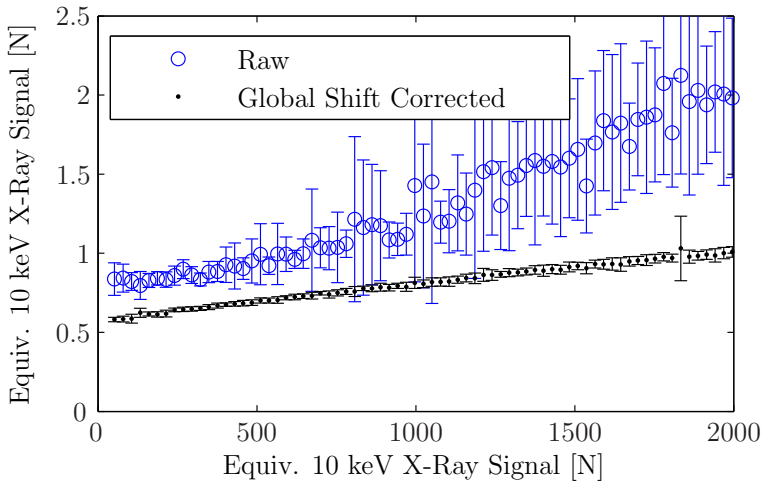

Figure 6.7: Detector noise as a function of accumulated diode leakage current with the hybrid maintained at +20 deg. C in the camera. To generate this figure, measurement statistics were calculated from sets of 25 images acquired at 1000 integration times randomly distributed from 1 ms to 1 s. The range of signal observed was divided into 75 evenly spaced bins into which the mean per-pixel RMS values, based on a Gaussian fit as described in the text, were divided based on their corresponding mean signal. The data point plotted then indicates the mean, mean per-pixel RMS in each bin and the error bars indicate the RMS fluctuations about this mean. The units on the horizontal and vertical axes are given in equivalent 10 keV x-rays (assuming a 1 mV = 1 keV conversion gain) to make the comparison to an experimental signal more straightforward, although the ordinate axis could equivalently have been labeled in time spanning up to 1 s.

For signals weaker than this, we may look to figure 6.8 to understand the contribution of electronic noise to the total measurement uncertainty. This figure describes the fluctuations in measurements of the diode leakage current with integration time, where the detector was operated at -25 deg. C in a dark environment.

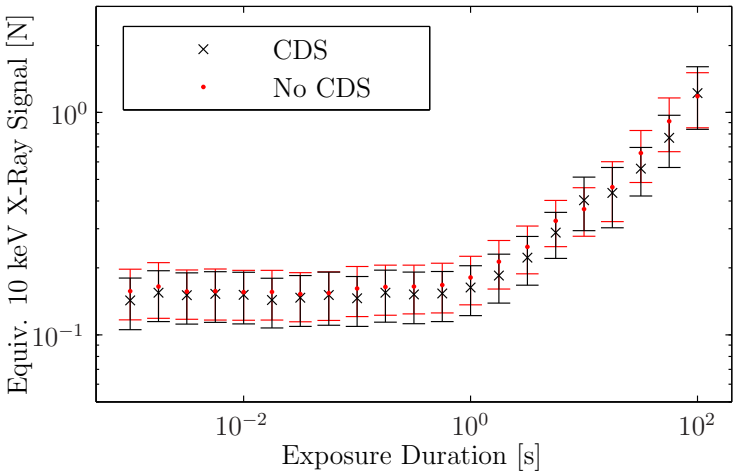

Figure 6.8: Detector noise observed over a series of diode leakage current integration extending from 1 ms to 100 s, taken with and without CDS. For these measerements the hybrid was maintained in the camera housing at -25 deg. C and exhibited an average leakage level 29.5 x-rays/s (assuming a 1 mV = 1 keV conversion gain).

The points on this curve each represents data from a series of 25 images. As discussed earlier, a global correction was applied to each image based on the mean signal observed. From this corrected image set, RMS fluctuations for each pixel were calcuated and histogrammed. This histogram was then fit to a Guassian, with the resulting centroid and width of this fit used for the data points and error bars in figure 6.8. Data sets were taken both with and without analog CDS. In agreement with our discussion from section 4.3.2, the benefits of analog CDS are small but present up to time scales of a few seconds.

If one could assume that the full signal from each x-ray was observed by a single pixel, then, based on our discussion in section 2.3, the uncertianty contributed by fluctuation in the signal from each x-ray will be negligable, for the weak x-ray fluxes we are considering, in comparison to this noise source. Consequently, for a dark hybrid at -25 deg. C, the uncertainty in a measurement of N x-rays over t s will be the electronic noise fluctuations from figure 6.8. In this idealized case, one

should be able to confidently measure quantized sets of x-rays up to integration times of nearly a minute or farther, in the event the analog oversampling discussed in the conclusion to chapter 4 were implemented.

Unfortunately, this picture is complicated by the fact that the x-ray signal may be distributed between more than one pixel as a result of the spreading of the photocurrent as it drifts through the detector diode. The characterization and analysis of this spreading are the subject of the next two sections (sections 6.3 and 6.4) followed by the impact of this effect on the measurement accuracy in section 6.5.

6.3 Charge Collection

Having verified the linearity of individual pixels and developed a picture of the noise performance that can be expected from the pixel electronics, we may begin to look at the quality of images produced by the imager. This begins with an assessment of the detectors signal collection characteristics.

As discussed in section 2.2.2, when x-rays convert to charge carriers within the detector diode, diffusion of the generated carriers causes a notable lateral spreading of the photocurrent signal, the extent of which depends predominantly on the detector's detector diode bias and, to a lesser extent, on the temperature of the diode layer (through its effect on the diffusion coefficient). This phenomena is illustrated in figure 6.9, which shows the raw x-ray beam from a Cu rotating anode source, masked to produce a roughly circular illumination field, slightly larger than a typical diffraction spot, at differing high-voltage settings. Figure 6.10 offers a more direct illustration of the effect's magnitude, showing the changes in the amplitude profile along a line taken through the spot center at differing bias voltages. It should be noted that above $\sim$100 V the profile of the spot changes very

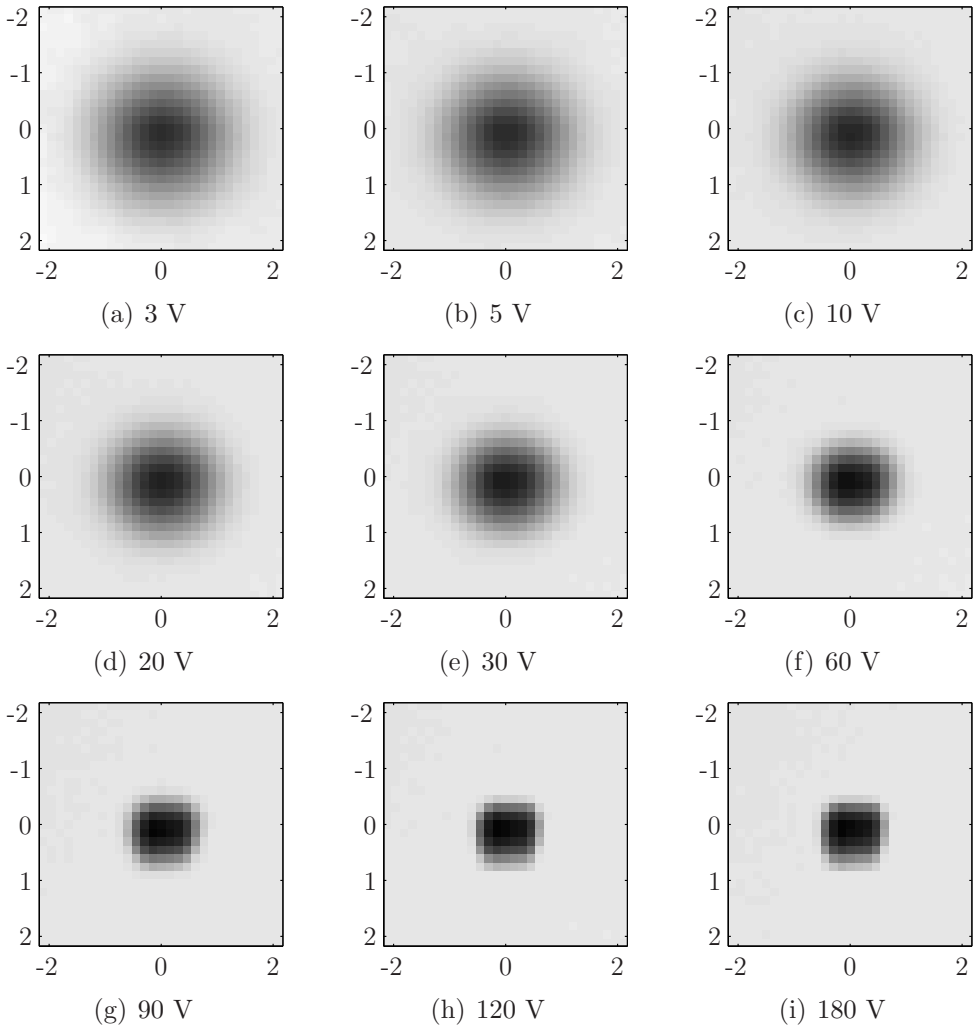

Figure 6.9: Multipixel x-ray spot generated by a Cu target rotating anode source, imaged at differing detector diode reverse bias voltages. Images were acquired with identical integration times and are shaded using the same logarithmic grey scale, to bring out both faint and intense features. Vertical and horizonal axis units are mm.

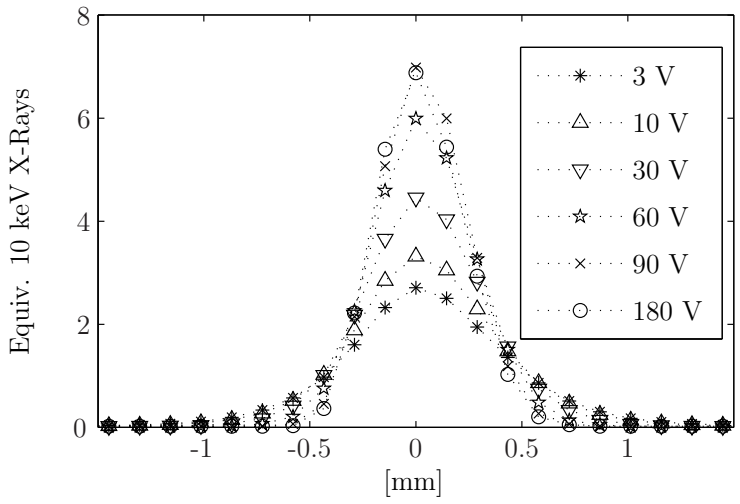

Figure 6.10: From selected images in figure 6.9, x-ray spot intensity profile taken along a vertical line through the center of the pixel, for differing detector diode reverse bias voltages.

little with increasing bias. Calculations suggest that, at this reverse bias voltage, the detector diode layer is nearly fully depleted.

While reducing the high voltage blurs the incident signal by redistributing it over a larger number of pixels, the total amount of charge collected is not significantly diminished. This is shown in figure 6.11, which depicts the total integrated dose from a flood field source as a function of the detector diode bias. The decrease in efficiency as the detector diode bias falls below the full depletion level of ∼100 V is explained by the growing size of the detector diode's undepleted region. Recalling the discussion of the diode structure from section 2.3.1, the diode's P/N junction occurs on the side of the wafer opposite the face where x-rays enter. As a result, when the diode is not fully depleted, the undepleted region is located between the x-ray source and the depletion region. This undepleted region lacks the strong electric field, present in the depleted region, that rapidly sweeps charge carriers to the integration node of the pixel. Instead, the carriers diffuse until they either fall into the depletion region or recombine. As the size of the undepleted

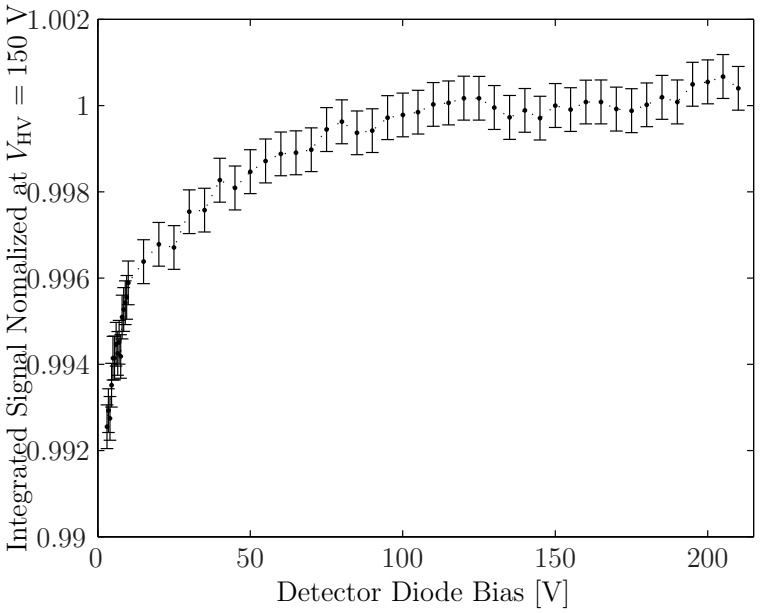

Figure 6.11: Total acquired dose, integrated across the full detector, of a flood field as a function of detector diode bias. The flood field was generated by a Cu x-ray tube biased at 25 kV and the integration time was held constant over all measurements. Results are normalized to the dose measured at a bias of $V_{\mathrm{HV}} = 150$ V.

region grows with reduced bias voltage the likelihood that a carrier will recombine before reaching the depletion region increases, resulting in the small reduction in charge collection efficiency shown in figure 6.11.

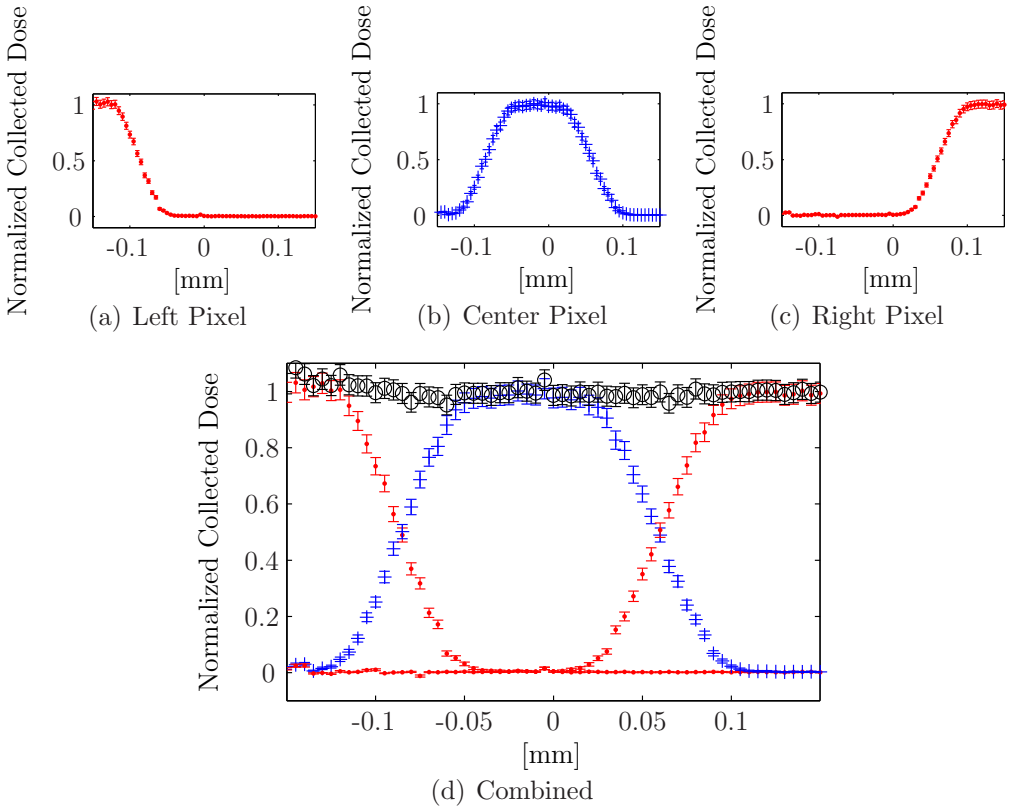

Figure 6.12: Mixed–Mode PAD charge collection. These figures show the charge collection from a 75 μm spot source of x-rays as it is translated along a line near the bisector of three pixels sharing the same row. Panels (a), (b), and (c) show the dose collected in each pixel normalized against the average of the sum of the dose measured in the three pixels at each spot location. These individual measurements are combined in panel (d) along with the sum of the dose measured in the three pixels at each spot location (denoted by the open circles with error bars). This measurement indicates that no charge is lost in the regions between pixels.

Although charge may not be lost by the diode layer the total carrier yield from a single x-ray conversion may not be completely collected by a single pixel. Between any two, three, or four adjacent pixels there is a region wherein photocurrent may split between these pixels. This effect is illustrated in figure 6.12, which shows

the collected signal within a pixel and two of its neighbors as a 75 μm spot is translated in 5 μm steps across the detector. If conversion gain and distortion effects are correctly accounted for, then the total signal collected will be constant.

6.4 Spatial Response and Resolution

The preceeding discussion of signal collection is sufficient to describe the imager spatial response for much of the work that will be done with the Mixed–Mode PAD. This is because, for many of its intended applications, most notably protein crystallography, the role of the Mixed–Mode PAD will be to measure the position and intensity of discrete spots or rings of x-ray signal. For these applications, one generally does not need an understanding of the imager spatial response that extends beyond the extent to which flux incident at a point on the diode above one pixel is detected by its neighbors. However, applications producing extended and complex patterns, as one can find in x-ray tomography, require a more detailed understanding of the imager's spatial response.

This section presents an evaluation of the spatial response of the Mixed–Mode PAD directed at both these concerns. It begins with a discussion of the effect that discrete sampling and limited imager active area have on the observed signal, particularly its spectral representation. We then discuss data collection and refinement methods used to study the pixel spatial response with sub-pixel resolution. This leads to an evaluation of the spatial and spatial frequency response of individual pixels, concluding with a discussion of inhomogeneity in the response of different pixels.

6.4.1 Discrete Sampling of Limited Active Area

Pixel Array Detectors, as fixed-grid discrete sampling devices, present a problem for imaging applications because of the discrepancy between the extent of the diode impulse response and the pixelation effects of averaging and discrete sampling. Traditional metrics for evaluating the imager's spatial response, such as the point spread function and modulation transfer function, are poorly defined for sampled image systems [74]. These metrics were created to characterize continuous imaging systems, *e.g.* film [54], where the imager response is assumed to be translation invariant (isoplanatic), so that the image of a point source is independent of its location. In the Mixed–Mode PAD and other discrete sampling imagers, this symmetry is broken by the pixel grid. While this effect is not new, with extensive discussion in the literature extending back into the early 1980s ([74], [31], [81], and references therein), the problem is particularly acute for the Mixed–Mode PAD, because the lateral extent of the detector diode's impulse response (20–30 μm under typical bias conditions) is much smaller than the pixel geometry (150 μm $\times$ 150 μm). As a result, the pixel geometry and the effects of discrete sampling are the dominant factors determining the imager spatial response.

To evaluate the effects of pixelation on the imager spatial response we begin with an analysis of the impact the discrete sampling grid has on the spectral transfer function (h_{det}) of the imager. From introductory digital signal processing theory, it is well known that the spatial limit of a discrete set of uniformly spaced samples of a waveform is given by the *Nyquist Frequency* (f_{nyq}) of the set:

$$f_{\mathrm{nyq}} = \frac{1}{T_{\mathrm{samp}}}, \tag{6.8}$$

where the sampling period (T_{samp}) is twice the distance between samples (ℓ_{pix}). At frequencies above this limit, the imager may have sensitivity, but the sampling grid is not capable of accurately representing this response. This occurs because

uniform sampling makes the Fourier Transform of the sampled data set periodic with period f_{nyq}, causing signal aliasing with spatial frequency beyond the Nyquist Limit into the sampled spectral range.

Whether or not aliasing is an issue depends on how responsive the imager is to spatial frequencies beyond the Nyquist Limit of the sampling grid. This, in turn, depends on the continuous response of the system (h_{det}) which may be written, in one dimension under the assumption that the system is separable, as

$$h_{\text{det}}(x) = h_{\text{pix}} \star h_{\text{dio}}(x), \tag{6.9}$$

where h_{dio} represents the impulse response of the diode detector layer and h_{pix} represents the binning effect of pixelation. To remove discrete sampling effects, h_{det} is defined continuously in x; effectively equivalent to the spatial response of a composite image built up by merging frames with detector translations so that there is always a pixel whose center is at x. In this case, h_{pix} will be

$$h_{\text{pix}}(x) = \Theta_{\frac{\ell_{\text{pix}}}{2}}(x), \tag{6.10}$$

where Θ is the boxcar function defined in section 4.2.2.2. From the convolution theorem of Fourier Analysis we know that the spectrum of h_{det} ($\widehat{h}_{\text{det}}$) is the product of the spectra of h_{pix} ($\widehat{h}_{\text{pix}}$) and h_{dio} ($\widehat{h}_{\text{dio}}$), *i.e.,*

$$\widehat{h}_{\text{det}}(f) = \widehat{h}_{\text{pix}}(f) \cdot \widehat{h}_{\text{dio}}(f). \tag{6.11}$$

The Fourier Transform of h_{pix} is

$$\widehat{h}_{\text{pix}}(f) = \frac{\sin(\pi \ell_{\text{pix}} f)}{\pi \ell_{\text{pix}} f}, \tag{6.12}$$

which has significant amplitude beyond the Nyquist Limit. The exact form of h_{dio} is generally quite complicated and thankfully unnecessary for the present work. Knowing that the spatial extent of this response (ℓ_{q}), effectively the spread of the

x-ray photocurrent within the diode, is much less than the dimensions of the pixel (ℓ_{pix}) tells us that

$$f_{\text{nyq}} = \frac{1}{2\ell_{\text{pix}}} < \frac{1}{\ell_{\text{pix}}} < \frac{1}{\ell_{\text{q}}}, \tag{6.13}$$

which implies that h_{dio} must have significant frequency response at least out to $\frac{1}{\ell_{\text{q}}}$. This lets us conclude that h_{det} will have significant frequency response beyond the Nyquist Limit of the sampling grid, and, thus, aliasing effects should not be neglected.

A second problem arises, at the other end of the spatial spectrum, when one is interested in obtaining a spatial frequency spectrum from the imager. If the spatial signal extends beyond the boundaries of the imager active area, acquiring an image effectively imposes a sampling window onto the data [78], [44]. Mathematically this sampling window is represented by the boxcar function (Θ_τ) mentioned earlier in this section. Generally the signal that is discretely sampled may be written as the product of a sampling window (w) and the distribution of x-ray flux (Φ). Again appealing to the convolution theorem we have that

$$\widehat{w \cdot \Phi} = \widehat{w} \star \widehat{\Phi}, \tag{6.14}$$

which shows us that windowing results in a convolution of the spectra we are interested in with the spectra of the sampling window. This convolution can cause the leakage of power within one spectral frequency into neighboring frequencies leading to distortion of the spectra of interest.

6.4.2 Data Collection and Refinement

We employ a method adapted from [81] to study the response of the Mixed–Mode PAD beyond the Nyquist Limit of its sampling grid. A series of images are taken in which the detector is translated in sub-pixel steps across features in the

illumination field over a distance greater than the width of a pixel. Data from these images are then combined, in two steps, to produce a map of pixel position vs detected intensity. The first step of this combination procedure is to divide each image into sets of pixels representing the same local illumination field, *e.g.* rows of pixels perpendicular to a knife edge or individual spot profiles from a pinhole mask, taking into account the translation offsets. In the second step, a reference is determined for each local image subset allowing these local images to be overlaid, producing a combined image of the local illumination field with finer sampling and good statistics.

The resulting single image has high spatial resolution but nonuniform sample spacing and noise with too high a spatial frequency to transform without aliasing. To remedy these problems, some of the spatial resolution was sacrificed to reduce noise through convolution with a filter, in the process resampling the data onto a uniform grid. Our choice of filter is governed by a desire for a broad spectral response, to minimize the distortion of the signal, with a sufficiently rapid fall off at and above the Nyquist Frequency to prevent aliasing. The filter used for this purpose is

$$h_{\text{flt}}(x) = \frac{1}{2\lambda} \exp\left\{-\frac{|x|}{\lambda}\right\}, \tag{6.15}$$

where λ is the filter parameter, generalizing to 2D by replacing x with the distance between the resampling point and the data point. The 1D spectral filtering this yields is

$$\widehat{h}_{\text{flt}}(f) = \frac{1}{1 + (2\pi\lambda f)^2}. \tag{6.16}$$

Practically, the filter convolution was carried out in Matlab by defining a new sampling grid and estimating the value of the integrated product of the filter, centered on that point of the grid, with the imager response through a discrete integral of the large number of samples in the combined image.

The final step of image refinement involves removing the spectral leakage effects associated with the limited imager active area. The typical technique for removing spectral leakage is to use a windowing function with more desirable characteristics [44]. A variety of alternative windows [44] are available each with merits in terms of minimizing leakage, amplitude fidelity, and frequency resolution. As we are not trying to isolate peaks, the frequency resolution of the window is not as important as the amplitude accuracy and minimizing spectral leakage, so the frequency analysis presented here employs the Blackman Window [44], which offers good response in these areas.

Throughout this work, Matlab was used as an analysis platform.

6.4.3 Spatial Characterization Measurements

To experimentally determine the imager's spatial response, a uniform x-ray source is occulted by various masks to produce an illumination field on the imager with known spatial and spectral characteristics. For these experiments, a flood field was generated by a Cu x-ray tube[3] operated at a 25 kV bias with a 0.4 mA tube current.[4] The occultation masks used were a commercial resolution target,[5] offering a knife edge mask and bar-target series, and a custom 50 μm tungsten sheet with a widely spaced grid of 75 μm holes. Images of the illumination fields were taken with the detector translated relative to the field by a high-precision translation stage.[6] Data sets were acquired, merged, and refined according to the method outlined in section 6.4.2.

[3]Model TFS-6050Cu tube with TCM-5000M power supply (Truefocus–Watsonville, CA).

[4]At these settings, the bremsstrahlung spectra will extend up to 25 keV, which will result in some flat-field transition. The effect of this is a very small, uniform increase in the background level across the chip. This effect may be easily corrected for during image processing by applying a global mean correction to the nonilluminated regions.

[5]Model 07-525 (Nuclear Associates division of Victoreen–Carle Place, NY), lead thickness 0.03 mm.

[6]Model ESP300 controller with ILS150 and ILS100 linear travel stages (Newport–Irvine CA).

In mask measurements such as these, it is important to consider the effects of a non-ideal flood field, resulting from an extended x-ray source and limited separation of this source from the pattern mask and imager, upon the resulting illumination field. A non-ideal flood field distorts the image measured by the detector in two ways: first, as discussed in section 2.3.1, inclined x-rays will produce a response in the imager with larger spatial extent than normally incident x-rays; second, the conjunction of the non-ideal field and the mask will result in parallax effects.

To evaluate the first effect, we begin by noting that all of the measurements made in this section were taken within 4 cm of the central axis of a 1 m collimator. As a result, no x-ray should have entered the detector diode at an incidence angle of greater than 2.3 degrees. Within the Cu tube spectra the signal from the K_α emission line dominates both the flux spectral distribution and charge yield profile, as discussed in section 2.3.1, measured by the Mixed–Mode PAD. Using the modeling method discussed in section 2.3.1, we can estimate the broadening of the Cu K_α peak (8.05 keV) resulting from this incidence angle to be ~ 1 μm, which is below the resolution of our experiment.

The effects of parallax may be estimated by considering the geometry of: the flood field, in terms of source size and collimator length; the mask, in terms of hole size and position, and the imager distance. To simplify this problem, one may reduce it to one dimension by considering the illumination field in any plane intersecting both the source and a hole in the pattern mask, as shown in figure 6.13.[7]

[7]This model makes some assumptions that will not hold for a real mask; specifically, it assumes ideal edges of the hole and a perfectly opaque mask material. If we were to introduce these additional elements, they would effect a small blurring of the edges of the d_O and d_I regions along with an overall background shift dependent on the mean transmission of the mask. Given the material and dimensions of the masks used for these measurements, the blurring of the edges of the d_O and d_I regions is expected to be negligible and we have already discussed how background shifts induced by uniform transmission through the mask may be corrected for. Thus, this model should suffice for the measurements presented here.

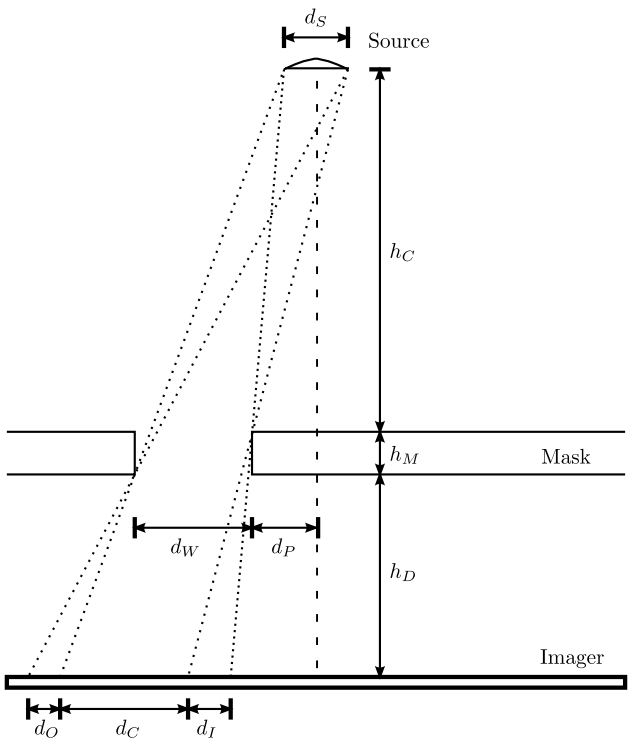

Figure 6.13: Parallax model used for the derivation of equations 6.17, 6.18, and 6.19.

From this, one sees that the masked image will have a central region (C) of full illumination as well as an inner and outer boundary regions of partial illumination (I and O resp.). The size of these regions is given by

$$d_O = \frac{h_D}{h_C + h_M}\, d_s, \tag{6.17}$$

$$\begin{aligned} d_C &= d_W + \frac{h_D}{h_C + h_M}\,(d_W - d_S) \\ &\qquad - \frac{h_M}{h_C + h_M}\left(d_P + \frac{d_S}{2}\right) - \frac{h_M(h_M + h_D)}{h_C(h_C + h_M)}\left(d_P + \frac{d_S}{2}\right) \\ &\cong d_W + \frac{h_D}{h_C + h_M}\,(d_W - d_S)\,, \end{aligned} \tag{6.18}$$

$$d_I = \frac{h_M + h_D}{h_C}\, d_S, \tag{6.19}$$

where the variables are defined by figure 6.13. It should be noted that these formulas do not provide the total parallax distortion from consideration of a single plane. A complete evaluation of the distortion would require an integral over all planes that intersect both the hole and the mask, which is a difficult task to parameterize much less undertake and, thus, not what these formulas are useful for. The utility of these formulas instead comes from their ability to set limits on the total parallax distortion based on the worst-case geometry of the hole.

For the flood field measurements presented here, the collimator length (h_C) was 1 m and the x-ray source size (d_S) was $\sim$ 180 μm. With each measurement presented in the following sections, we will use these values along with the relevant mask parameters and the formulas above to compute the maximum parallax distortion and thereby estimate the effect of parallax on the measurement.

6.4.3.1 Spatial Response Curves

For diffraction and other applications where the minimum feature size of interest is larger than twice the pixel spacing, so that one does not need to consider Nyquist Limit sampling effects, the real space response curves provide the most

useful metric of the imager's spatial resolution. In these response curves, the main consideration is the sensitivity of a pixel centered at one location within the array to x-rays incident on the imager at various relative positions. In an idealized PAD, 100% of the photo-current signal yielded by each x-ray would be detected by the pixel below, assuming the x-rays are incident from above, the point where the x-ray enters the diode. In a practical PAD, the variation in interaction depth and incidence angle, discussed in section 2.3.1, and diffusion of the charge as it travels through the detector diode layer, discussed in section 2.2.2, result in pixel sensitivity outside of geometric pixel boundaries. Measurement of two basic response curves, the pixel spot response and the edge-spread response, may be used to derive the fundamental spatial response of the Mixed–Mode PAD.

The purpose of the first measurement is to examine the structure of the ensemble-average pixel response to x-rays incident on the imager at a given relative location. The measurement is made by occulting a flood field with a mask containing an array of widely spaced holes, relative to the pixel width, to produce spots whose width (~ 75 μm) is less than half that of the pixel. The imager is then rastered across this illumination field in translations smaller than the spot width (25 μm) over an area that spans more than a full pixel (400 μm $\times$ 400 μm). An example of the spot pattern produced at a signal location is shown in panels (a) and (b) of figure 6.14. To evaluate the effects of parallax in this experiment, we note that the imager was well centered along the source axis, so that the entire image was within $\pm$1.4 cm of this axis, in both the horizontal and vertical. Using 2.0 cm as a limit on the maximal displacement of any spot (d_P), a imager to mask separation of 2.0 cm (d_D), and the other parameters from section 6.4.3, we find that $d_C = 71.4 - 1.5$ μm, where the 71.4 μm value represents the spot size reduction due to the source size and is common to all pixels within the array and 1.5 μm repre-

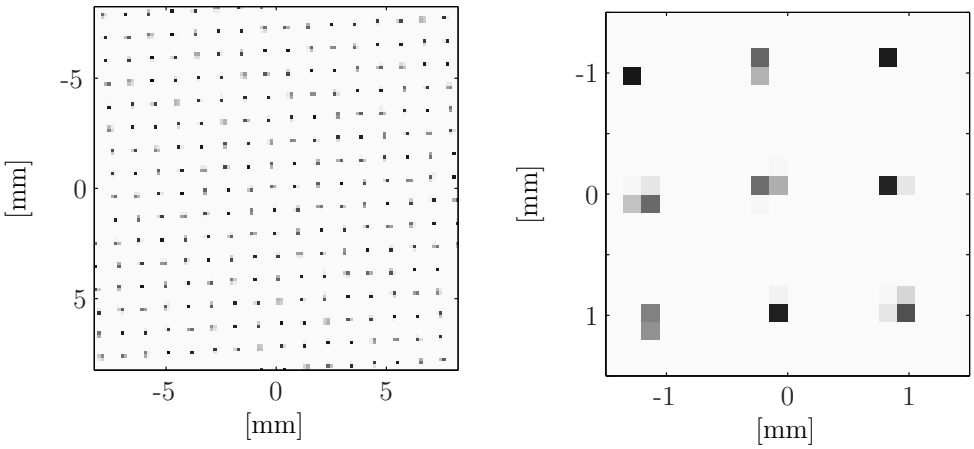

(a) Single Image (full)

(b) Single Image (zoom)

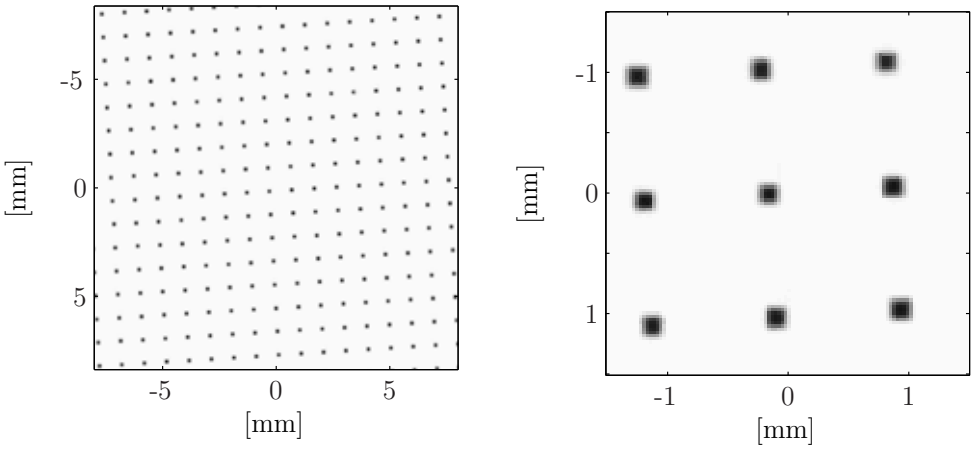

(c) Merged Image (full)

(d) Merged Image (zoom)

Figure 6.14: Spot pattern images taken with the Mixed-Mode PAD when illuminated with an x-ray flood field occulted by pin-hole mask. In these images, the holes on the 50 μm thick Tungsten mask are 75 μm in diameter arranged in a grid with 1 mm $\times$ 1 mm spacing. Panels (a) and (b) depict a single image while panels (c) and (d) represent a filtered combination of many images in which the detector was translated in sub-pixel steps relative to the pattern of spots.

sents the additional reduction due to the spots displacement from the normal axis. The other parallax parameters, d_O and d_I, are also constant across the array with value $d_O \cong 3.6\,[\mu\text{m}] \cong d_I$. This level of distortion is at the limit of the accuracy of the measurements presented in this section and, thus, may be neglected.

Using the method outlined in section 6.4.2, we combined these images to produce a composite image, shown in panels (c) and (d) of figure 6.14, in which individual spots may be easily located and analyzed to identify reference points. These reference points are then used to overlay the raw data from spots across the imager, creating an ensemble of responses. From this, the ensemble-mean response is extracted using the filtering operation described in section 6.4.2, producing the result in figure 6.15.

Because of the finite extent of the spots in the illumination field, this image does not directly represent the point source response of the pixel. Instead, the pixel point source response is blurred by convolution with the stimulating spot, producing the Pixel Spot Response (PSR). Methods exist to deconvolve the effects of this stimulus, either with Fourier Analysis techniques or direct deconvolution algorithms. However, both these approaches have limitations.

The Fourier Analysis method is particularly sensitive to the model used for the stimulating spot. To illustrate this point, consider the response of a pixel when the spot is moved across a line that bisects its vertical axis. This lets us represent the stimulating spot with only the horizontal location (x) as a variable. If we assume that the spot is a perfect circle of illumination of radius $R = 36$ μm, then the profile in x will be $\Phi(x) = 2\cos\left(\frac{\pi x}{R}\right)$, for $x \in [-R, R]$, and zero otherwise. This is equivalent to windowing the $2\cos\left(\frac{\pi x}{R}\right)$ response with a boxcar function, *i.e.* $\Phi(x) = \Theta_R(x) \cdot 2\cos\left(\frac{\pi x}{R}\right)$. Following similar arguments to those in the discussion of windowing presented in section 6.4.1, we see that the effect of

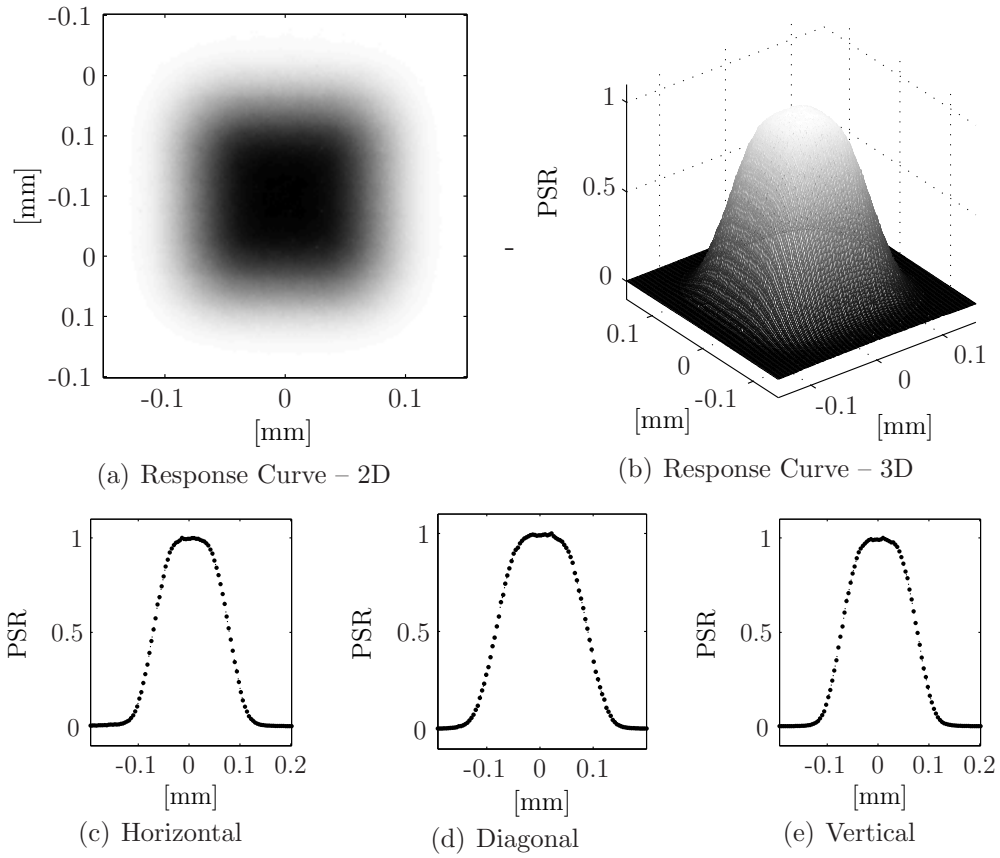

(a) Response Curve – 2D

(b) Response Curve – 3D

(c) Horizontal

(d) Diagonal

(e) Vertical

Figure 6.15: Pixel Spot Response (PSR) to illumination with a flood field occulted by a 75 μm circular aperture. The ordinate axis of both panels (a) and (b) are along the imager's row/column axes. Slice profiles taken horizontally (c), diagonally (d), and vertically (e) through the 2D response function illustrate the symmetries of the system.

this is

$$\begin{aligned}\widehat{\Phi}(\omega) &= \widehat{\Theta}_R \star 2\delta\left(\omega \pm \frac{\pi x}{R}\right) \\ &= \frac{\sin\left(\omega + \frac{\pi x}{R}\right)}{\left(\omega + \frac{\pi x}{R}\right)} + \frac{\sin\left(\omega - \frac{\pi x}{R}\right)}{\left(\omega - \frac{\pi x}{R}\right)},\end{aligned} \tag{6.20}$$

which is essentially the sinc response of the boxcar function's fourier transform, reproduced at two points via the convolution. An alternative windowing function (w), offering perhaps smoother response at $x = \pm R$, would still transform as $\widehat{\Phi}(\omega) = \widehat{w} \star 2\delta\left(\omega \pm \frac{\pi x}{R}\right)$. In this way, the transform of our choice of window becomes the dominant element of the deconvolution, making it difficult to produce an objective result.

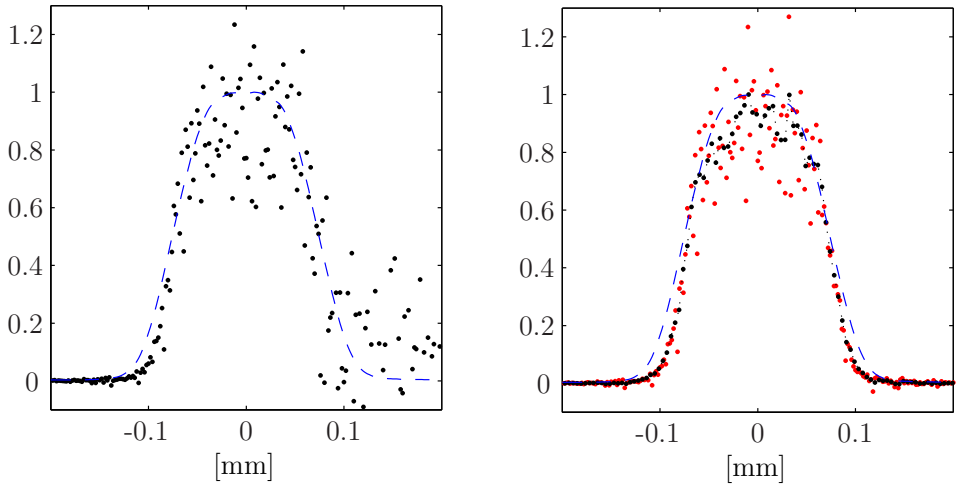

(a) Single Sided Deconvolution

(b) Double Sided Deconvolution

Figure 6.16: Discrete deconvolution of the pixel spot response. Panel (a) shows a deconvolution initiated at the left and carried out over the whole data set. The increasing fluctuations in the resulting deconvolution are due to noise amplification effects that result from the recursive form of the algorithm. Panel (b) shows the result of two half deconvolutions initiated from either side of the data set. This method gives an accurate representation of the extent of the pixel point source response; however, it still suffers from error amplification in its interior region. In both panels, a dashed line is included to indicate the pixel spot response.

Discrete deconvolution algorithms are less sensitive to the choice of illumination profile yet more sensitive to error within the data. This is a result of the recursive

form of standard discrete deconvolution algorithms. Generally if the convolution is defined by $\rho = h \star \Phi$, where ρ is the measured data and Φ is known, then the discrete deconvolution of h is given by the recursion relation

$$h(k) = \frac{\rho(k) - \sum_{l=1}^{N-1} \Phi(l) \cdot h(k+l-N)}{\Phi(N)}, \tag{6.21}$$

where $\rho(1)$ is the first non-zero sample of the measured response and $h(i) = 0$ for all $i \leq 0$. Because of this recursive definition, error introduced in early terms is propagated and magnified in later terms, as illustrated in panel (a) of figure 6.15. This problem may be partially overcome by merging deconvolutions begun at either end of the data set, as shown in panel (b) of figure 6.15. However, this only improved the data quality near the edge of the deconvolution, leaving the interior region contaminated with propagated error.

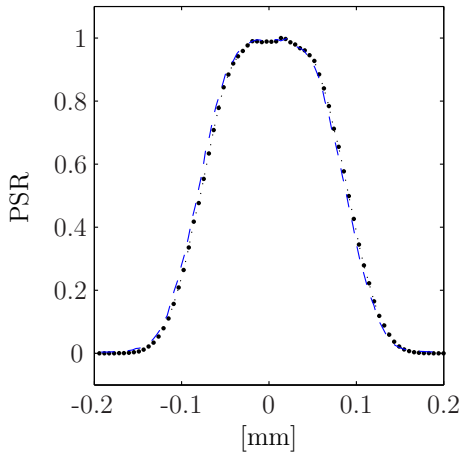

Figure 6.17: Comparison of the measured diagonal pixel spot response profile (dashed line) with the form calculated under the assumption of separability from the vertical and horizontal response profiles (dots).

Even without yielding a complete pixel point source response, the pixel spot response still tells us much about this function. For example, symmetries in a convolution reflect symmetries within the convolved functions; the left–right, horizontal–vertical symmetry exhibited in figure 6.15 indicate that the pixel point

source response also possess these same symmetries. Similarly, separability of the convolution will imply separability of the convolved functions. We test whether or not the pixel spot response is separable by comparing the product of the horizontal and vertical spot response profiles, shown in figure 6.15, with the diagonal profile. The result, shown in figure 6.17, very closely matches the direct measurement, indicating that the assumption of separability that we have been working under is valid. Noting the symmetries and separability of the pixel spot response allows us to simplify our problem from deconvolving the complicated 2D pixel point source response to retrieving this information directly from the much easier to measure Line Spread Response (LSR).

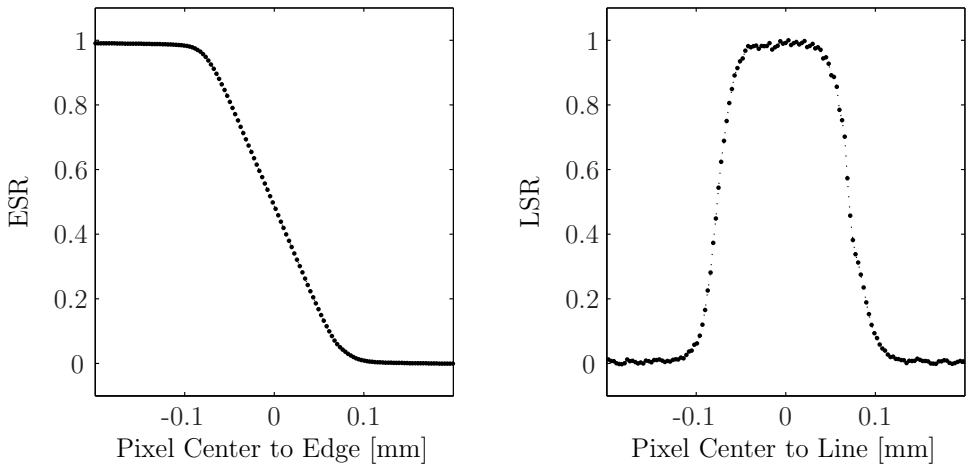

(a) Edge-spread Response (b) Line-spread Response

Figure 6.18: Linear averaged pixel response curves. The ordinate axis of both panels (a) and (b) are along the row/column axes and assume that the knife edge of the occultation mask is perpendicular to this axis.

The line spread response involves a second measurement where the imager is illuminated with a binary step illumination field using the knife-edge mask, discussed at the beginning of section 6.4.3. For this measurement, the knife edge was aligned to within 1 deg of the imager's vertical (column) axis. As translations are only necessary along the axis perpendicular to the knife edge, fine sampling trans-

lations (5 μm) over a large range (250 μm) are reasonable. As this measurement, in particular, is very sensitive to parallax effects an effort was made to get the imager as close as possible ($d_D \cong 1$ cm) to the mask. As a result, the interior source distortion (d_I) was only 1.8 μm, which is below the accuracy of this measurement. The other two parallax parameters (d_O and d_C) are not relevant, as the knife edge data only has one edge.

The imager response to these translated illumination fields is the imager Edge Spread Response (ESR), depicted in panel (a) of figure 6.18. Differentiation of the edge-spread response gives the imager LSR, effectively the imager's response to an infinitely thin line of illumination located at the edge of the knife-edge mask. From the symmetry and separability arguments presented earlier, we can conclude that the LSR is equivalent to one dimension of the pixel's point source response, a conclusion that is supported by comparing the LSR with the direct deconvolution of the pixel's point source response from figure 6.16. From the LSR, we find that the pixel's response, at a nominal 150 V diode bias, extends at the greater than 10% level nearly 22 μm and at the greater than 1% level nearly 45 μm.

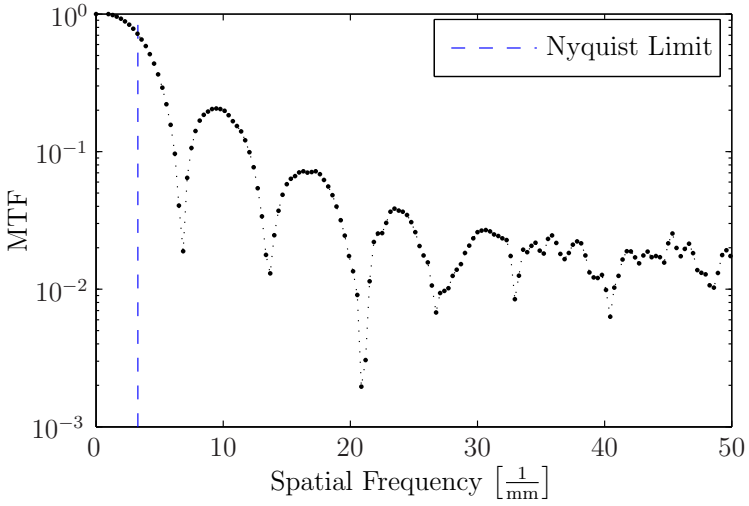

Figure 6.19: Modulation Transfer Function (MTF) calculated from the imager LSR.

6.4.3.2 Modulation Transfer Function

For systems with spatial features that extend beyond the imager's Nyquist limit, as can be found in x-ray tomography, a more complete description of the imager's spatial response is needed. This additional information is typically encapsulated in figure of merit known as the imager's Modulation Transfer Function (MTF). The MTF describes how the amplitude of individual spatial frequencies within an illumination field are altered, typically attenuated, by an imaging device. As remarked at the beginning of section 6.4.1, the MTF is poorly defined for discretely sampled systems where the extent of the imager's impulse response is comparable to or smaller than the spacing of the sampling grid. Thus, care needs to be taken in interpreting the results of this section. Within a single image, the Mixed–Mode PAD is only capable of providing spatial information up to its Nyquist Limit. The MTF presented here indicates how responsive the imager is to spatial frequencies beyond this limit. Within a single image, signals present at these higher frequencies are aliased back into the Nyquist Range, distorting the image. Only when the detector is translated relative to the illumination field so as to provide a finer sampling grid, as discussed in section 6.4.1, is this MTF response attainable.

The MTF is a derivative of a more general concept known as the Optical Transfer Function (OTF) that defines the spatial-filtering effect of an optical system on an illumination field's spectral distribution [74]. The OTF, though, is more of a theoretical tool than a measurable quantity as imaging systems are rarely able to capture the phase component of an illumination field. Yet, analytically it is a much easier starting point from which to begin an analysis of the Mixed–Mode PAD MTF. Recalling the discussion from section 6.4.1, the spatial transfer function of the Mixed–Mode PAD, *i.e.* the imager's pixel point source response, may

be written as

$$h_{\text{det}} = h_{\text{pix}} \star h_{\text{dio}}. \tag{6.22}$$

The OTF of the imager is then given by

$$\widehat{h}_{\text{det}} = \widehat{h}_{\text{pix}} \cdot \widehat{h}_{\text{dio}}. \tag{6.23}$$

The preceding section tells us the 1D form of h_{det} directly from the LSR measurement. Fourier transforming this returns the 1D OTF whose magnitude, shown in figure 6.19, is the imager MTF.

This spectral response has some features that are worth noting. To begin, as predicted by our discussion in section 6.4.1, the Mixed–Mode PAD MTF extends well beyond the imager's Nyquist Limit of $0.3^{-1}\frac{1}{\text{mm}}$. Within the MTF, periodic oscillations in intensity are evident. These oscillations are a direct result of the pixelation of the imager. As was discussed earlier, the OTF of the Mixed–Mode PAD may be written as $\widehat{h}_{\text{det}} = \widehat{h}_{\text{pix}} \cdot \widehat{h}_{\text{dio}}$, where in section 6.4.1 we derived the form of $\widehat{h}_{\text{pix}}$ to be a sinc function with period ℓ_{pix}. This response will have zeros at the spatial frequencies $f = \frac{n}{\ell_{\text{pix}}}$ for $n = 1, 2, \ldots$

6.4.3.3 Contrast Transfer Function

Due to the potential pitfalls in any discrete Fourier Analysis, it is useful to have an independent check of our MTF result. This can be provided by a measurement of the imager Contrast Transfer Function (CTF), which describes an imaging system's response to a series of at least three- or four-line targets of equal spacing and width, with binary transmission characteristics. The CTF measures the peak-to-trough intensity variation observed by the imager at the fundamental spatial frequency of the lines (s, in line pairs per mm), or explicitly

$$\text{CTF}(s) = \frac{T_{\max} - T_{\min}}{T_{\max} + T_{\min}}, \tag{6.24}$$

where $T_{\max}$ and $T_{\min}$ the maxima and minima of the detected amplitude, respectively.

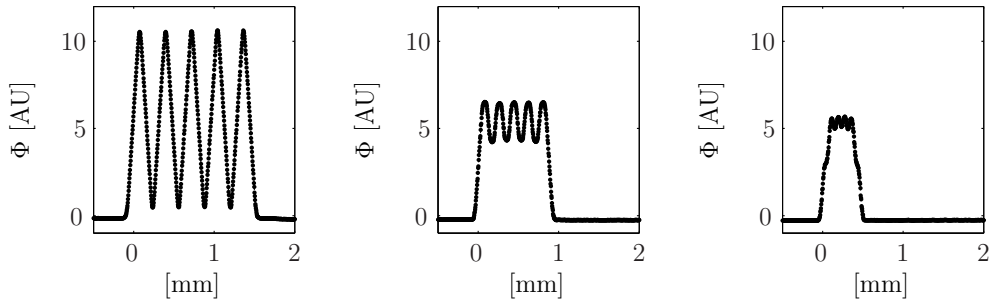

(a) 3.10 $\frac{\text{LP}}{\text{mm}}$ (b) 5.48 $\frac{\text{LP}}{\text{mm}}$ (c) 12.42 $\frac{\text{LP}}{\text{mm}}$

Figure 6.20: Real space Contrast Transfer Function (CTF) response at particular spatial frequencies.

For this measurement, the bar series on the commercial x-ray target discussed in section 6.4.3 was used. As with the knife-edge measurement, the detector was stepped along the axis perpendicular to the bar series in 5 μm translations over a full range of 250 μm. Parallax is not as significant an issue for bar series measurements, as it mainly distorts the higher-frequency harmonics instead of the fundamental harmonic of the series, unless one is interested in trying to reconstruct the MTF components from the CTF. Still, its effects are worth estimating. In this measurement, the farthest bar series was within 3.4 cm (d_P) of the normal axis with a width of 35 μm (d_W) and the mask to imager spacing was $\sim$ 1 cm (d_D). Based on this, the edge distortion was $d_O = 1.8\,[\mu\text{m}] = d_I$ and the size of the fully illuminated region was $d_C = 35 - 1.8 - 2.6$ μm, where the 1.8 μm term results from the source size and is applicable to all the bar series and the 2.6 μm term is due to the line displacement. These distortions are at the limit of the resolution of this experiment and, as such, will be neglected.

The CTF method is appealing in its simplicity as targets are easy to obtain and the data analysis is straightforward. It has the limitation, though, that it is not a direct measurement of the imager MTF. At a given fundamental frequency, the

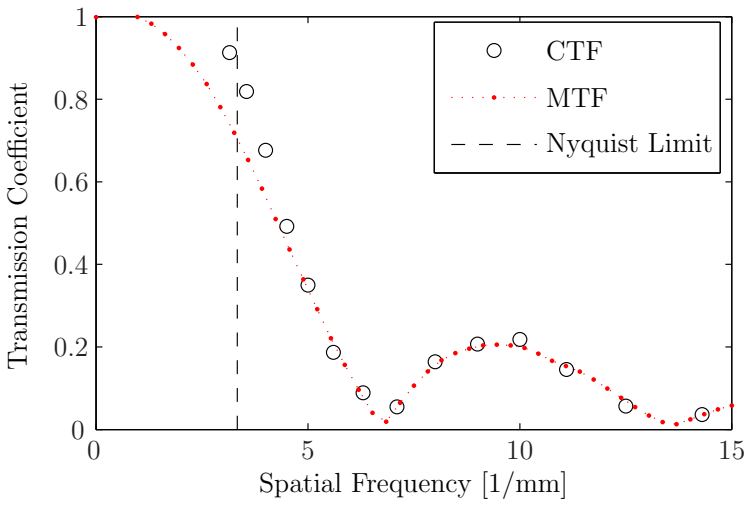

Figure 6.21: Comparison of the Mixed–Mode PAD Modulation Transfer Function (MTF), discrete Contrast Transfer Function (CTF) measurements, and the Nyquist Limit imposed by the imager sampling grid.

CTF response will be greater than the MTF response due to the additional harmonic components present in the binary-response spectra. However, as illustrated in figures 6.20 and 6.21, as the fundamental frequency of the lines is increased, higher-frequency harmonic components are suppressed, causing the CTF to converge to the MTF.

6.4.4 Spatial Response Inhomogeneity

Up to this point, our discussion has focused on the ensemble-mean pixel response. In the ideal, the data from a pixelated imager like the Mixed–Mode PAD should represent the signal acquired from a two dimensional array of equally-spaced and equally-sized pixels where each pixel responds according to this ensemble-mean. Practically, this ideal is never achieved in raw data due to distortion effects within the imaging system. For the Mixed–Mode PAD, image distortions are primarily the result of inhomogeneities within the detector diode combined with gain and

offset variations in the signal processing electronics.

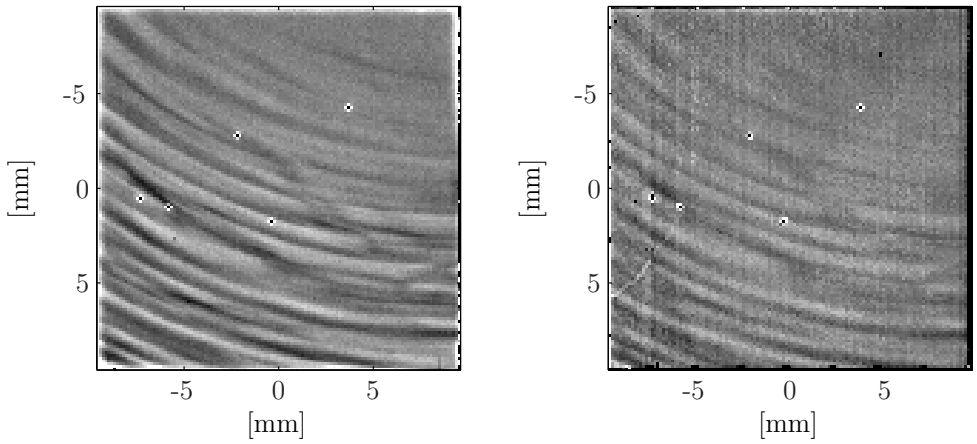

(a) Cu X-Ray Tube (b) Mo X-Ray Tube (Al attn)

Figure 6.22: Background subtracted and mean intensity normalized flat field response from Cu and Mo x-ray tube sources (25 keV and 30 keV tube bias resp.) as measured with the same Mixed–Mode PAD hybrid biased at 150 V. Both images are shown on a gray scale spanning ±10% of the mean intensity. A ∼1 m collimator separated the imager from the x-ray source. In addition, a 794 μm Al attenuator was used to suppress the bremsstrahlung component of the Mo spectra with the main effects evident below 10 keV.

To elucidate the image distortion created by the detector, flood field response measurements were taken using x-ray tubes with Copper (Cu) and Molybdenum (Mo) targets separated from the imager by a ∼1 m, air filled, collimator. This setup produces an x-ray flood field that is uniform to within a $\leq$ 0.3% linear gradient across the imager.[8] The operating conditions under which this data was taken utilized a tube current of 0.4 mA with a high voltage bias of 25 kV for the Cu tube and a tube current of 0.4 mA with a high voltage bias of 30 kV for the Mo tube. With these parameters, the Cu tube produces a spectra dominated by the 8.05 keV Cu K_α emission line, with a significant bremsstrahlung component. The Mo tube produces a spectra dominated by the 17.5 keV Mo K_α emission line, also with a significant bremsstrahlung component. To suppress the lower energy

[8]Determined by rotating the source 180 deg. and comparing the image difference.

element of the bremsstrahlung spectrum from the Mo illumination field, a 794 μm Al attenuator occulted the portion of the field on the imager. This effectively removes the majority of the bremsstrahlung radiation below 10 keV.

The response to a flat field illumination pattern, background subtracted and normalized to the mean pixel intensity, of both sources as measured by the same hybrid operated at the nominal detector diode bias of 150 V, is shown in figure 6.22. Clearly evident are arcs of intensity variation that are independent of relative detector to source position. As we will discuss later in this section, the magnitude of these intensity variations is strongly dependent on the bias voltage of the detector diode. However, from a figure 6.11, the accumulated signal integrated over the full array is essentially independent of the detector diode bias, allowing for a small decrease due to recombination resultant from longer drift times at lower bias voltages. This indicates that these arcs are due to redistribution of charge within the detector as opposed to local variations in the detector quantum efficiency or conversion gain per x-ray.

The explanation commonly given and accepted for these arcs is that they arise due to electric fields within the plane of the imager, caused by variations in the bulk doping profile. These variations are, in turn, a result of the process used to grow high resistivity silicon wafers [100]. Typically, the high-purity, high-resistivity wafers used for optoelectric detectors are fabricated in a float zone process. In this process, a rod of polycrystalline material is held above a crystal seed in the growing chamber. The bottom of the rod is then gradually heated to the point of melting, causing a 'float-zone' of molten silicon to form between the seed crystal and the polycrystalline rod. This zone gradually advances up the rod allowing a single crystal to grow behind. Dopants are typically introduced by allowing controlled amounts of gaseous dopants into the growth chamber, or less frequently through

neutron irradiation after growth [103, 80].

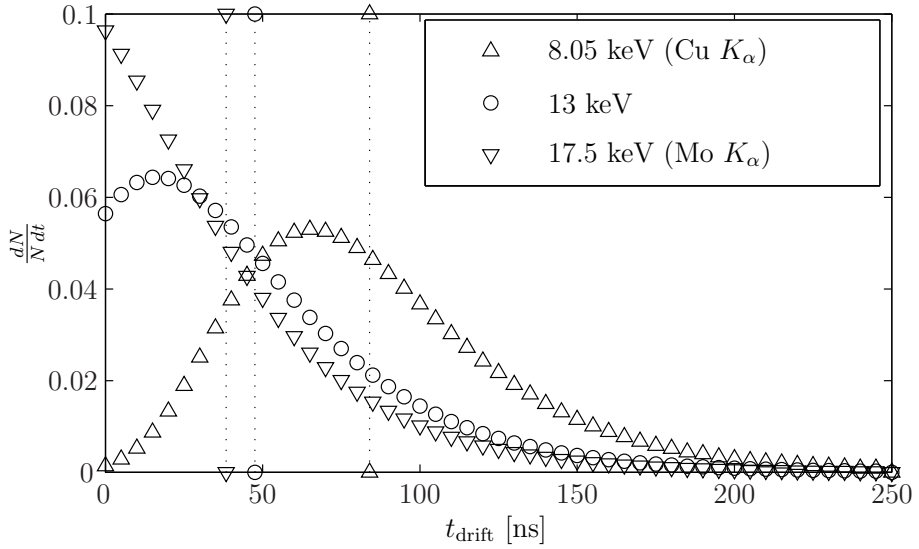

Figure 6.23: Calculated profile of the drift time (t_{drift}) for holes, in a 500 μm diode biased at 150 V, generated by normally incident x-ray beams of 8.05 keV, 13.0 keV, and 17.5 keV, based on equations 2.19 and 2.30. The dotted vertical lines denote the mean drift time for the curve denoted by their end points (38 ns, 48 ns, and 84 ns resp.).

To understand the effect of these lateral fields consider that the mean velocity ($\langle\vec{v}\rangle$) of a charge carrier (*e.g.* electron) moving in an electric field ($\vec{\mathcal{E}}$) decomposes simply as[9]

$$\langle\vec{v}\rangle = \mu_{\mathrm{e}}\vec{\mathcal{E}}: \qquad \langle v_x\rangle = \mu_{\mathrm{e}}\mathcal{E}_x, \qquad \langle v_y\rangle = \mu_{\mathrm{e}}\mathcal{E}_y, \qquad \langle v_z\rangle = \mu_{\mathrm{e}}\mathcal{E}_z, \tag{6.25}$$

where μ_{e} is the electron mobility (the equations are analogous for holes with the exception of a change in sign). The motion of carriers in the plane of the detector is separable from motion through the detector. The magnitude of the arcs then depends on the magnitude of the fields and the time spent under their influence. There is little that can be done to reduce the magnitude of these transverse fields

[9]This assumes that the induced velocity is below the velocity saturation limit. For detector diode bias voltages within our operating range of 0 to 200 V, this assumption should hold.

in commercial wafers—apart from selecting diodes, after hybridization, that exhibit minimal distortion. However, as alluded to earlier in this section, there are controllable parameters that affect the time spent drifting in these fields. From equation 2.19, the time spent drifting within a diode depletion region (t_{drift}) is dependent on the depth into the diode at which the x-ray is absorbed, the width of the diode, and the potential across the diode.

As discussed in section 2.3, the profile of absorption depths for a monochromatic beam of x-rays is strongly affected by the energy of the x-rays. As a consequence, the profile of drift times depends on the energy; however, this effect is complicated and nonlinear. In figure 6.22, there is little perceptible difference between the magnitude of intensity fluctuations in the two flat field images though the spectral composition and hence absorption profile of the two illumination patterns is quite different. The explanation for this is a combination of the compression effects of the diode's roughly linearly increasing, with depth into the diode, field strength on the drift times and the limited total thickness of the diode. As illustrated in figure 6.23, although the attenuation length of 8.05 keV x-rays (λ= 70.89 μm) and 17.5 keV x-rays (λ= 699.02 μm) differs by nearly an order of magnitude, the mean drift time differs by only 46 ns, roughly a factor of 2.2. Due to the nonlinearity of the dependence of drift time on x-ray energy, it is possible to create scenarios where the spectra of the illuminating field would strongly affect the magnitude of these intensity variations, *e.g.* very thick diodes, weak reverse bias, or extreme energy differences. However, within the operating range of the Mixed–Mode PAD under typical operating conditions, an energy dependence in these distortions will be difficult to detect.

A more dominant, though less tunable, factor in the magnitude of this distortion is the thickness of the diode layer. When manufacturing the diodes, it

is typically possible to select from a number of different wafer thicknesses.[10] At the same energy and depletion potential, substantially less distortion will occur in thinner diodes. Thus, in applications where the additional radiation tolerance of thicker diodes, to be discussed in section 6.7, is not needed, thinner diode layers are preferable.

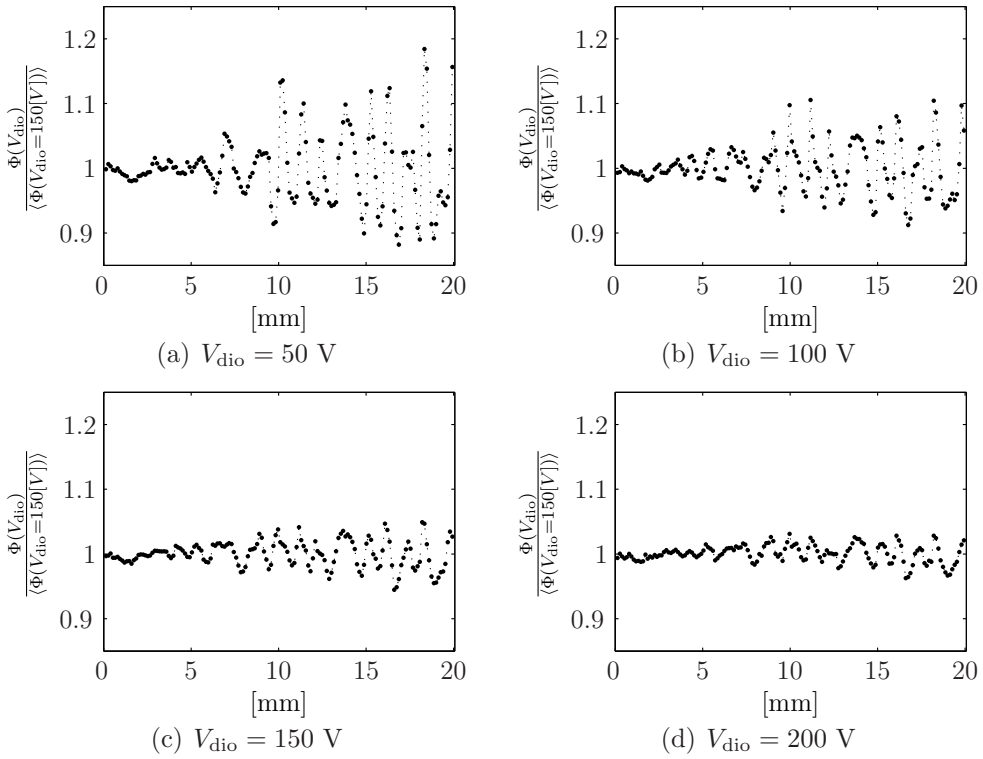

(a) $V_{\rm dio} = 50$ V (b) $V_{\rm dio} = 100$ V

(c) $V_{\rm dio} = 150$ V (d) $V_{\rm dio} = 200$ V

Figure 6.24: Intensity profile drawn across the same line on the same hybrid showing the variation in distortion with detector diode bias, normalized to the mean intensity at a bias of 150 V. The line shown here was chosen to be roughly normal to the arcs of intensity distortion.

The final external parameter affecting these drift times, and thereby the magnitude of distortions, is the potential across the diode. The impact of this parameter is illustrated in figure 6.24, which shows four plots of the same normalized line profile at differing detector diode bias voltages ($V_{\rm dio}$). While the impact of the diode bias setting can be strongly seen from these plots, the extent to which this

[10] Our diode manufacturer (SINTEF–Trondheim, Norway) offered a choice of 300 μm or 500 μm diodes.

can efficiently minimize image distortion is limited. As we recall from equation 2.19, the relationship between the bias voltage and drift time is inversely linear, *i.e.* doubling the diode bias reduces the drift times by a factor of two. Practically though, there are reasonable limits to how high one may safely increase the detector diode bias voltage[11] and thus the extent to which this may be used to suppress distortion.

As a consequence, some degree of distortion is nearly inevitable within the raw data. Section 6.6 discusses the implication of this distortion on the data as well analysis methods to reduce its impact.

6.5 Detector Quantum Efficiency

Having assessed the performance of the Mixed–Mode PAD at the pixel level and having considered its spatial response and spatial distortions, we are now in a position to evaluate the importance of calibration to the detector image quality. To do this, we employ the Detector Quantum Efficiency (DQE), a figure of merit commonly used to characterize x-ray imagers [41]. Defined as the squared ratio of the observed signal to noise over the ideal signal to noise of the source

$$\mathrm{DQE}\left(\int \Phi\right) = \frac{\left(\frac{\left\langle \int \Phi \right\rangle}{\sqrt{\mathrm{Var}\left(\int \Phi\right)}}\right)^2_{\mathrm{obs}}}{\left(\frac{\left\langle \int \Phi \right\rangle}{\sqrt{\mathrm{Var}\left(\int \Phi\right)}}\right)^2_{\mathrm{src}}}, \tag{6.26}$$

it can offer insight into the detector's ability to reproducibly record an x-ray signal in terms of the intrinsic fluctuations within the signal. An ideal detector, one that observes every x-ray produced by the source without contributing any noise to the measurement, would have a DQE of 1. The DQE is then reduced by the

[11] Although the breakdown voltage is silicon is high, $\sim 3 \times 10^5$ V/cm, one must also consider the attainable stability of the bias at high voltages, see discussion at the end of section 2.2.3.1, and the potential for arcing in a vacuum environment.

extent to which the detector is capable of absorbing the incident radiation, *i.e.* the Quantum Efficiency (QE) of the detector, and the extent to which it contributes to the variation within this measurement. By its very definition, the DQE is a complicated beast that effectively merges many elements of a detectors response into a single performance metric.

One most straightforward and telling methods for measuring the DQE of a detector involves taking a representative point of x-ray flux and repetitively sampling the signal observed from this spot at different locations on the detector (by translating the detector relative to the illumination field). The choice of spot size, however, has a nontrivial impact on the results of the DQE measurement due to the effects of discrete spatial sampling. When compared at equivalent dose per unit area, large spots, extending over multiple pixels, will exhibit higher DQE at a given dose level than small spots due to averaging effects that suppress pixel-to-pixel gain variations and the smaller percent variations in the number of pixels sampling a spot.[12] A sub-pixel spot will exhibit the worst possible DQE, at constant dose per unit area, as this case maximizes the effects of pixel to pixel gain variations while potentially dividing the x-ray signal between as few as one and as many as four pixels.

For the Mixed–Mode PAD DQE measurement a Mo x-ray tube source was used, operated with a high voltage bias of 30 kV and tube current of 0.4 mA. A 791 μm Al attenuator was also used to suppress the low-energy bremsstrahlung radiation of the source, leaving a spectrum dominated by the Mo K_α line at 17.5 keV (example spectra from this source are shown in section 7.1). After an air filled flight path of 1 m this flood field was occulted by a 50 μm thick tungsten sheet with a 1 mm $\times$ 1 mm grid of 75 μm diameter holes, resulting in an illumination

[12]Though, if the same comparison is made strictly in terms of total dose, small spots will perform better at low total dose levels, due to their smaller read noise contribution, though large spots will still outperform small ones once the read noise effects become negligible.

pattern on the detector comprised of sub-pixel spots over a very weak ($< 0.1\%$) flat background, due to transmission through the tungsten sheet. Exposure times were varied logarithmically from 1 ms to 100 s, with sets of 4 images taken at a total of 25 different locations of the detector relative to the spot field. These displacements were chosen at random, to remove systematic bias, from a flat distribution in a 2 mm $\times$ 2 mm region of the plane perpendicular to the x-ray path.

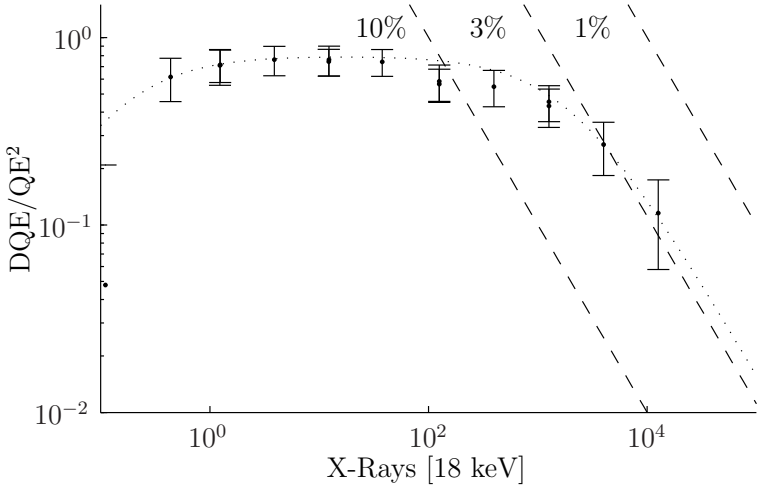

Figure 6.25: Detector Quantum Efficiency normalized to Quantum Efficiency of the detector. The error bars included with the data indicate the distribution of RMS computed from the individual illumination spots. Due to systematic variation between the different spots this RMS is much larger than the fluctuation between DQE/QE^2 measurements, as evident by the four repeated measurements. The dashed lines included on the plot represent curves of constant precision, indicating where fixed pattern noise is at a level that the precision of the measurement ceases to improve with dose. Curves of 10%, 3%, and 1% precision are shown.

To analyze the data, the set of four images at each camera position were used to identify unintended radiation events.[13] Following identification of a clean image, any residual global fluctuation in the level of the un-illuminated portion of the detector, *i.e.* the valleys between the peaks, was subtracted off, in accordance with the discussion of section 6.2. Within each combined image, illumination peaks were

[13] A.K.A. Zingers. These may be produced by terrestrial radiation, cosmic rays, etc.

identified and correlated with the position of the same illumination field with the camera at another location. The intensity of each peak was then calculated in each image. As the source itself exhibited non-negligible, non-Poisson fluctuations at the level of ∼5%, over the time periods considered, it was necessary to normalize the intensity of each peak by the total intensity of all peaks in each image. The RMS of the integrated intensity of each peak, computed over all locations of the peak from the same illumination spot, was then scaled by the average total intensity of all peaks over all images. This series of transformations minimizes the error introduced by the systematic source fluctuation. The average RMS variation of the peak-by-peak intensity was then taken as the observed intensity variation and the source variation was then approximated as monochromatic at 17.5 keV with Poisson statistics. While not completely accurate, this estimate is, at worst, a lower limit on the noise and thus a lower limit on the resultant DQE. The final curve, normalized by the quantum efficency of the detector, is shown in figure 6.25.

In interpreting the DQE curve it is useful to consider at a second order polynomial model of the observed noise,

$$\mathrm{Var}\left(\int \Phi\right) = a_0 + a_1 n + a_2 n^2, \tag{6.27}$$

where n is the mean number of x-rays emitted by the source. In this model, the 0^{th} order term (a_0) represents the read noise of the detector, the 1^{st} order term (a_1) represents the random noise of the detector and source, and the 2^{nd} order term represents the systematic noise of the detector. Under this model, a Poisson source will yield

$$\frac{\mathrm{DQE}}{\mathrm{QE}^2} = \frac{n}{a_0 + a_1 n + a_2 n^2}. \tag{6.28}$$

This model allows for up to three regions of accumulated dose [95]. In the first region, the noise of the detector is dominated by the fixed readout noise, with improving $\mathrm{DQE}/\mathrm{QE}^2$ with increasing dose. In the second region, the readout noise

becomes negligible compared to the combined random noise of the detector and source. Then, in the third and final region, the systematic noise of the detector becomes dominant leading to a fall of in the DQE/QE^2 response. Referring these regions to the normalized DQE curve shown in figure 6.25, we observe the initial readout noise dominated portion rising to the second region. Here, the fact that the detector's response peaks at slightly below one tells us that the electronic contribution to the variance in the observed signal is small in comparison with the intrinsic fluctuation in the x-ray signal, particularly given that we are likely underestimating our source fluctuations. Finally, we see the fall off introduced by systematic error in the measurement.

From the fall off introduced by the detector systematic error, we may estimate the attainable level of precision for an uncalibrated Mixed–Mode PAD. This is done by extrapolating the curve in the high-dose region to predict its limiting curve of constant precision. From the fit shown in figure 6.25, we find that the uncalibrated imager is limited by fixed-pattern noise to a precision of $\sim$2.5% for 17.5 keV x-rays. It should be noted, though, that the systematic gain is a function of the total charge measured, while the number of x-rays observed will ideally obey Poisson Statistics. Thus, a systematic limit of 2.5% at 17.5 keV corresponds a systematic limit of $\sim$1.9% at 10 keV. Improvement beyond this requires post-acquisition image correction, to suppress or remove these systematics, that will be addressed in the next section.

6.6 Detector Calibrations and Corrections

From the preceeding sections, an understanding of the characteristics of the Mixed–Mode PAD has been developed. For the detector to perform beyond this level requires post acquisition image processing based on careful calibration of the de-

tector. This section offers an outline and discussion of potential corrections that can be applied to Mixed–Mode PAD images and how the detector can be calibrated to measure these correction factors. However, as the Mixed–Mode PAD is the first example of a large-area pixel array detector utilizing an integrating front end and, therefore, this is the first attempt to calibrate such a device,[14] this should be seen as only a starting point for optimizing the detector rather than a complete recipe book. Future experience with the detector, through regular use, can be expected to refine and improve upon these techniques.

For most work, correcting the raw data for the Mixed-Mode PAD will entail:

- Merging the analog data (V_{res}) and digital data ($N_{\Delta Q}$) into a single measurement for each pixel, on each exposure, with an appropriate scaling factor (g_{dig}).
- Removal of pedestal offset variations (V_{ped}).
- Applying an absolute gain calibration (g_{abs}) to relate observed signal to deposited x-ray energy.
- Removal of detector-induced image distortion.

The first three elements relate to the total integrated signal observed by the detector through the relation

$$\int dt\, \Phi = g_{\mathrm{abs}} \cdot \left(g_{\mathrm{dig}} \cdot N_{\Delta Q} + V_{\mathrm{res}} - V_{\mathrm{ped}} \right), \tag{6.29}$$

where g_{abs}, g_{dig}, and V_{ped} are terms that must be individually calculated for each pixel. The remaining element, the distortion correction, is distinctly different and will be addressed separately.

[14] While integrating PADs bear some resemblance to their digital counterparts as well as to phosphor coupled CCD detectors the techniques used to calibrate these devices are difficult to apply to PADs for reasons that will be discussed.

In the following sections, each of these correction steps is taken up in detail, presenting both typical correction data along with a brief discussion of how this data is obtained.

6.6.1 Analog and Digital Data Combination.

The first calibration step to perform is to combine the analog and digital data from each pixel into a single measurement. This amounts to choosing a scaling factor, for each pixel, that converts the number of charge removal operations performed to an equivalent voltage shift in the analog output. In an ideal environment, this would simply be the difference of V_{ref} and V_{low}; however, due to contributions from differences between pixels and their respective analog output chains, there is systematic pixel-to-pixel variation in the analog equivalent voltage of each pixel's digital counts. Calibrating the individual contributions to this variation would be a difficult, time consuming, and tedious task. Instead, it is better to lump these effects into a direct measurement of the scaling factor (g_{dig}), a similar technique to that, explained in section 6.1, used to demonstrate the low-end linearity of the pixel response.

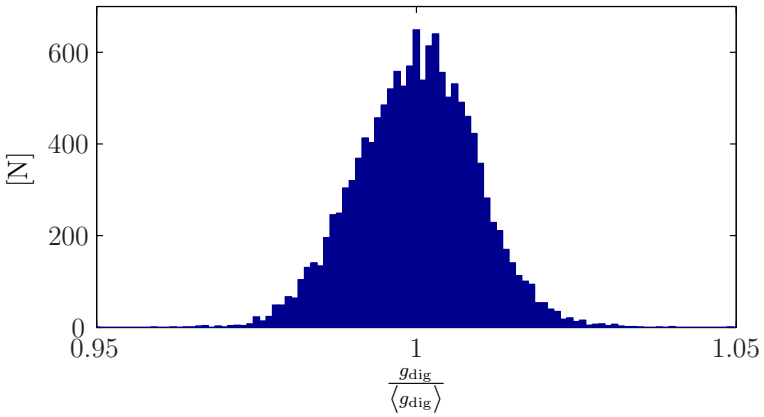

Figure 6.26: Normalized distribution of scaling factors for combining analog and digital data from the Mixed Mode PAD.

Explicitly, to perform this calibration a series of dark integrations are taken in which the exposure time is varied randomly to remove systematics due to low-frequency drift in the leakage current. For the data taken in section 6.1, the maximum duration of the dark exposure was chosen to be sufficiently long to allow for at least 10 charge removal operations in every pixel and the number of trials to be large enough that there was sufficient resolution to map out analog residual voltages in the range where no removals occur. Consequently, over 1000 separate integration periods were needed (with repetition for statistics) to clearly map out the slope of the analog and digital data, a quantity that is impractical for a standard calibration method. For standard calibration, instead of the flat distribution, we recommend the following a simple algorithm (algorithm 1) for generating a series of exposure periods, where t_{ana} is chosen to cover the span of the analog residual when no charge removals occur and t_{dig} is chosen to span at least 10 (though preferably more) digital charge removals:

Algorithm 1 Exposure Times for Mixed–Mode PAD Calibration

procedure ExpoTime($t_{\mathrm{ana}}, t_{\mathrm{dig}}$) ▷ Generate one exposure period.
    $r = \mathrm{random}([0, 1))$ ▷ Random number: $[0, 1)$, flat dist.
    if $r < 0.5$ **then**
        $t_{\mathrm{exp}} = t_{\mathrm{ana}} \cdot \mathrm{random}((0, 1]))$
    else
        $t_{\mathrm{exp}} = t_{\mathrm{dig}} \cdot \mathrm{random}((0, 1])$
    end if
    return t_{exp}
end procedure

With this simple algorithm, less than 100 exposures should provide good sampling of the slopes of the analog and digital data curves. As in section 6.1, the ratio of the slope of the analog data to the slope of the digital data yields the scaling factor (g_{dig}) for the pixel. Figure 6.26 shows the distribution of normalized scaling factors from one Mixed–Mode PAD detector hybrid. It should be noted that,

along with the pedestal offset, the scaling factor will vary over long, on the scale of weeks or months, time periods. Thus provisions should be made to recalibrate this portion of the detector on a regular basis.

6.6.2 Pedestal Offset

The pedestal offset (V_{ped}) is a shift in the measured signal that occurs independent of the duration of the integration. As with many of the other systematic errors, it is derived from many sources including clock feedthrough from transistor switches, such as the reset, CDS, and pixel sample and hold switches, as well as offsets and gain differences in the various amplification and buffer stages. There are two methods to easily remove the pedestal offset. If a difference of two images is taken, *e.g.* if a background is subtracted from the exposure, then the static variation within the pedestal will be removed. Where single images must be analyzed the constant term from the fit of the analog residuals preformed in section 6.6.1 may be used to remove the static pedestal. Figure 6.27 shows the distribution of pedestal offsets from one Mixed–Mode PAD detector hybrid.

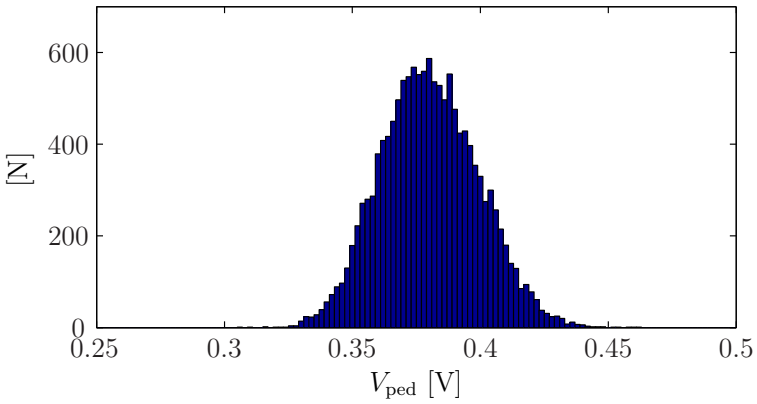

Figure 6.27: Distribution of pedestal offsets from one Mixed–Mode PAD detector hybrid.

6.6.3 Absolute Gain

The calibration techniques presented up to now focus on correcting individual pixels, independently. From the DQE work in section 6.5 this degree of correction is good to ∼2.5% (for a 17.5 keV source). To increase the detector's precision beyond this point, the level of fixed pattern noise must be diminished through calibration with an external source. In addition, for some experiments, we will need to be able to relate the data to an absolute quantity—the total x-ray energy deposited in the detector diode layer.

Unfortunately, measuring the absolute gain of a pixel is non-trivial because of the coupling between spatial distortions and gain variation in the total non-uniformity of the imager response. Therefore, in order to correct for these effects, they must first be decoupled and characterized independently. In devices like phosphor coupled CCD systems, the spatial distortions vary on long (*i.e.* mm) length scales relative to the pixel size. As a result, the canonical calibration method is to measure an image distortion map by employing a, relatively sparse, grid illumination pattern and computing the displacement of pattern elements relative to their expected positions [11]. Once the spatial distortions are corrected, the phosphor coupled CCD system may be gain calibrated through simple flat field exposures. In a PAD, however, the spatial distortion occurs through pixel-to-pixel variation in the total active area of each pixel, as discussed in section 6.4.4, on length scales near to or less than the pixel size; thus, the traditional method of utilizing a flat field illumination to determine relative, and sometime absolute, gains breaks down because there is no way to determine if excess signal seen in a pixel is due to a variation in the gain of the pixel or variation in total active area of the pixel.

It was proposed that, despite this, it should be possible to measure the gain

using a flat field illumination by taking multiple short frames in which the expected x-ray occupancy per pixel is less than one. Following the procedure outlined in section 7.1, the individual reading from each pixel could be histogrammed, with the idea that two peaks corresponding to zero and one would form. As we will show in section 7.1, the Mixed–Mode PAD is capable of this level of sensitivity; however, when the illumination field is broader than a single pixel, charge sharing between pixels becomes an issue. Studies into this issue indicated that, this is possible if a high energy, highly monochromatic source is available. However, this work also indicates that this method is incapable of yielding a calibration that would improve the detector beyond the 2.5% precision of the uncalibrated fixed pattern noise. To understand this, consider that the error in this two peak gain calibration is the sum of the error in the position of the zero and one x-ray peaks divided by the separation of the two peaks. As our observed, average conversion gain is only ∼0.7 mV/keV, even the Mo K_α line (17.5 keV) only yields ∼12 mV/x-ray. Consequently, calibration at the level of 2.5% would require peak measurement at an accuracy of 0.21 mV in each peak which would be challenging as our analog digitization is binned in 1 mV steps. Calibration to the precision of 0.25% desired of the final detector would require an accuracy of 0.021 mV in each peak, not strictly impossible, but very difficult with this technique.

A better alternative would be the construction of a dedicated calibration field. This setup would require a flood field of monochromatic x-rays occulted by a mask of sub-pixel-width, diameter holes spaced to match a multiple of the pixel grid, *e.g.* 25 μm holes on a grid of 750 μm × 750 μm centers. If the source is sufficiently monochromatic, then this setup should allow x-ray spectra to be recorded with multiple photon peaks. Recording spectra with multiple peaks dramatically reduces the accuracy required in determining the location of any given peak, making

spectral calibration much simpler. In section 7.1, this technique is demonstrated, although only a single illumination spot was used as a masks with an appropriate grid spacing was not available. In addition to this mask, for this setup to be practical, it would be require two linear as well as one angular degree of freedom between the mask and detector—for mask alignment and translation.

6.6.4 Distortion Correction

The image distortion introduced as the charge drifts through the diode layer is potentially a substantial source of measurement error. Spatial distortions and gain variations are closely coupled sources of fixed pattern noise in any imaging detector. For the Mixed–Mode PAD, detector calibration requires decoupling these two noise sources via independently measuring correction factors for one, either the spatial distortion or gain variations may be chosen, after which the other correction may be determined from flood field analysis. The canonical method of correcting phosphor coupled CCD imagers is to first measure the distortion correction, by illuminating the image with a known pattern of spots and mapping the displacement of these spots relative to their expected location [11]. The length scale on which distortions occur in these systems is sufficiently large that correction on the mm length scale is acceptable. This approach is not well suited for PAD detectors, however, as the image needs to be corrected on length scales that are close to the pixel dimensions. Because we have a very accurate means of measuring the absolute gain per pixel, as was discussed in the previous section, we may, instead, measure the gain variation independently of the spatial distortion and then use an analysis of the flood field response to correct for spatial distortion.

How to deal with this distortion, however, strongly depends on the type of measurement being performed. To illustrate this, consider a diffraction exper-

iment where one needs to measure the intensity of diffraction spots at various locations across the detector against a background of diffuse scattering. The inhomogeneities in the detector are unlikely to substantially distort the spots directly but they can have a significant effect on the background that can, in turn, distort the spot measurement. To see this, suppose that one were trying to measure a spot whose intensity was half the flat background level. Since spatial distortions introduce up to ∼3% variations in a flat field, the signal measured from this spot will systematically vary by 6%, depending on where it lands on the detector. Similarly, if the diffraction were twice the intensity of the background, the systematic variations from spatial distortions would be reduced to 1.5%. For this problem, the distorted field is effectively a background and so removing it through subtraction is preferable. Alternatively, if the data of interest extend continuously across the detector, as for example with x-ray tomography or ring diffraction, then a simple subtraction will likely not be sufficient, requiring instead a correction transformation. There are different choices for these, however, that themselves preserve and distort different aspects of the image as a consequence of their correction.

6.6.4.1 Image Correction Transforms

For an image with extended continuous regions of interest, the problem is more difficult, as one must contend with a distortion correction that will alter the distribution of signal between pixels. Because illumination fields and their resulting images are not isomorphic, it is not possible to construct a correction transform that preserves all aspects of the original image. Instead, one faces trade-offs between image aspects, such as the spatial resolution and the signal fidelity of the detector. In the remainder of this section, we will present two image correction transforms. The first, which relies on a normalization of the detector's quantum

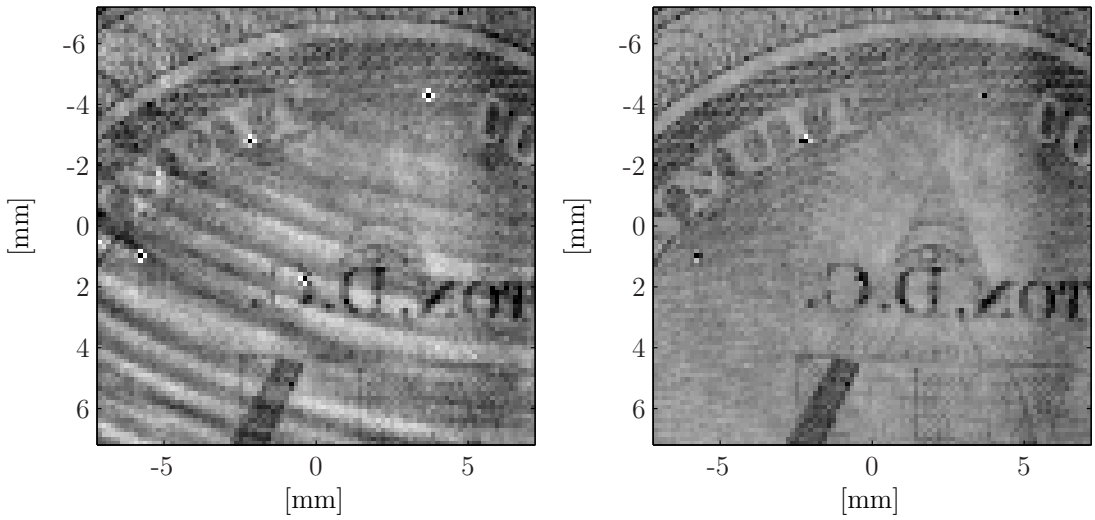

(a) Raw Image (b) QE Normalized

Figure 6.28: Tomography of a section of a $1 bill illustrating the quantum efficiency normalization distortion correction transform. Panel (a) shows the image original image, a merger of 10 100 second exposures using a Cu x-ray tube operated at 25 keV and 0.4 mA. Panel (b) shows the effect of applying the quantum efficiency normalization. There are a four dead pixels in the imager used to generate this image. The same data set was used to generate these image as was used to generate those shown in figure 6.29.

efficiency, effectively preserves the spatial resolution of the image at the expense of accurate measurement of local signals. The second, which uses a signal redistribution technique, accurately preserves the total signal observed, though, at the expense of some spatial resolution. Other correction transforms are certainly possible, but beyond the limited scope of this work.

The first correction transform that we will consider involves a quantum efficiency normalization. The idea behind this correction is that the differing total active areas among pixels may be viewed as differing quantum efficiencies, at least in the case of uniform illumination. So long as there is sufficient dose, this effect may then be removed by a simple rescaling; in essence, this means that measurements from pixels that collect from a large active area are suppressed while those that collect from a small active area are amplified, to the end that the detector yields a uniform response to uniform illumination.

To calculate this correction, one begins by calculating a correction map ($\mathfrak{C}$) from an ensemble mean, background subtracted flat field illumination ($\mathfrak{I}_{\text{flat}}$)[15] by way of

$$\mathfrak{C}(r,c) = \frac{\mathfrak{I}_{\text{flat}}(r,c)}{\langle \mathfrak{I}_{\text{flat}} \rangle}, \tag{6.30}$$

where r and c are row- and column-wise indices, respectively, of pixels in the image. The correction is then applied to a background subtracted, yet distorted, image ($\mathfrak{I}_{\text{dist}}$) through

$$\mathfrak{I}_{\text{qen}}(r,c) = \frac{\mathfrak{I}_{\text{dist}}(r,c)}{\mathfrak{C}(r,c)}, \tag{6.31}$$

to produce the final, quantum efficiency normalized, result ($\mathfrak{I}_{\text{qen}}$).

As figure 6.28 shows, this correction is an effective means to generate a nice picture, as it does a good job preserving the spatial resolution of the image. However,

[15]For all the correction transforms, the reference images need to have sufficient statistics to be limited by fixed-pattern noise. Otherwise, the error in the reference image will introduce unnecessary uncertainty into the distortion corrected image.

there are a few caveats that must be recognized with this transformation. The first is that the normalization only corrects the mean value of the illumination pattern, it does not correct the noise. Thus, pixels that collect from a larger total active area will show greater signal variation than pixels that collect from a smaller total active area, as a simple consequence of counting statistics, and as a consequence the level of signal in the pixel is no longer a reliable indicator of the uncertainty of the measurement. In addition to this, and perhaps more important, the quantum efficiency normalization does not ensure conservation of signal. As such, it can introduce systematic error into measurements, particularly those requiring accurate quantification of a signal deposited in an area spanned by only a few pixels.

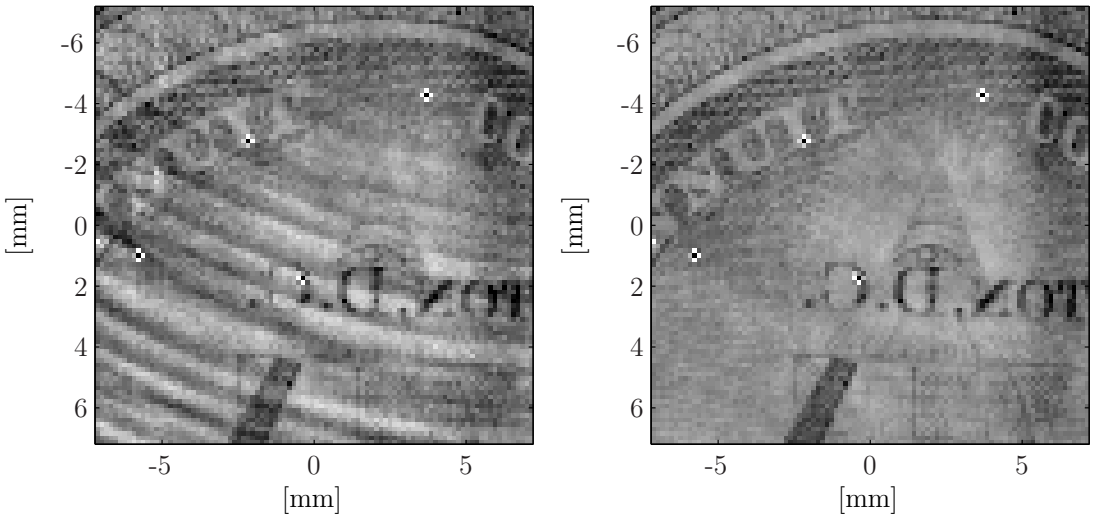

(a) Raw Image (b) AF Signal Shift

Figure 6.29: Tomography of a section of a $1 bill illustrating the charge-shifting adaptive filter correction transform. Panel (a) shows the image original image, a merger of 10 100 second exposures using a Cu x-ray tube operated at 25 keV and 0.4 mA. Panel (b) shows the effect of applying the adaptive filter correction. There are a four dead pixels in the imager used to generate this image, these pixels and their nearest neighbors are excluded from the adaptive filter. The same data set was used to generate these image as was used to generate those shown in figure 6.28.

An alternative distortion correction transform, that conserves the total signal

across the image, may be performed using signal redistribution. Here, the idea is, effectively, to rebin the image into an array where each element represents an equivalent total active area by shifting signal between nearest neighbor pixels. While this may sound simple, determining the appropriate amount of charge to shift and where to shift it is nontrivial. To accomplish this, we turn to a method developed in the field of system theory called *adaptive filtering* [105], where a recursive algorithm is used to determine a series of parameters (called weights) that control the correction transform. Based on these weights, a distorted image ($\mathfrak{I}_{\text{dist}}$) is transformed into an image where the distortion has, in principle, been removed ($\mathfrak{I}_{\text{af}}$).

For the algorithm developed for the Mixed–Mode PAD, these weights can be thought of as corresponding to extent to which one pixel extends beyond, or does not reach, its ideal boundary with its nearest neighbor. With one weight for each shared edge, this transform uses 128×127 (for row edges) and 127×128 (for column edges) arrays of weights to perform the requisite correction. In determining these weights, the distorted image from a flat field illumination ($\mathfrak{I}_{\text{flat}}$), after pixel gain calibration and background subtraction, is processed with an iterative training algorithm that uses the absolute value of the row- and column-wise gradient as an error measure to step the image towards an idealized flat response. In other words, we posit that the gradient of the corrected image should be zero and then recursively execute an adaptive loop on the training image ($\mathfrak{I}_{\text{flat}}$) that seeks to minimize this gradient by shifting charge between nearest neighbor pixels.

More rigorously, the training process of the adaptive filter is performed by letting $\mathfrak{I}_{\text{flat}} = \mathfrak{I}_0$ denote the initial flat-field image and $\mathfrak{I}_i(r, c)$ be the pixel value of the r^{th} row and c^{th} column element of the i^{th} recursive image. The update

algorithm is then[16]

$$\begin{aligned}\mathfrak{I}_{i+1}(r,c) = \mathfrak{I}_i(r,c) + q_{\text{stp}} \cdot \big[&\text{sgn}\left(\mathfrak{I}_i(r+1,c) - \mathfrak{I}_i(r,c)\right) \\ &+ \text{sgn}\left(\mathfrak{I}_i(r-1,c) - \mathfrak{I}_i(r,c)\right) \\ &+ \text{sgn}\left(\mathfrak{I}_i(r,c+1) - \mathfrak{I}_i(r,c)\right) \\ &+ \text{sgn}\left(\mathfrak{I}_i(r,c-1) - \mathfrak{I}_i(r,c)\right)\big], \end{aligned} \tag{6.32}$$

where sgn is the sign operator[17] and q_{stp} is a step size that determines the rate of conversion for the training algorithm as well as the resolution of the correction. Collecting the terms generated by the recursive algorithm we may define the weights after the n^{th} iteration of the adaptive filter as

$$a^n_{(r,c)} = q_{\text{stp}} \sum_{i=0}^{n} \text{sgn}\left(\mathfrak{I}_i(r+1,c) - \mathfrak{I}_i(r,c)\right), \tag{6.33}$$

$$b^n_{(r,c)} = q_{\text{stp}} \sum_{i=0}^{n} \text{sgn}\left(\mathfrak{I}_i(r,c+1) - \mathfrak{I}_i(r,c)\right), \tag{6.34}$$

where the subscripts a and b denote row- and column-wise weights. Then

$$\begin{aligned}\mathfrak{I}_{n+1}(r,c) = \mathfrak{I}_0(r,c) &+ a^n_{(r,c)} \cdot \left[\left(a^n_{(r,c)} > 0\right) ?\ (-1)\ :\ (+1)\right] \\ &+ a^n_{(r+1,c)} \cdot \left[\left(a^n_{(r+1,c)} > 0\right) ?\ (+1)\ :\ (-1)\right] \\ &+ b^n_{(r,c)} \cdot \left[\left(b^n_{(r,c)} > 0\right) ?\ (-1)\ :\ (+1)\right] \\ &+ b^n_{(r,c+1)} \cdot \left[\left(b^n_{(r,c+1)} > 0\right) ?\ (+1)\ :\ (-1)\right], \end{aligned} \tag{6.35}$$

where we make use of the ternary conditional operator[18] to simplify our notation. As n grows, the resultant image should converge, at a rate dependent on the size of q_{stp}, to a flat distribution to within $\pm 4 \cdot q_{\text{stp}}$ at each pixel, thus, in the following,

[16] In the remainder of this section we will be a little loose with our indices, to keep the description as succinct as possible, operating with the understanding that where an indexed element may not exist, *e.g.* at the edges of the the image, the indexing term will be taken as zero.

[17] Definition of the sign operator: $\text{sgn}(x) = \begin{cases} 0, & \text{if } x = 0, \\ \frac{x}{|x|}, & \text{otherwise.} \end{cases}$

[18] Definition of the ternary conditional operator: $(A)\ ?\ (B)\ :\ (C)$ equals B if A is true and C otherwise.

we will discard the superscript n with the understanding that $\lim_{n\to\infty} r^n = r$ and $\lim_{n\to\infty} c^n = c$.

To normalize these weights so that they may be applied to a general image, we divide each by the initial signal in the pixel which charge is shifted from,

$$\mathfrak{a}_{(r,c)} = \frac{a_{(r,c)}}{\big(a_{(r,c)} > 0\big)\ ?\ \big(\mathfrak{I}_0(r+1,c)\big)\ :\ \big(\mathfrak{I}_0(r,c)\big)}, \tag{6.36}$$

$$\mathfrak{b}_{(r,c)} = \frac{b_{(r,c)}}{\big(b_{(r,c)} > 0\big)\ ?\ \big(\mathfrak{I}_0(r,c+1)\big)\ :\ \big(\mathfrak{I}_0(r,c)\big)}. \tag{6.37}$$

Then the distortion corrected form ($\mathfrak{I}_{\text{af}}$) of a distorted image ($\mathfrak{I}_{\text{dist}}$) is given by

$$\begin{aligned}
\mathfrak{I}_{\text{af}}(r,c) = \mathfrak{I}(r,c)&\\
&+ \mathfrak{a}_{(r-1,c)} \cdot \big[\big(\mathfrak{a}_{(r-1,c)} > 0\big)\ ?\ \big(-\mathfrak{I}(r,c)\big)\ :\ \big(+\mathfrak{I}(r-1,c)\big)\big]\\
&+ \mathfrak{a}_{(r,c)} \cdot \big[\big(\mathfrak{a}_{(r,c)} > 0\big)\ ?\ \big(+\mathfrak{I}(r+1,c)\big)\ :\ \big(-\mathfrak{I}(r,c)\big)\big]\\
&+ \mathfrak{b}_{(r,c-1)} \cdot \big[\big(\mathfrak{b}_{(r,c-1)} > 0\big)\ ?\ \big(-\mathfrak{I}(r,c)\big)\ :\ \big(+\mathfrak{I}(r,c-1)\big)\big]\\
&+ \mathfrak{b}_{(r,c)} \cdot \big[\big(\mathfrak{b}_{(r,c)} > 0\big)\ ?\ \big(+\mathfrak{I}(r,c+1)\big)\ :\ \big(-\mathfrak{I}(r,c)\big)\big]\,.
\end{aligned} \tag{6.38}$$

An example of the results this transformation may produce is shown in figure 6.29, which utilizes the same data sets for the reference flat-field image and the image to be distortion corrected as was used to produce figure 6.28.

6.7 Radiation Tolerance

As was discussed in section 2.3, integrated circuits are susceptible to degradation when exposed to x-ray radiation. Because the signal processing electronics of the Mixed–Mode PAD pixel lie in the beam path, directly behind the detector diode, they will be subject to levels of radiation that will often exceed, during a single experiment, the lifetime dose of most terrestrial electronics, by many orders of magnitude. Consequently, the Mixed–Mode PAD was designed with

radiation tolerance in mind, as was detailed during our review of the pixel design in chapter 4. Yet, since the most effective means of radiation hardening involve component layout techniques that incur a substantial cost in terms of the circuit area they require, brute force hardening of the entire circuit through its layout was not an option. Instead, it was necessary to assess the impact of radiation damage on individual pixel components and, then, balance the need to harden these component with the constraints of the pixel area, developing a design that used a mixture of mitigation techniques to fit within the limited pixel area a circuit that was sufficiently tolerant of radiation damage for long-term use at synchrotron light sources.

This section presents an evaluation of the radiation tolerance of the Mixed–Mode PAD pixel based on laboratory measurements on the bare ASIC (*i.e.* no detector layer) and synchrotron measurements on the hybridized device. It should be understood that these measurements focus on the radiation tolerance of the pixels within the hybrid rather than the long term radiation hardness of the device as a whole. In particular, they do not evaluate the radiation tolerance of support structures on the periphery of the array. The reason for this is primarily pragmatic, as an evaluation of the radiation tolerance of these peripheral elements is much more involved and the fault mechanisms more global than the comparatively isolated situation of individual pixels within the ASIC. Furthermore, the radiation hardness of these elements is not as critical as that of the pixel as they can be shielded by building appropriate masks into the camera housing.

6.7.1 Comments on Units and Dose

Evaluating the radiation hardness of a detector such as the Mixed–Mode PAD can be confusing because of the nontrivial relation between the flux incident on the

imager and the total dose accumulated in radiation-sensitive areas of the device. Typically, when discussing radiation tolerance, one speaks in terms of the Total Ionizing Dose (TID), as defined by the energy absorbed per unit mass of the absorber,[19] absorbed by the detector or a portion there of. Most often the reference selected for these TID calculations is a portion of the detector that is particularly sensitive to radiation effects. The problem with this approach is that it does not make clear the relationship between this dose and the ultimately important scientific quantity, the integrated flux on the surface of the detector. To illustrate this point, consider that phosphor coupled CCD systems, discussed in chapter 2, are rarely investigated for their radiation tolerance, although CCDs are known to be susceptible to a variety of radiation damage mechanism, very similar to those that affect PADs [62]. The reason for this is that the phosphor and fiber optic effectively shield the CCD from damage by the x-ray beam, reducing the question of radiation tolerance in the CCD itself to an academic exercise. Similarly, in considering the radiation tolerance of a PAD detector it is not enough to only speak in terms of the effect that a given dose of radiation has on a region of the device. The relationship between this dose and the integrated flux on the surface of the detector must be clear for the dose to have meaning.

As discussed in section 2.3, the dominant effects of radiation damage within the ASIC layer of the hybrid are due to x-rays absorbed within the oxide layer. To relate this to experimental quantities, *i.e.* the scattered intensity in a diffraction peak, the shielding from the material within the diode layer must also be considered. As was also discussed in section 2.3, the absorbed fraction of x-rays follows an exponential decay with path length through a material. For the 500 μm Si

[19]Traditionally, the absorbed does has been quoted in the unit of rad; however, continuation of this practice is discouraged in favor of the SI unit, the Gray: 1 [Gy] = 1 [J] / 1 [kg] [98]. We will adopt the SI standard in this thesis. Fortunately, conversion between the Gray and rad is not difficult as 1 [Gy] = 100 [rad].

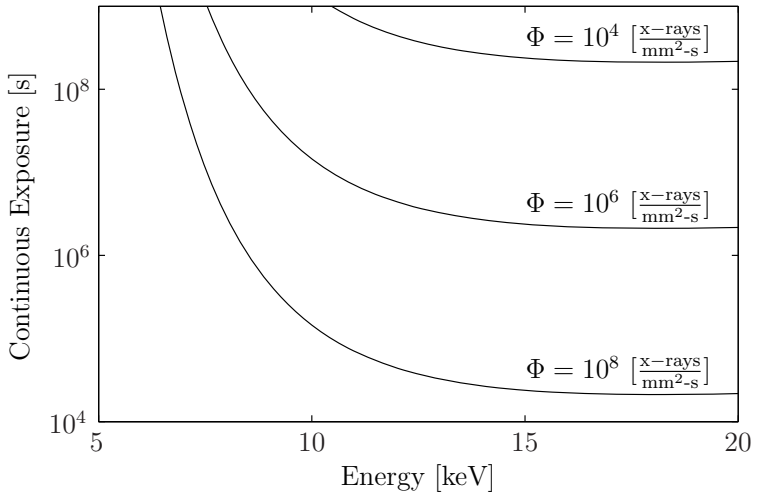

Figure 6.30: Estimation of the continuous exposure times required for a total dose of 1 kGy(SiO_2) for this Mixed–Mode PAD for three different flux densities incident on the detector. These times are as calculated based on equation 6.39 assuming that the flux density (Φ) incident on the detector is attenuated by 500 μm of silicon (the depth the Mixed–Mode PAD diode layer). Notably, these estimates do not include the additional protection the ASIC layer receives from the bump bonds.

detector diode used in the Mixed–Mode PAD, this means that the incident flux (Φ) is attenuated by a factor of $\exp\left(-\frac{500[\mu\text{m}]}{\lambda_{\text{Si}}(E_x)}\right)$, where $\lambda_{\text{Si}}(E_x)$ is the attenuation length in silicon at the x-ray energy E_x, as given by figure 2.5. With the attenuated flux after the diode (Φ_{dio}), we may estimate the TID in the SiO_2 layer of the ASIC.[20] Since the depth of the SiO_2 is small (5–7 nm for gate oxides and on the order of 100 nm for field oxides) relative to attenuation length of x-rays within the design range of the Mixed–Mode PAD (on the order of 10 to 100 μm), we may use the small x approximation of the exponential ($\exp(-x) \approx 1/x$) allowing us to remove the depth dependence from the absorbed dose through cancellation with the depth term in the mass factor. It is then possible to derive the TID in the

[20]This estimate offers an upper bound, as it neglects the absorption from the metal interconnect layers and bump bond that separate the diode from the Si surface of the ASIC. The metal interconnect layers are extremely thin and composed of weakly absorbing Al interconnects or SiO_2 passivation, so we may consider them to be effectively transparent. The bump bond offers significantly more protection, however, only to a limited portion of each pixel; hence, this calculation will express the TID expected in areas that it does not shield.

oxide independent of its thickness

$$\mathrm{TID}(\mathrm{SiO_2}) = \Phi_{\mathrm{dio}} \frac{E_x}{\lambda_{\mathrm{SiO_2}}(E_x) \cdot \rho_{\mathrm{SiO_2}}}, \tag{6.39}$$

where $\lambda_{\mathrm{SiO_2}}(E_x)$ is the energy dependent attenuation length and $\rho_{\mathrm{SiO_2}}$ is the density of SiO_2 ($\sim 2.2 \times 10^3$ kg/m^3). Combining these allows us to calculate a TID related to an integrated flux on the surface of the detector. This is expressed in figure 6.30, which shows, against axes of continuous exposure time and x-ray energy at different flux levels, the contours of 1 kGy(SiO_2) total absorbed dose.

6.7.2 Bare ASIC Damage

TID measurements were performed on bare, *i.e.* unhybridized, ASICs from the AE190 submission,[21] 16×128 pixel test chip, and AE207 submission, final 128×128 pixel chip.

For the AE190 ASIC, dosing was performed at room temperature on inactive devices that were removed from the dosing mount and tested at a probe station in stages as dosing progressed. The x-ray source used was an Enraf Nonius X-Ray Generator[22] located in the Gruner laboratory at Cornell University. This rotating anode source is equipped with a Cu target and was operated with electron beam settings of 40 kV and 60 mA. An Osmic confocal multi-layer mirror[23] was used to collimate and monochromize the beam at the Cu K_α line (8.05 keV). The flux was measured to be 4.9×10^7 x-rays/s/mm^2, corresponding to a dose rate, integrated over the beam area, of 0.216 Gy(SiO_2)/s. In these measurements, no signs of failures within digital pixel elements were seen up to 10 kGy(SiO_2) TID. Within the analog circuitry, a decrease in the number of digital counts for fixed

[21] Measurements on the AE190 were performed by Lucas Koerner and are summarized in [57].
[22] Model FR571 (Enraf Nonius/USA–Bohemia, NY).
[23] Model CMF15-165Cu8 (Osmic–Troy, MI).

integration time was observed within the dosed area. This may be attributed to a reduction in the current sourced by the test current source resulting from shifts in the threshold voltage of the current mirror transistor in this circuit induced by the ionizing damage (see section 4.3.1 for a description of this circuit). In addition, the one-shot[24] pulse duration and the voltage retention time of the sample and hold circuit were investigated. The former increased as damage accumulated, in accordance with expectations. The latter degraded as a result of increased leakage through the CMOS switch separating the sampling amplifier from the sampling capacitor (see section 4.3.3 for a description of this circuit). As a result of these measurements, in subsequent submissions, all CMOS switches connected to charge sensitive nodes were changed to utilize an enclosed layout architecture for their nMOS component.

To verify that the radiation tolerance observed in the AE190 devices extended to the final AE207 submission a series of TID measurements was performed on an unhybridized AE207 ASIC. Dosing was performed on active devices maintained at a temperature of -25 deg. C so that these measurements would be more representative of actual operating conditions. The x-ray source utilized was the Christine beamline in the Gruner laboratory at Cornell University, which is fed by a Rigaku rotating anode source.[25] This source supplies a Ni filtered Cu spectra dominated by the Cu K_α line at 8.05 keV with negligible bremsstrahlung and K_β components. The source was operated with beam settings of 40 kV and 50 mA. A small hole in a lead mask ($\sim$50 μm thick) was used to isolate the beam to a roughly 3×5 pixel region, providing an integrated flux of 21.6×10^6 x-rays/s for a dose rate of 0.278 Gy(SiO_2)/s.

[24]The predecessor of the gated oscillator from section 4.2.2.3. Operationally, this circuit is very similar to the gated oscillator except that it could initiate only one charge removal cycle per comparator cycle.

[25]Model 4151C6 (Rigaku/USA–Danvers, MA).

Unlike measurements on the AE190 submission, it was observed that ∼100 Gy(SiO_2) was sufficient to induce problems with the digital circuitry, evident initially in an inability to program the CSR in portions of the ASIC and ultimately, at the ∼1 kGy(SiO_2) level, in null digital data from the ASIC—evidence of a failure within the in-pixel counter. In unbonded detectors, there is no efficient method of reliably testing the front-end response without the CSR, thus it was not possible to evaluate the condition of the analog circuit while this digital failure persisted. Annealing the active ASIC overnight, within the camera housing under vacuum, through use of the thermoelectric as a heat source to raise the hybrid to an elevated temperatures of +35 deg. C to +42 deg. C removed the problems seen on the digital system. Subsequent repetitions of the dosing and annealing cycles yielded similar damage and recovery patterns.

Regarding the pixel analog front end, at intervals of ∼10 Gy(SiO_2) during the dosing and then again after low temperature annealing for ∼10–15 hrs, when it was once again possible to program the CSR, the analog front-end was examined for indications of damage to analog pixel elements. Apart from the expected decrease in the test current, corresponding to shifts in the threshold voltage of the test current source transistor, no degradation in other analog parameters, specifically the charge removal response of the amplifier and the sample and hold retention time, were detectable.

The initial component to fail, the CSR, was present with an identical layout in both the AE190 and AE207 devices. The absence of digital failure in the AE190 devices relative to the AE207 is postulated to be due to a combination of the lower temperatures ($\Delta T \approx -50$ deg. C) at which the AE207 devices were operated and the fact that the AE190 devices were not biased during operation. The relationship between ionizing dose, temperature, and the operating state of the device

is complicated, but not intractable. As discussed in section 2.3.2, the dominant long term effect of radiation damage is the trapping of holes in the passivation oxide. However, the situation is not as simple as the outline in that section might suggest. A complete discussion, reviewing the topic in full nuance, requires an extensive multi-volume series [9, 10]. Here, we will offer only a focused look into the topic of the transport, trapping, untrapping, and ultimate removal of holes within the oxide, intended to explain the discrepancies between the TID measurements on the AE190 and AE207 submissions as well as provide guidance for future radiation tolerance assessments of PAD devices.

A key point that needs to be appreciated, although it is glossed over in many discussions of oxide ionizing radiation damage, is that holes are not trapped at creation but rather have a low, and temperature dependent, mobility.[26] They may therefore drift or diffuse, depending on the the local electric field conditions, until they recombine, are caught at a trap site, or leave the oxide. In most cases, the absence of free carriers in the oxide means that when recombination occurs it is between initial electron–hole pairs created by the ionization event [16]. The likelihood of this is, however, strongly dependent on the presence and strength of electrical fields within the oxide [16]. As an illustrative example, studies done in the late 1980s on the TID effects of 10 keV x-rays reported the following empirical relationship between the oxide field strength ($\mathcal{E}_{\text{ox}}$) and the fractional hole yield (f_h) after short-time-scale recombination [23],

$$f_h(\mathcal{E}_{\text{ox}}) = \left(\frac{1.35}{\mathcal{E}_{\text{ox}} \left[\frac{\text{MV}}{\text{cm}} \right]} + 1 \right)^{-0.9} . \tag{6.40}$$

Thus, in an inactive device, where the oxide field is $\sim$0, the fractional hole yield will be practically zero. In comparison, a 0.5 V gate voltage applied across a 5 nm thick oxide, yielding oxide fields of 10 $\frac{\text{MV}}{\text{cm}}$, increases the fractional hole yield

[26] Typically holes in SiO_2 have a mobility of 1.6×10^{-5} $\text{cm}^2/\text{V} \cdot \text{s}$ at room temperature as compared with 20 $\text{cm}^2/\text{V} \cdot \text{s}$ for electrons [16].

to 89%. Consequently, much greater TID levels are required to induce equivalent levels of damage in inactive, relative to active, devices.

Compounding the problem for nMOS devices is the issue that, when the detector is operating, gate bias and channel current generate electrical fields in the oxide that direct the drift of holes towards the channel [72]. This situation presents three possibilities: either the hole will be trapped in a bulk trap, it will reach the Si/SiO_2 interface and become trapped in an interface trap, or it will enter the Si and recombine with a electron supplied by the channel.[27] As trap sites are typically due to oxide defects[28] the cross section for trapping in the bulk is strongly dependent on the quality of the oxide. However, even within very high quality oxides with low bulk trap densities a substantial concentration of traps will occur at the Si/SiO_2 interface as a result of mismatching lattice parameters in the two materials—which, incidentally, is also the location where trapped holes will have the greatest effect on the channel. Thus active nMOS devices exhibit dual effects related to electrical fields in their passivation, increased hole yield and a distribution of trapped holes biased towards the channel, that amplify the rate of device degradation with total ionizing dose.

Cooling the detector also contributes to long term TID effects in MOS devices. Initially this comes via decreased mobility in the charge carriers. To a limited extent cooling will suppress the ionization yield fraction by making recombination more likely. The significance of this is, however, highly dependent on the presence of electrical fields in the oxide, as outlined above. Generally more significant is that, with decreasing mobility, the the time a hole spends drifting through the SiO_2,

[27] One may wonder what happened to the accompanying electron. While it is possible that it could also be trapped in the oxide, the cross section for this is typically two to three orders of magnitude smaller than that for holes. As a result most treatments ignore electrons in the irradiation portion of the TID assessment [16].

[28] A veritable menagerie of oxide defects exist, many too many to go into here. Those interested in details on oxide defects are referred to [62].

before reaching the Si channel, is increased, thereby increasing the likelihood that the hole will become trapped.

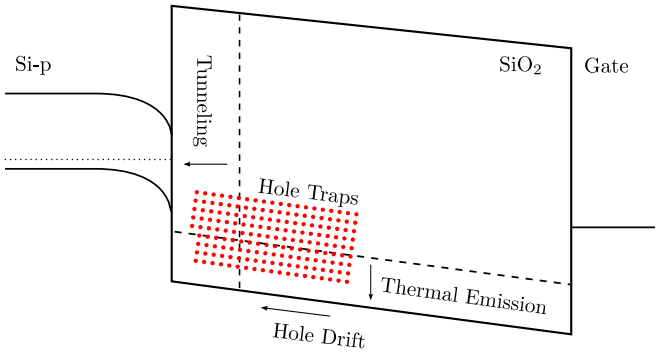

Figure 6.31: Illustration of TID recovery mechanisms for SiO_2 adjacent to the channel of an nMOS or parasitic nMOS device. The two dominant radiation damage recovery mechanisms are tunneling, in which holes tunnel directly through the SiO_2 into Si and as such is strongly dependent on the distance between the trap and the channel, and the thermal emission, in which holes are thermally emitted from low energy traps into the valance band of the SiO_2 and drift towards the channel under the influence of fields in the oxide (assuming an active device). Adapted from [68].

From another perspective cooling the detector increases the TID effects by depressing the rate at which holes are able to escape the oxide. Radiation damage recovery is typically modeled as having two primary components, illustrated in figure 6.31, one due to tunneling of trapped holes through the oxide and a second due to thermionic emission of holes from traps into the SiO_2 valance band and subsequent drift or diffusion to the channel or gate where recombination may occur. Tunneling through the oxide is strongly dependent on the separation of the trap from the channel and to a lesser degree the electric field in the oxide with only a very weak dependence on temperature [68]. Thermal emission, on the other hand, is strongly dependent on temperature. This has been found to be well modeled by an Arrhenius process with a second order pre-exponential factor; in

other words, a process with a rate constant ($R_{\rm thm}$) governed by an equation of the form

$$R_{\rm thm} \propto T^2 \exp\left\{-\frac{E_{\rm trp}}{kT}\right\}, \tag{6.41}$$

where $E_{\rm trp}$ is the energy barrier required to tunnel into the oxide [8]. Without detailed and difficult studies of the oxide to determine the distribution and energy levels of the trap states, it is not possible to predict the fractional importance of these effects, except perhaps in the limiting case of very thin gate oxides where one expects recombination through tunneling to dominate at reasonable temperatures. However, these considerations offer a compelling explanation of why digital failures were seen at the relatively low dose of 100 Gy(SiO_2) in a cooled and active ASIC while the same structures showed no damage when dosed up to 10 kGy(SiO_2) in an inactive, room temperature device.

Device recovery following annealing is expected based on detailed, transistor-level studies of TID effects on deep sub-micron technologies [3, 58, 59, 28]. These studies have shown that, at the transistor level, damage from TID levels of > 10 kGy(SiO_2) may be recovered from through annealing at temperatures of 100 deg. C over 10 to 15 hours. Our work indicates that the 100 Gy(SiO_2) damage threshold of the Mixed–Mode PAD may be recovered from on similar time scales and at lower temperatures ($\sim$40 deg. C) that are straightforward to obtain within the detector housing by using the camera thermoelectric as a heat source rather than a heat sink. The vacuum environment of the cryostat is advantageous in this regard as it prevents degradation of the oxide through hydrogen binding during the annealing process. Because of this, regular low temperature annealing is not expected to damage the detectors.

This work, however, indicates that a part of the regular operating procedure of a Mixed–Mode PAD camera should entail a period wherein the temperature of the

detector hybrids should be raised while the hybrid is active and time allowed to remove damage. If low-temperature annealed overnight, once a week, it is unlikely that damage will accumulate to appreciable levels. In the event of more significant damage, *e.g.* if a beam stop were to fall off and the detector were to be exposed to the main beam for an extended period of time, this too should be recoverable in a short period of time by further elevating the detector temperature (*e.g.* to 100 deg. C).

6.7.3 Hybridized ASIC Damage

While the Cu targets of the rotating anode sources within the Gruner lab are an effective means to deliver dose to the oxide of a bare ASIC, figure 6.30 reminds us how much longer their predominantly 8.05 keV beams, whose flux is on the order of 10^7 x-rays per second, will take to deliver a similar dose to the oxide of a hybridized detector. At a synchrotron light source, though, one is not restricted to characteristic emission lines and flux intensities are three to six orders of magnitude higher, offering a much more practical means to massively irradiate hybridized detectors.

Tests in April of 2007, at the CHESS F2 beamline, a description of which will be given in section 7.4.3.1, systematically examined dosing of a Mixed–Mode PAD hybrid by placing various regions of the imager into the unattenuated, monochromatic 13 keV main beam for extended periods of time. The particular hybrid chosen for this experiment was a defective device, due to unconnected bump bonds and scratches on the surface that generated excess leakage current, the most intriguing of which resembled the letters "W4."[29] The ASIC, however, was perfectly

[29]Early in the fabrication of the first large-area hybrids, one of our collaborators scribed wafer identifiers onto a few detector diodes in an attempt to make inventory management easier. Hence the 'W4' indicates that the detector layer for this hybrid came from wafer #4. Needless to say, this practice was halted shortly after the first hybrids were tested.

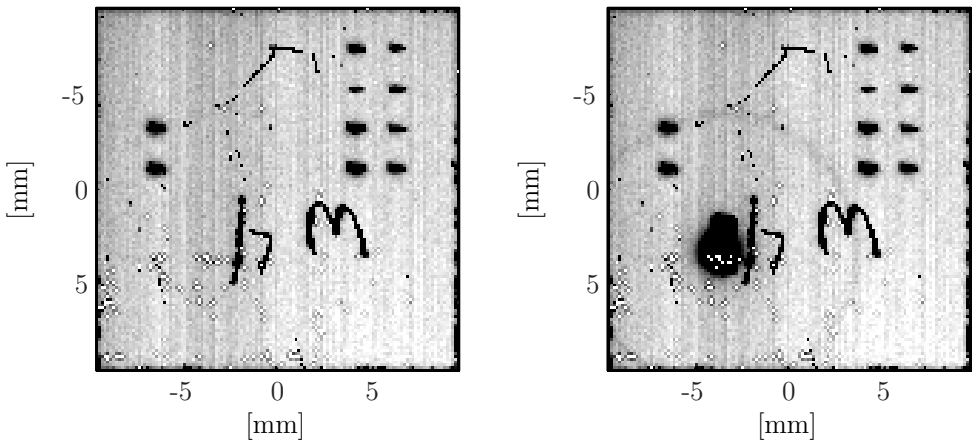

(a) Background

(b) Exposure

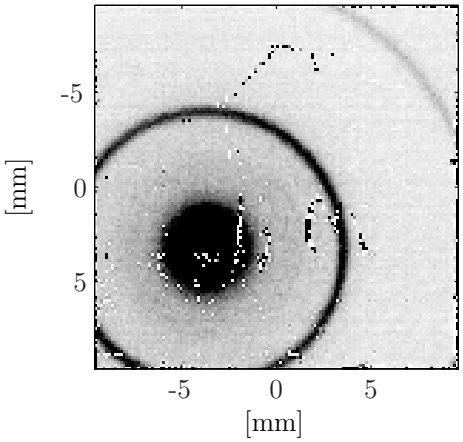

(c) Difference Image

Figure 6.32: Silver Behenate diffraction (neg.) from a hybrid used in the Synchrotron-based radiation-tolerance experiment. Panel (a) shows a combination of ten 1 s background images. Panel (b) shows a combination of ten 1 s exposures of a Silver Behenate powder sample, with no beam stop. The difference of panels (a) and (b) is shown in panel (c), where the intensity scale of the difference image is an order of magnitude smaller than that used in the exposure and background images. The radiation induced damage to the diode can be seen by the 10 large ($\sim$1 mm^2) spots of greater intensity in the background and exposure images, two in a column in the upper left quadrant of the image and eight in two columns of four spots in the upper right quadrant of the imager. From left to right, by column of damage locations, the exposure times were: (first column) 1440 s, 1920 s; (second column) 360 s, 30 s, 720 s, 960 s; and (third column) 60 s, 120 s, 240 s, 480 s.

functional and defects mentioned were isolated to specific regions of the hybrid. Using the defect free regions, a total of 10 doses were taken with exposure times of varying from 0.5 minutes to 32 minutes. The total flux on within the beam was monitored using the CHESS standard I_{short} ion chamber, placed immediately in front of the detector. As the beam profile and, hence, the dosing profile were not flat, images of the beam taken with a 2.3 mm Al absorber to estimate the beam profile.

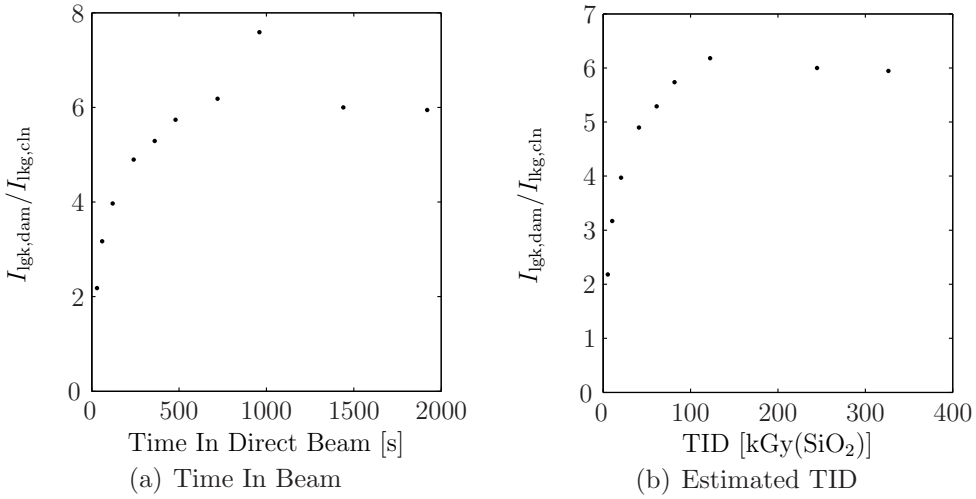

(a) Time In Beam (b) Estimated TID

Figure 6.33: Fractional leakage increase in primary beam region as a function of time in the main, uncollimated, F2 beam, panel (a), and estimated TID, panel (b). In the estimated TID plot, the point corresponding to 960 s was removed due to suspected beam fluctuations, as discussed in the text.

Analysis of the effects this dose has on the pixel is complicated by the fact that it is difficult to accurately estimate the total dose acquired by particular portions of the detector. This is a consequence of drift in the beam, both in terms of total intensity and spatial intensity distribution, on long time scales. The drift in beam intensity and profile were expected and measures were taken to mitigate it; specifically, the total intensity was monitored by the by chamber mentioned previously and images of the beam profile were taken between exposures.

Beam intensity fluctuations are somewhat mitigated by the design of the F2 beamline. Typically, the intensity of a synchrotron beamline decays over the duration of a fill (*i.e.* the period in which a single group of electrons or positrons is held in the storage ring) due to a gradual loss of beam current. The F2 beamline incorporates a feedback mechanism to minimize this effect, practically de-tuning the beamline at the start of the run and tuning it up as the run progresses, so that the beam intensity is maintained as constant as possible over the full duration of the run. As a result, the variations in beam intensity were less than 10% from the start to the end of this data set.

The impact of fluctuations in the beam profile was more insidious in this regard, as there is no easy way to monitor it while the dosing is occurring and the impact of this damage exhibits a nonlinear response. This point is illustrated by the aberrant 960 s data point in panel (a) of figure 6.33. The raw leakage in the peak region of this portion of the device was substantially higher than that exhibited by other elements, even those receiving significantly greater dose. As the damage profile of this spot also differs significantly from those of the other spots, it is likely that the beam was less stable during this exposure, providing greater damage to pixels neighboring the peak; a consequence of the nonlinear leakage increases with TID exhibited in figure 6.33.

Despite the increased leakage, the defect-free regions of the detector are usable for imaging. This is illustrated by the Silver Behenate powder diffraction pattern shown in panel (c) of figure 6.32. This data was taken on the Christine rotating anode beamline at the Gruner lab. For this measurement, the x-ray generator was operated at 40 kV and 50 mA. To produce this image, the sum of ten 1 s backgrounds, shown in panel (b) of figure 6.32 was subtracted from ten 1 s exposures, shown in panel (a) of figure 6.32, with the hybrid operated at -25 deg.

C. Despite the fact that the leakage was substantially stronger than the scattered x-ray signal, it has significantly cleaner statistics, as was discussed in section 6.1, that allowed it to be removed through background subtraction, without appreciably distorting the desired signal.

6.7.4 Conclusions on Radiation Tolerance

Based on our radiation tolerance assessments, the digital circuits in the Mixed–Mode PAD, specifically the CSR and the counter, show susceptibility to radiation damage near 100 Gy(SiO_2) TID. No degradation can be seen in the analog response up to this level, nor at any post-annealing level where the CSR functions.

The damage to the digital electronics is first apparent in the CSR due to the smaller transistor widths (at minimum length) used in its structures, compared to those in counter. By design, however, damage to this circuit does not degrade the performance of the pixel. This failure point is therefore quite useful as it allows the CSR to act as a "miner's canary" indicating, through an inability to program past the damaged area, when radiation damage has reached a level that the functionality of the pixel is itself at risk.

That said, there is no reason that this failure should ever be observed during normal operation of the Mixed–Mode PAD. This is because the level of radiation damage at which problems in the digital circuitry become apparent is quite high, with 100 Gy(SiO_2) representing nearly 83 continuous hours in a 13 keV beam with a flux of 10^6 x-rays/mm^2/s or more than 8,000 hours at the same flux in a 8.05 keV beam. This level of flux is much higher than what one expects from all but the brightest spots in a conventional scattering experiment. Consequently, one could expect months of operation before any failures would become evident on the digital system and therefore, a regular (weekly or biweekly), overnight, low-temperature

anneal should prevent the accumulation of damage from ever reaching a level where it may be observed.

6.8 Conclusions

Before attempting to use an imager like the Mixed–Mode PAD in scientific experiments, it is essential to gain a sufficient understanding of the device to confidently interpret the images it produces. In this chapter, we have delved into the details of the Mixed–Mode PADs performance, striving to offer a characterization of the detector that meets these ends.

Our investigations began with individual pixels, demonstrating their basic linearity, when the correct scaling between the analog and digital data is used, and was followed by a discussion of the read noise that showed that single x-ray sensitivity could be expected over a wide range of exposure times. From individual pixels, we moved to the imaging properties of the device as a whole, first considering the collection of charge from the diode layer and noting how the profile of measured signal was strongly dependent on the detector diode bias but the total signal collected was not. Our focus then moved to a careful analysis of the ensemble mean spatial response of the detector, both in the analog limit and considering the effects of pixelation on the device. From here, we concluded our investigation of the imaging properties of the detector with an examination of the causes and effects of spatial distortions within the device. Having looked at individual pixels as well as the imager as a whole, we used a measurement of the Detector Quantum Efficiency to predict the sensitivity of the uncalibrated imager. Here a conscientious effort was made not to make claims about attainable sensitivity, just what one could expect from a 0^{th} order device.[30] Instead, a discussion was presented

[30]This was done because measurements of the sensitivity of a tuned detector would only reflect

reviewing how the instrument can be calibrated and corrected to improve upon the sensitivity of the uncalibrated imager. Finally, the radiation tolerance of the device was assessed and found to be acceptable.

the effort put in to characterizing a particular device and as such would not be representative of what could be generally expected from the imager—though, there is a good chance that they would be enshrined as "typical performance" in future detector literature.

CHAPTER 7

FIRST MIXED–MODE PAD EXPERIMENTS

The first experiments with the Mixed Mode PAD that are presented here can be divided between those intended to demonstrate particular characteristics of the Mixed–Mode PAD and those intended to highlight fields where we believe the Mixed–Mode PAD has the potential make a significant scientific contribution. In the former category, we present three experiments that illustrate, respectively, the sensitivity, well depth, and spatial resolution of the Mixed–Mode PAD. In the later category, we strive to illustrate the scientific impact a detector like the Mixed–Mode PAD can be expected to have through actual experiments. To this end, studies were conducted of protein diffraction and self assembly of atomically thin monolayer films. This work is limited in scope by the relatively small active area of the single module prototype camera and the reasonable tenure of a graduate student. Still, it serves to illustrate the potential a full-sized, 512×512 pixel or larger for crystallography, or 128×512 pixel or larger for surface studies, has for these fields. Despite these limitations the prototype Mixed–Mode PAD camera was able to acquire data sets that would have been very difficult, if not impossible, to take with competing x-ray imaging technologies.

7.1 Spectral Lines

One of the touted performance characteristics of the Mixed–Mode PAD is its wide dynamic range, which we have claimed extends from single x-ray sensitivity, at least for 10 keV x-rays, to a full well depth of more than 2×10^7 10 keV x-rays. To demonstrate the capacity of the Mixed–Mode PAD to measure very weak x-ray signals, we performed a series of experiments where a single pixel was illuminated with x-rays, predominantly of one characteristic energy, and then taking multiple

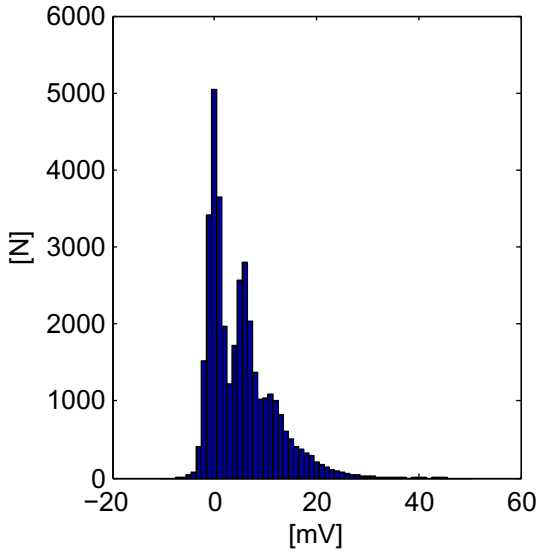

Figure 7.1: Observed spectrum from 1 ms exposures of a single pixel within the Mixed–Mode PAD, operating at -35 deg. C, illuminated by an unfiltered Cu x-ray tube operated at a bias of 25 kV and 0.4 mA of tube current. A 75 μm pinhole mask was used to isolate the x-ray beam to the interior of the pixel so that charge sharing effects were negligible. The spectrum of the Cu source will be dominated by the Cu K_α characteristic emission line. However, there will also be a significant bremsstrahlung component extending up to the tube bias voltage of 25 keV. Because of this, it is very difficult to distinguish quantized K_α peaks beyond 0, 1, and 2 x-rays.

short exposures to build a histogram of the distribution of x-ray induced signal, *i.e.* an observed spectrum.[1] The goal in these experiments was to show that the sensitivity of the Mixed–Mode PAD made it possible to observe quantized aspects of predominantly monochromatic source spectra.

To see how this works, note that if the source spectrum is given by $\Phi(E_x)$, where $\Phi(E_x)$ is the x-ray flux at energy E_x, then the likelihood ($P(E_x, n)$) of observing n x-rays of energy E_x within an exposure of duration $t_{\rm exp}$, assuming a Poisson source, is given by

$$P(E_x, n) = \frac{(\Phi(E_x) \cdot t_{\rm exp})^n \exp\left(-\Phi(E_x) \cdot t_{\rm exp}\right)}{n!}. \tag{7.1}$$

[1]Note that, since we are using an integrating device, the observed spectrum for an integration of a given duration is not equivalent to the source spectrum but a derivative thereof.

Now, suppose that a large number of constant duration exposures are taken. Each exposure will contain a number of x-rays of varying energy. We are interested in the probability that the sum of the energy of these x-rays ($P(E_{\text{dep}})$) will be a given energy (E_{dep}), as knowing $P(E_{\text{dep}})$ for all possible E_{dep} is equivalent to knowing our observed spectrum. If we let $\{E_x, n\}$ describe the results of a given measurement, where one reads this as the set of x-ray energies (E_x) within the source spectrum and their respective observed occupancies (n) during a particular measurement, then the probability of this result occurring is given by

$$P(\{E_x, n\}) = \prod_{\{E_x,n\}} P(E_x, n). \tag{7.2}$$

The total energy deposited in the detector by $\{E_x, n\}$ is then

$$E_{\text{dep}}(\{E_x, n\}) = \sum_{\{E_x,n\}} n \cdot E_x. \tag{7.3}$$

This, then, gives us the machinery to describe the spectrum of results that we expect to measure from a pixel within the Mixed-Mode PAD as

$$P(E_{\text{dep}}) = \sum_{(\{E_x,n\}|E_{\text{dep}}(\{E_x,n\})=E_{\text{dep}})} P(\{E_x, n\}). \tag{7.4}$$

The complexity of this final form is indicative of how difficult it can be to observe aspects of the source spectrum by looking at integrated quantities rather than individual x-rays, particularly when the source contains a complicated spectrum or there are effects like charge sharing at pixel boundaries. However, as figure 7.1 illustrates, even with the relatively complicated spectra from an unfiltered Cu x-ray tube,[2] it is possible to distinguish elements of the source spectrum when t_{exp} is sufficiently short. With simpler source spectrum, *e.g.* those with only a single characteristic line, limited background radiation, and where care is taken to mitigate charge sharing effects, still more quantized aspects of the source spectrum may be observed.

[2]Model TFS-6050Cu with power supply TCM-5000M (Trufocus–Watsonville, CA).

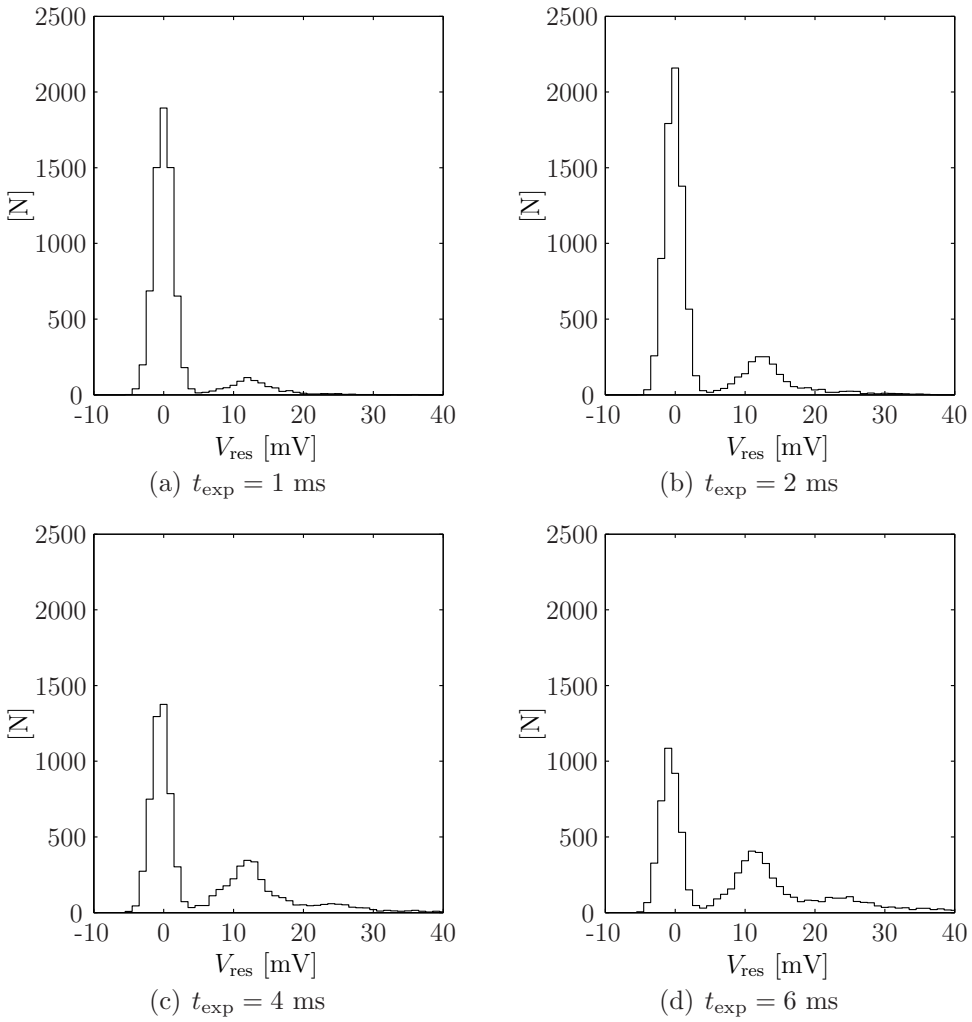

(a) $t_{\text{exp}} = 1$ ms (b) $t_{\text{exp}} = 2$ ms

(c) $t_{\text{exp}} = 4$ ms (d) $t_{\text{exp}} = 6$ ms

Figure 7.2: Spectra of the acquired signal from a series of short exposures with Molybdenum x-ray tube, operated at 30 kV with 0.4 mA tube current and attenuated by a 791 μm Al absorber, collimated with a 75 μm pinhole mask to restrict the beam to the interior of a single pixel, thereby eliminating charge sharing effects.

The first real attempt to measure multiple quantized x-ray peaks utilized a Mo x-ray tube[3] operated at a bias of 30 kV with 0.4 mA tube current. A 791 μm Al absorber was used to remove low-energy bremsstrahlung radiation, leaving a spectrum dominated by the Mo K_α characteristic emission line at 17.5 keV. With the detector at -35 deg. C, short integrations were used to limit the number of x-rays observed to 0, 1, 2, or (rarely) 3. To eliminate charge sharing effects, a 75 μm mask was used to isolate the incident flux onto a single pixel within the array. A large number, $\sim$10,000, exposures were then taken in this manner to produce the four spectra shown in figure 7.2. These spectra show clear peaks at 0 keV, 17.5 keV, and suggest peaks at 35 keV and 52.5 keV. Due to the presence of other x-ray energies within beam, specifically bremsstrahlung radiation above 10 keV that was not effectively removed by the Al absorber, the high energy peaks are notably blurred.

To improve this measurement, an essentially monochromatic source was employed. This source was derived from a Cu rotating anode[4] operated at 40 kV and 60 mA with multilayer optic[5] to isolate the Cu K_α line at 8.05 keV. In addition, a 50 μm thick Ni filter was used to reduce the intensity of the monochromatic beam and physically support a 50 μm thick Tungsten (W) mask with a 1 mm $\times$ 1 mm grid of 25 μm holes. By aligning a hole illuminated by the x-ray beam with a pixel in the imager, it was possible to completely isolate the x-ray signal to that single pixel. For a monochromatic source with no charge sharing, equation 7.4 then reduces to the basic Poisson distribution of equation 7.1, with the caveat that each peak is broadened by the read noise of the detector. In this experiment, it was then possible to observe a large number of x-rays without substantial blurring of the quantized x-ray peaks, as demonstrated by panel (a) of figure 7.3. Panel (b) of the

[3]Model TFS-6050Mo with power supply TCM-5000M (Trufocus–Watsonville, CA).
[4]Model FR571 (Enraf Nonius/USA–Bohemia, NY).
[5]Model CMF15-165Cu8 (Osmic–Troy, MI).

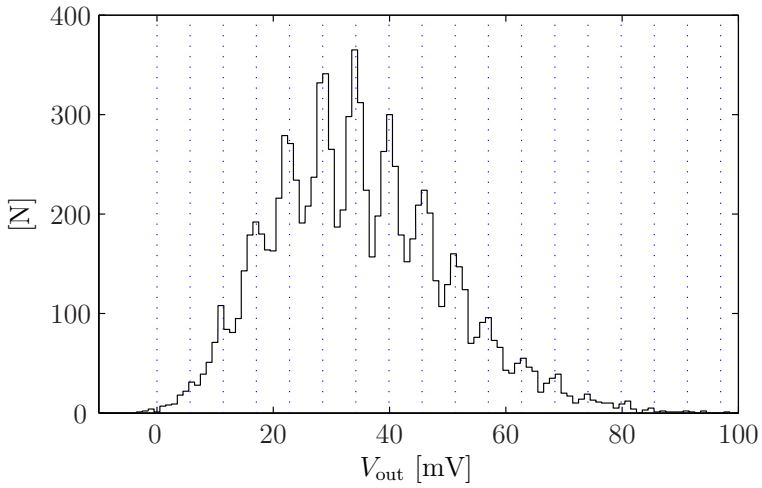

(a) Observed Spectra

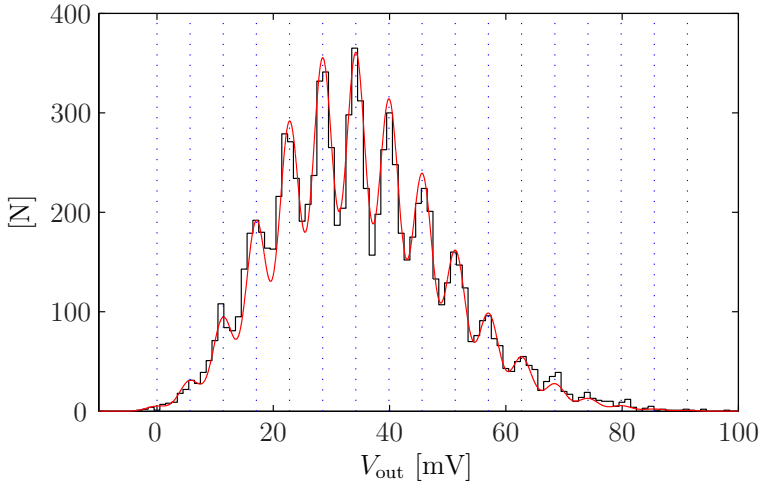

(b) Observed Spectra with Fit

Figure 7.3: Observed Poisson spectra for 1 ms integrations from a Cu rotating anode source, monochromatized at the Cu K_α line at 8.05 keV. A 25 μm pinhole mask was used to isolate the x-ray signal to the interior of a single pixel, thus preventing charge sharing. Panel (a) depicts the observed spectra, panel (b) shows the same result along with a three-parameter fit, where the scaling of the peak separation, the common width of each Gaussian peak, and the location of the zero x-ray peak are allowed to vary and be optimized.

same figure shows a fit to this data in which the scaling of the peak separation, the common width of each Gaussian peak, and the location of the zero x-ray peak are allowed to vary and optimized. Given the limited degrees of freedom, the quality of the resulting fit is quite good, with a reduced χ-squared error of 0.9. One very important consequence of this measurement is that it provides an exceptional tool for determining characteristics of the pixel, such as the pixel gain and the detector read noise, to an accuracy of better than a fraction of a percent.

It is noteworthy the there are very few x-ray imagers that are capable of producing an observed spectrum like that shown in figure 7.3. The fine spectral resolution evident in this figure is a result, not only of the low noise of the Mixed–Mode PAD front-end electronics, but of the merits of direct x-ray detection. As discussed in section 2.3, direct detection of x-rays is a Fano limited process yielding better than Poisson statistic in terms of photocurrent generation. Consequently, phosphor coupled CCDs, even with their lower read noise, would not be able to produce a figure like this due to their much larger uncertainty in signal yield that the indirect x-ray detection method, used by these imagers, introduces. It also goes without saying that a photon-counting PAD could not produce this spectra due to the information lost by photon discrimination. In purely analog PADs, there is a split due to the trade off between gain and total well depth such that analog PADs designed for sensitivity, *e.g.* the PAD being developed for the single-protein diffraction experiments [76], are capable of measurements with this level of sensitivity while deeper well depth devices, *e.g.* the Cornell 100×92 PAD [83], are not. No other analog PAD combines this level of sensitivity with the total well depth offered by the Mixed–Mode PAD.

7.2 Wide Angle Scattering From Sheet Aluminum

While diffraction from sheet aluminium may, at first glance, seem a rather benign experiment, this simple measurement serves to concisely illustrate some very impressive properties of the Mixed–Mode PAD—distinguishing it from other x-ray imagers in use or in development. Recall that it is the combination of analog and digital data in the Mixed–Mode PAD allows for a large total well-depth while simultaneously allowing a high maximum input flux per pixel. As discussed in chapter 4, the Mixed–Mode PAD combines 18-bits of digital data with the well-depth per pixel of an analog charge collector. The well-depth of the analog collector is a variable parameter that may be set within a range of ∼20 to ∼150 10 keV x-rays, though typically the detector is operated with a setting equivalent to ∼100 10 keV x-rays. This configuration yields a total system well-depth of more than 2.6×10^7 10 keV x-rays. The speed of the pixel circuitry is designed to allow a minimum quantized charge removal rate of 1 MHz which corresponds to a maximum input flux of at least 10^8 10 keV x-rays/pixel/s.

To succinctly demonstrate the performance of the Mixed–Mode PAD across a broad range of signal levels and flux intensity a single, one second, exposure of $1/32^{\text{th}}$ in (794 μm) sheet of rolled aluminium was taken, with no beamstop, using the collimated main beam from the Cornell High Energy Synchrotron Source (CHESS) F2 beamline. The intensity of the x-ray beam on the sheet aluminium was on the order of 10^{11} x-rays/mm^2. This configuration produced extreme flux conditions on the detector, near the limits of the Mixed–Mode PAD design specifications, as well as scattered x-ray intensities nearly five orders of magnitude weaker. The resulting diffraction pattern is described in figure 7.4 with views of the same image on four different intensity scales as well as cross-sectional line profile.

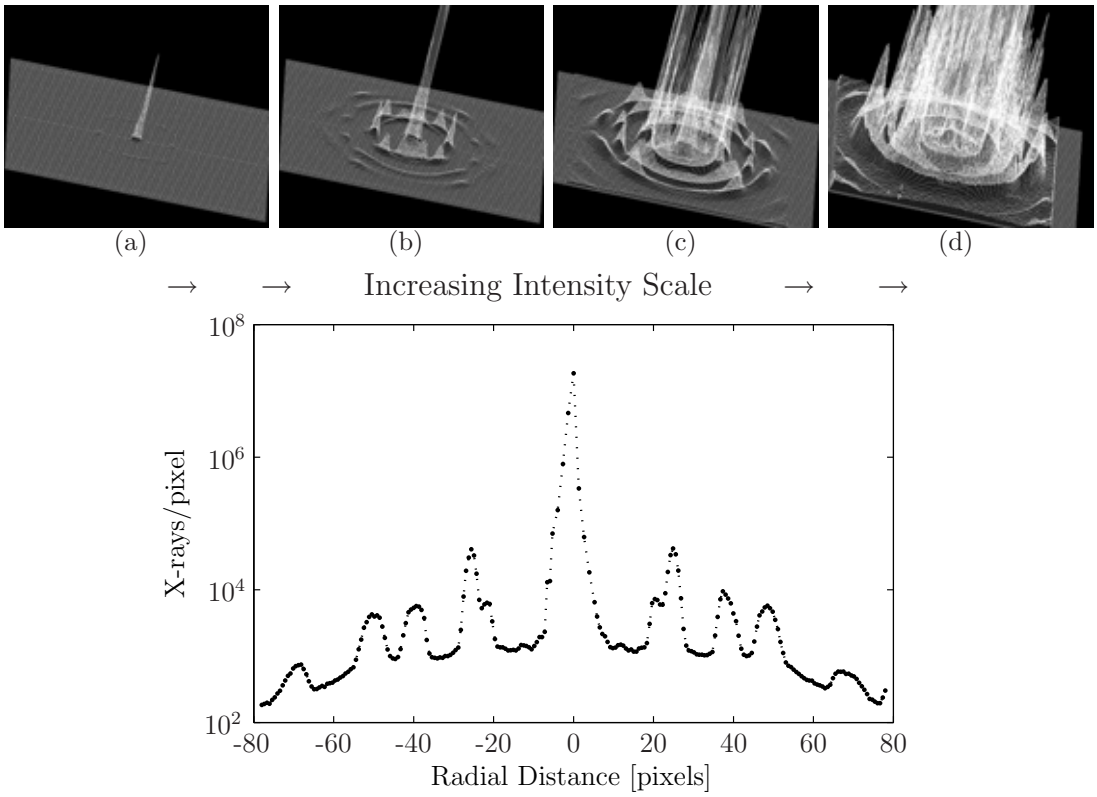

(e) Line Profile Through Diffraction Center

Figure 7.4: A wide-angle diffraction data set from a thin aluminium sheet is shown at increasing intensity scales from image (a) to image (d). An angular profile of this data is shown in panel (e); note that the vertical axis is logarithmic. The data set was acquired in a single, 1 s exposure and clearly illustrates the large dynamic range of the Mixed–Mode PAD. Both the signal of the attenuated main beam (shown in image (a) with a peak flux of 18 million x-rays/pixel/sec) and the sixth-order ring (shown just inside, though not at, the edge of images (c) and (d) or as the 5th peak from the center in panel (e) with a peak flux of ∼700 x-rays/pixel/s) are clearly visible and measured with good statistics although they differ in intensity by a factor of more than 25,000. The dynamic range of the Mixed–Mode PAD is, in fact, larger than this example would suggest, as even fainter rings should also be observable with a larger-area Mixed–Mode PAD detector.

Consequently, this range of flux allowed us to measure, in a single exposure, signal levels that spanned nearly the entire well-depth of the detector. In the resulting diffraction pattern, shown in figure 7.4, the brightest pixel, located in the transmitted image of the direct beam at the center of the pattern, recorded more than 1.8×10^7 x-rays while in the sixth order diffraction ring, visible at the highest level of magnification, the brightest pixel reports only ∼700 x-rays. The intensities recorded by these two pixels differed by more than a factor of 25,000, which notably is larger than one third of the total well-depth of typical phosphor-coupled CCD x-ray detectors, and yet, as was demonstrated in section 7.1, the sensitivity of the Mixed–Mode PAD extends well below the minimum presented by this illumination pattern, indicating that with a larger area detector it should be possible to see even fainter diffraction rings.

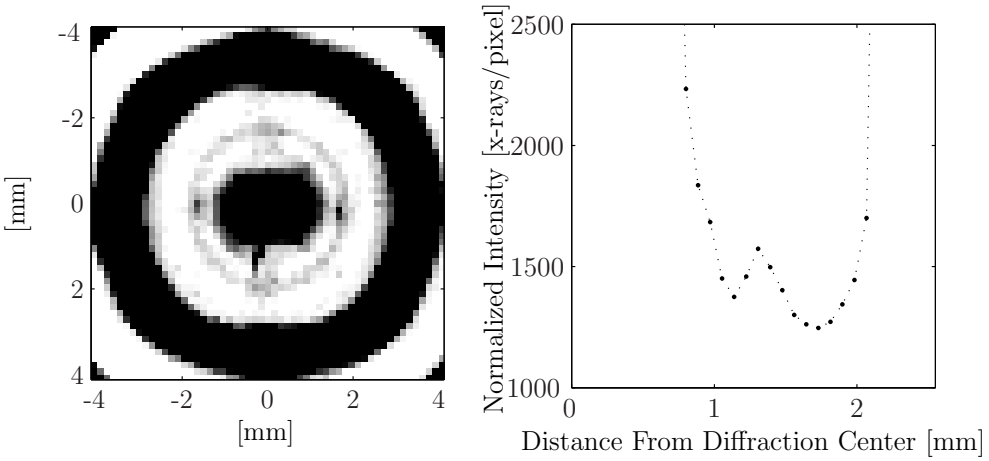

(a) Scaled Diffraction Image

(b) Avg. Intensity vs Radius

Figure 7.5: Panel (a) shows a zoomed in region of the Al WAX image from figure 7.4, scaled to more clearly display the diffraction from higher-order harmonics passed through the monochromator. A quantitative description of this scattering is show in panel (b), indicating an average intensity of 300 x-rays/pixel. What is remarkable about this image is that it is possible to see such a weak signal so near to the much more intense transmitted main beam and primary first-order diffraction ring.

Another remarkable feature about this image, illustrating a consequence of

the sub-pixel impulse response of these detectors, is shown in figure 7.5. Here, a comparably faint diffraction ring, presumably from a higher-order harmonic of the beam, is evident; encircling the transmitted main beam at a distance of ∼1.3 mm from the diffraction center, roughly one half the distance of the primary first order diffraction ring. That this signal is visible so near to the transmitted beam and primary first order diffraction ring, whose maximum intensity per pixel are, respectively, roughly five and two orders of magnitude larger, is a testament to the fact that the spatial spread signal from any x-ray is limited to, at most, nearest-neighbor pixels. Thus, even if one had a phosphor coupled CCD system that had sufficient well depth to record this diffraction image, the tails in its point spread function would prohibit observation of this feature.

7.3 Fine-sampled Image Resolution

As discussed extensively in section 6.4, the spatial response of the Mixed–Mode PAD is dominated by the detector pixelation. For diffraction experiments this is typically not an issue as the diffraction spot or ring, as in the case of powder diffraction or small angle scattering, normally extends over a region of pixels with spot or ring separations greater than the pixel spacing. In these measurements one is typically interested in the total signal in the peak or ring as well as the mean location of the scattering, so the detailed spatial response is typically not a concern. However, this effect must be considered for radiographic or other continuous imaging experiments.

To demonstrate the spatial distortion effect we begin by looking at a single radiographic image, shown if figure 7.6. This image is of a Canadian dime whose head side was filed away to avoid superposition of two images. It was produced by using the dime to occult the unattenuated flood field produced by a Molybdenum

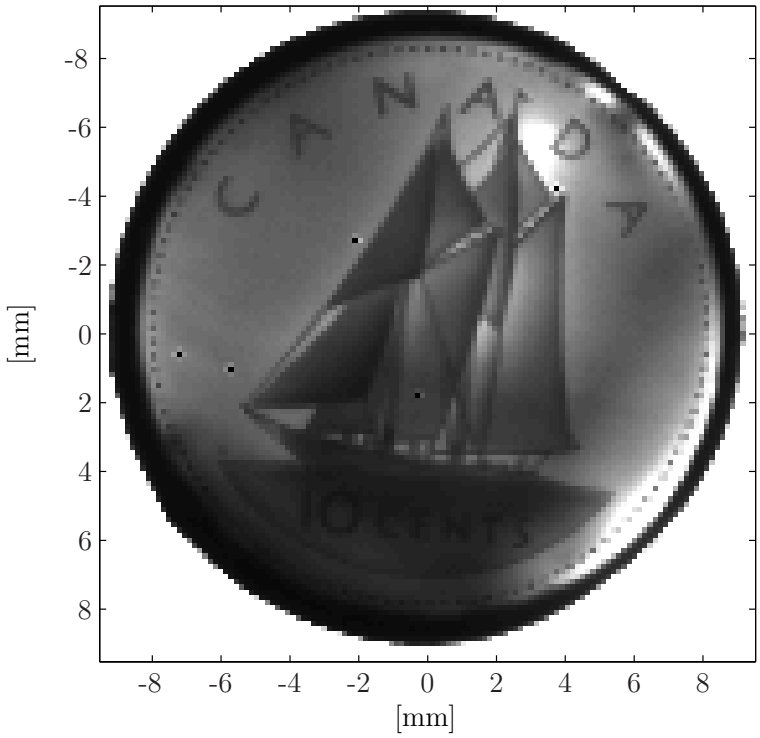

Figure 7.6: Single radiographic image of a Canadian dime taken with a Mo x-ray tube biased at 30 keV. The opposing face of the coin was filed off to provide a clearer image and increase transmission.

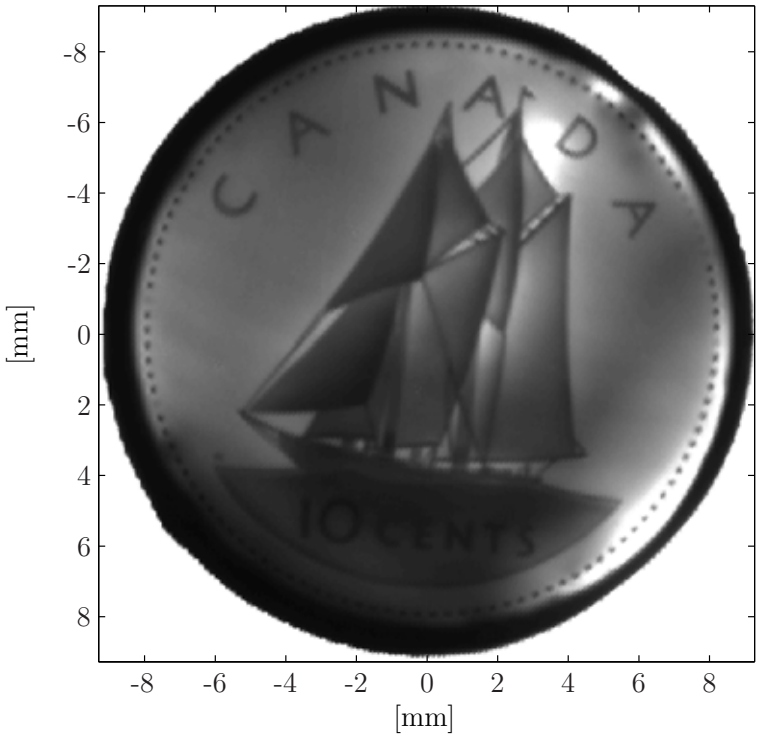

Figure 7.7: Fine–sampled radiographic image of a Canadian dime taken with a Mo x-ray tube biased at 30 keV. The opposing face of the coin was filed off to provide a clearer image and increase transmission.

x-ray tube[6], biased at 30 kV with 0.4 mA of tube current, after a ∼1 m air filled flight tube. The exposure duration was 1 s and 25 background subtracted images were combined to produce the image shown. As expected based on the discussion in section 6.4, the effects of pixelation are most evident along boundaries, where the local spatial frequencies are highest,[7] such as the edges of the sail, the rigging, the dots that border the interior of the coin edge, and the coin edge itself.

With a little extra work though it is possible retrieve the frequencies lost to aliasing, realizing the spatial response of the MTF shown in figure 6.19. To do this we use a method similar to that employed in section 6.4 to study the spatial response of the detector and recover information from spatial frequencies beyond the detector Nyquist Limit. Specifically, by translating the detector in sub-pixel steps it is possible approximate the pixel's continuous spatial response. This is shown in figure 7.7, which takes 100 single images, identical to figure 7.6, randomly located relative to each other and merges their response with the filter discussed in section 6.4.

To elucidate the difference between these images, zoomed regions of the jib sail and coin edge are shown in figure 7.8, with direct comparison between the pixelated and fine-sampled images. What is remarkable about these images is not only the blockage the pixel size imposes on the image, but more so the fine resolution that one is able to retrieve because of the sub-pixel analog impulse response of the detector. In contrast, were an analogous set of images taken with a phosphor coupled CCD, binned to offer similar sized pixels in terms of collection area in the phosphor, one expects very little improvement with fine sampling due to the substantially broader impulse response of the detector.

[6]Model TFS-6050Mo with power supply TCM-5000M (Trufocus–Watsonville, CA).

[7]Local spatial frequencies are a concept from Wavelet Analysis. For more information see [66].

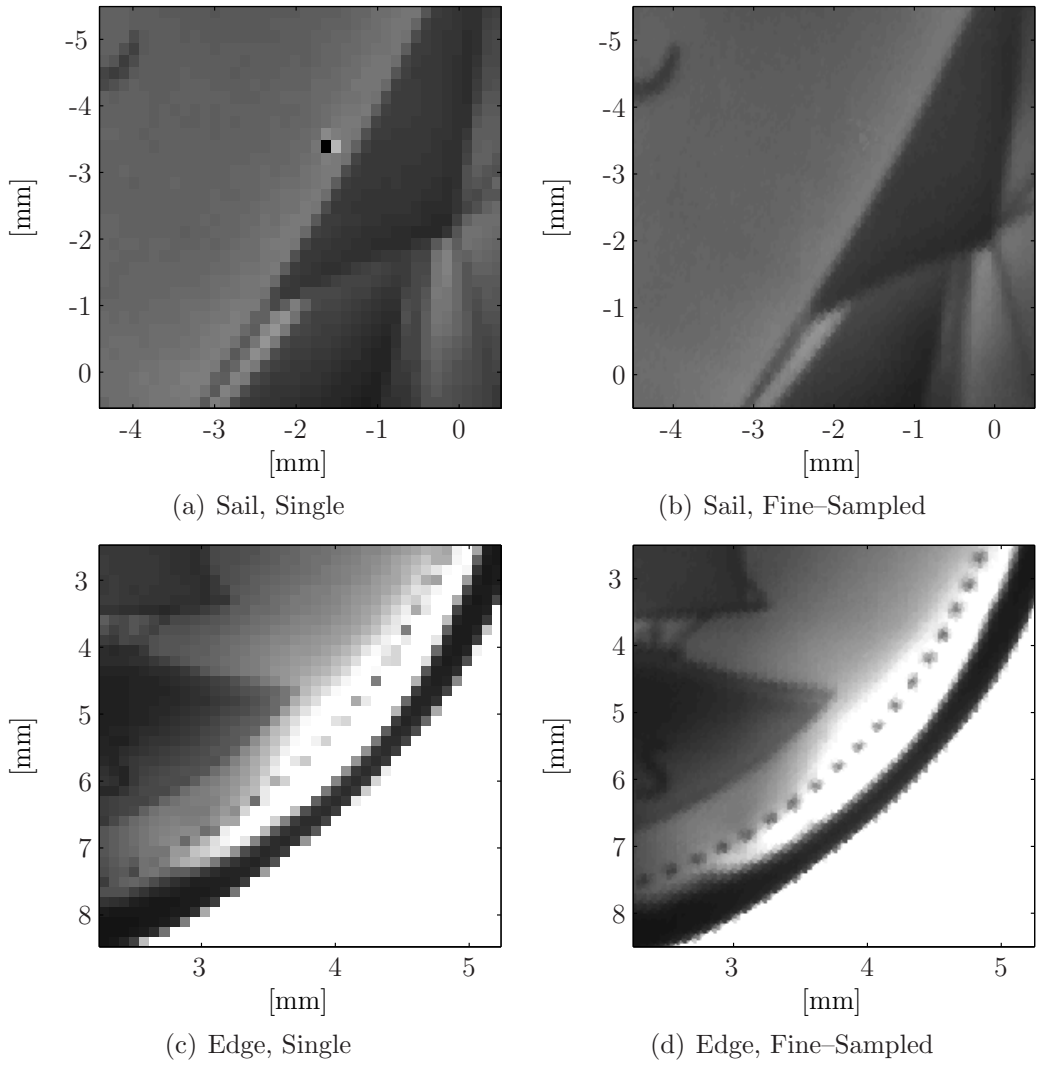

(a) Sail, Single

(b) Sail, Fine–Sampled

(c) Edge, Single

(d) Edge, Fine–Sampled

Figure 7.8: Comparison of magnified regions of figures 7.6 and 7.7. Panels (a) and (c) show sections of the single radiographic image, a portion of the sailboat jib and the right lower edge of the coin, selected to highlight the effects of pixelation on the image. One dead pixel is evident by the black square in panel (a). Panels (b) and (d) show the same regions, respectively, fine–sampled to remove the pixelation effects.

7.4 Protein Crystallography

The Mixed–Mode PAD project was funded by the National Institute of Health in 2003 to build an x-ray detector for Protein Crystallography possessing exceptional characteristics not found in the generation of x-ray imagers available at synchrotron beamlines. With this in mind, it is important that the inaugural synchrotron experiment for the Mixed–Mode PAD should be collecting diffraction from a protein crystal.

The usefulness of a single detector hybrid in this work is somewhat limited, due to its small active area—for comparison, the typical phosphor-coupled CCD system used on a crystallography beamline will have over 100× the active area of a single Mixed–Mode PAD detector hybrid. For this reason, the results presented in this section do not focus on the bread and butter work of most protein crystallographers, solving protein structures, but instead look towards new techniques and experiment possibilities that unique features of the Mixed–Mode PAD bring to this field.

7.4.1 Overview of Protein Crystallography

For a crystal diffraction experiment, a narrow collimated beam of x-rays is used to illuminate a crystalline sample. While the majority of the primary beam is transmitted through of the crystal, a fraction interacts with the electrons bound to each atom and are, as a result, scattered in different directions. If the scattering elements are localized in a structure and these structures are, in turn, arranged in a periodic two- or three-dimensional array, then portions of the scattered radiation may interfere constructively to produce pronounced beams of x-ray intensity [6, 1], whose locations are in accordance with Bragg's Law. A full derivation of x-ray scattering formulas is the proper subject of an extended text, such as the two

preceeding references or [79]. For our purposes though, the primary result from these sources is that the diffraction reveals a portion of the Fourier Transform of the electron density map of the crystal, *i.e.* a portion of the reciprocal space of the crystal. The goal of crystallography is to use diffraction to determine a sufficient portion of the crystal's reciprocal space to reconstruct the electron density map of the basic element of the crystal lattice, the unit cell. From this, then, the atomic structure of the crystal may be derived.

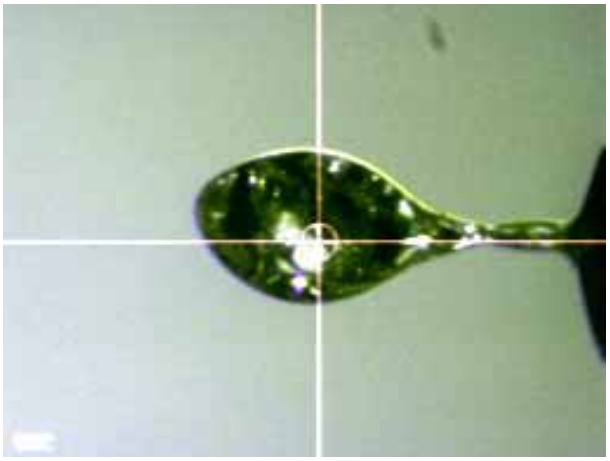

Figure 7.9: Image of the Thaumatin protein crystal used for the diffraction experiments reported in this section.

While this description holds for all forms of crystallography, there are a number of points that distinguish the challenges of Protein Crystallography. Foremost among these is the complexity of the diffraction patterns produced by macromolecular biological molecules. A direct result of the complexity of the protein molecules, and, thus, their electron density maps, is that the diffraction patterns have hundreds to thousands of unique peaks, as the image of a diffraction pattern from a Thaumatin protein crystal shown in figure 7.10 illustrates. Consequently, to obtain sufficient information about the diffraction to reconstruct the electron density map of the protein molecule a substantial portion of the crystal's reciprocal space must be mapped out through rotation of the crystal relative to the

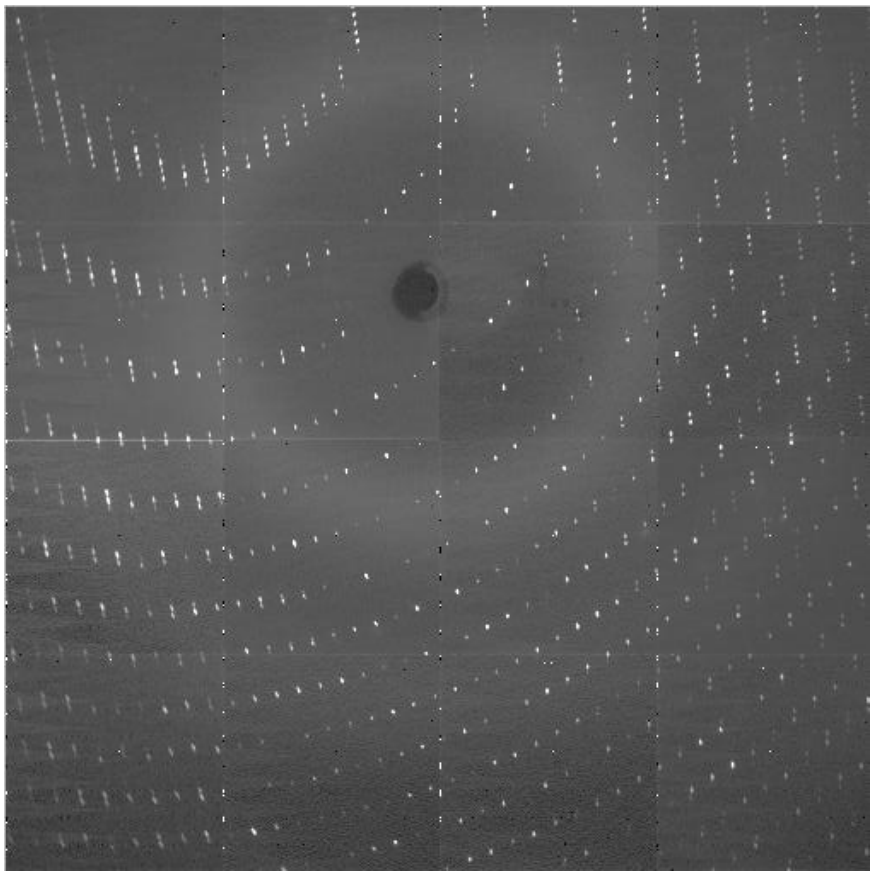

Figure 7.10: Mosaic image of the diffraction pattern from the Thaumatin protein crystal shown figure 7.9, when rotated through $\Delta\phi = 1$ deg. in 1 s. This image was made by combining sixteen separate images (*i.e.* tiles) of the same crystal rotation, acquired with the same single PAD hybrid at sixteen different detector displacements. In each tile a separate background image was subtracted and the tiles global scaling was adjusted to offset beam intensity variation. The border evident at the edge of each tile is due to a one pixel overlap region between images. The data in this edge region is of poor quality due to the high edge leakage of the uncooled detector. This was the first protein diffraction pattern taken with the Mixed–Mode PAD. The image is shown to scale.

incident x-ray beam. The exact number of images this requires will depend on the quality of the crystal, the complexity of the biological molecule, and symmetries in the crystal structure; however, data sets containing hundreds of images are quite common.

As an additional challenge faced by Protein Crystallographers, there are notable differences between even a high-quality protein crystal and an ideal crystal. Large, well-ordered protein crystals are difficult to produce because the large and irregular shaped protein molecules do not pack into crystals without forming large void areas and channels within the individual molecules [15]. As a result, the crystals used for protein crystallography depart from the ideal crystal response of in two ways. First they are rarely a single perfect crystal, but rather made up of many small blocks randomly misaligned with respect to each other. In addition, as each block has finite extent the diffraction spot from each block will have an intrinsic spread, or rocking width, that is inversely proportional to the size of the crystallite. In this "mosaic" model of a crystal, the intrinsic spreads of the separate crystal blocks and the misalignment factors combine to give a width of each diffraction spot on the order of 0.1 deg. [45].

For reasons that will be presented in more detail in the next section, the challenges presented by the large fraction of the crystal's reciprocal space that needs to be sampled and the mosaic spread of these diffraction spots places a strain on the current generation of synchrotron x-ray images. The large-area, phosphor-coupled CCDs typically used to acquire diffraction data sets at synchrotron beam lines are pressed in terms of their frame rate and spatial resolution by the need to acquire many images in as short a time as possible, in which the diffraction from individual spots is well resolved from that of neighboring spots. For these reasons, as well as others we will soon discuss, a new generation of PAD-based x-ray imagers is

being developed, among which the Mixed–Mode PAD, for its integrating rather than photon counting front end, stands out.

7.4.2 Data Collection–Towards Finely Slicing the ϕ

One area in which fast framing imagers like the Mixed–Mode PAD are expected to have a dramatic impact is the way in which synchrotron data sets are taken. To explain this, though, let us first review the canonical approach to acquiring a protein crystallographic data set at a synchrotron source. An example of a protein crystal is shown in panel (a) of figure 7.10. This crystal is mounted on a loop and pin system with a magnetic base. The experimental stage used for performing the diffraction measurement is described in figure 7.11. As this figure indicates, the sample is mounted on the goniometer via magnetic attachment to the mounting pin. While it is mounted in this apparatus, a steady stream of boil-off nitrogen is used to keep the crystal frozen. This stream is supplied by a device known as a coldstream, partially shown in figure 7.10 by the tube directly above the labeled diffraction center, offers a controlled flow rate of temperature-maintained boil-off nitrogen over the crystal to maintain the crystal temperature. Unless noted otherwise this device was used to keep the crystal temperature at $\sim$100 deg. K.

The canonical approach to taking protein crystallography data sets involves rotating the sample through a range of rotation angles ($\Delta\phi$), typically on the order of 0.5 deg. to 2.0 deg., while a single continuous exposure, typically 0.1 s to 20 s in duration, is taken. During this rotation, multiple diffraction spots will come into and then leave the plane of the imager, so that the final image contains an amalgamation of diffraction spots from different rotation angles (ϕ), though for the data set to be reconstructible, the angular steps must be small enough so that individual diffraction spots are distinct and disjoint. The motivation behind this

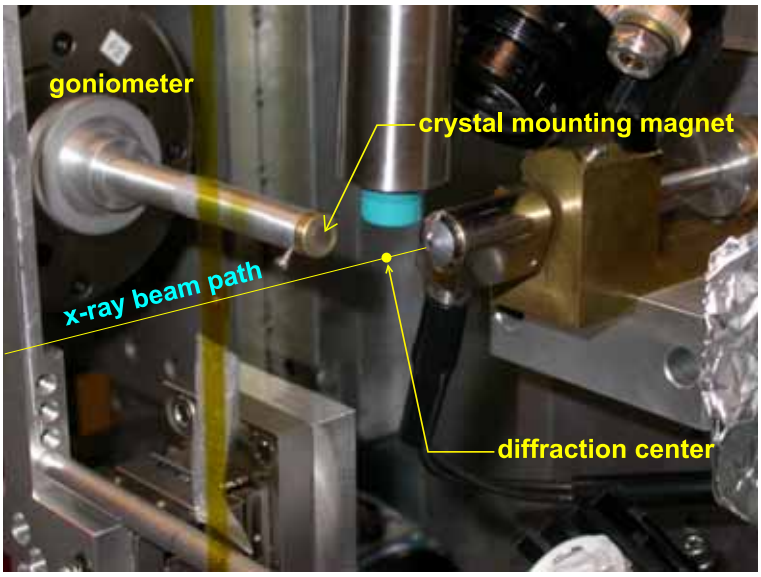

(a) Component Description

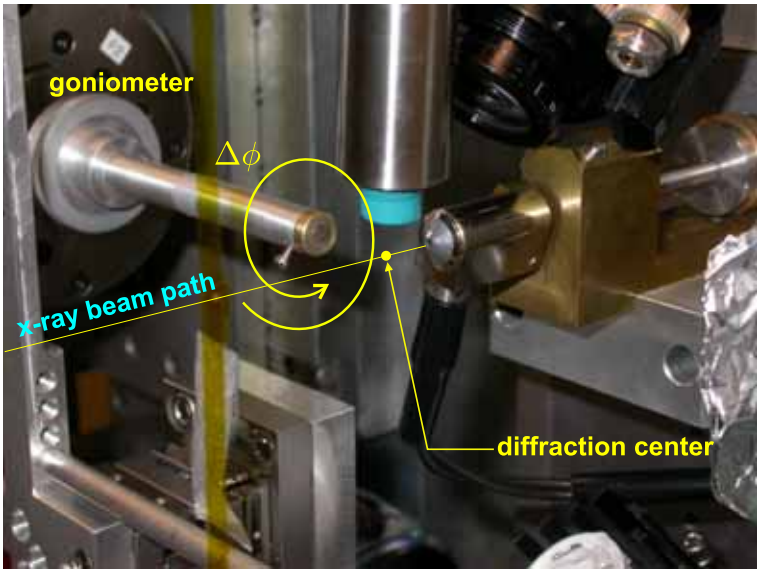

(b) Operational Description

Figure 7.11: Annotated protein crystal diffraction stage at the CHESS F2 beamline. Panel (a) illustrates the main components of this setup while panel (b) illustrates how the crystal is rotated to produce a diffraction series.

method is rooted in pragmatism.

To see why this is, one needs to first appreciate that, with the current generation of x-ray imagers, there is a premium placed on acquiring a diffraction data set in as few frames as possible. This is due to a combination of a need to minimize systematic errors in the data set along with a need to acquire the data set as rapidly as possible. In terms of systematic error, the need to limit the number of frames results from the long, 1 s or more, dead time required to readout the imager. To acquire a data set covering a continuous span of rotation angles, this dead time requires the rotation of the sample be stopped and the sample returned to a reference point during the imager read out between each rotation step. Even with the high-precision goniometer used in these experiments, the stopping and starting of the crystal rotation introduces a small error into the measurement, an error that is compounded by the number of $\Delta\phi$ steps used to complete the rotation series. Adding to this uncertainty, the relatively slow, order of 10 ms response with jitter on the order of ms, x-ray shutters used in these setups contribute an uncertainty to the rotation angle at which the exposure begins and ends. Thus the experimenter is faced with the need to cover a sufficient angular span (180 deg., 90 deg., 45 deg., etc., depending on the symmetries of the crystal) and a trade-off between angular resolution and absolute angular accuracy in choosing the total number and size of the angular steps.

A second consideration comes with regards to acquiring the data set as quickly as possible to preserve the protein crystal quality. Protein crystals are typically frozen to liquid Nitrogen temperature and then maintained near to this temperature during data set collection to reduce the rate of radiation damage [35]. Generally; this damage is taken to be caused by free radicals generated within the crystal and its preserving liquid, under the influence of radiation. These free rad-

icals diffuse through the protein crystal and interact chemically with the protein molecules, altering the structure of individual crystal elements and thereby degrading the diffraction pattern. Keeping the crystal cold reduces the rate of damage, but does not eliminate it, effectively setting a dose limit in which a full data set must be acquired. This, along with the simple fact that synchrotron beam time is limited and precious, means that the number of frames needs to be chosen so that the total time spent reading out the detector is as small a faction as possible of the total time spent acquiring the data, given the limits imposed by the density of diffraction spots in the crystals.

The Mixed–Mode PAD offers a solution to the problems, outlined above, with the canonical approach to data taking. First, thanks to the fast readout capabilities of this detector, it is possible to acquire a complete data set from a protein crystal in one un-broken rotation. This is possible because the Mixed–Mode PAD's front end acts as an electronic shutter, rejecting signal detected in the detector diode layer when an exposure is not active.[8] This dramatically reduces the time required to collect a full crystallographic data set while at the same time improving the quality of the data set by eliminating sources of systematic error. In addition, this data collection method introduces a, potentially very useful, time correlation between frames. Synchrotron sources are not always stable, exhibiting potentially substantial drift due to fluctuations in the beam position and beam intensity. Time correlation of the image frames provides a good opportunity to characterize and remove some of this variability through post acquisition filtering.

Finally, the fast framing capabilities of the Mixed–Mode PAD makes it relatively easy to sample the diffraction data set in angular steps that are smaller than the angular width of the diffraction spots within the set. Known as fine ϕ-slicing,

[8]For low-flux circumstances, $\leq \sim$100 10 keV x-rays equiv. per pixel acquired during the dead time of the detector, or a flux of $\sim 10^5$ x-rays per second in the final detector, one may actually acquire signal during the detector readout, for an effectively dead-time-less device.

this technique has seen some limited application in the past (*e.g.* [45]), but is generally unaccessible. It is expected, though, that this will change once fast readout and fast framing Pixel Array Detectors start to become common enough that the effort required to update aspects of the crystallographic data analysis packages to make use of this new data becomes worthwhile. However, the advantage of ϕ-slicing over the current data collection methods are many.

7.4.3 Synchrotron Measurements

All protein crystallography measurements were made the F2 station at CHESS, during two runs—one in early April of 2007 and one in mid July of the same year. In this work, the CHESS F2 beamline scientist, Marian Szebenyi, was extremely helpful in modifying the crystallography station control software to interface with the Mixed–Mode PAD. Also, the other members of the Cornell PAD group[9] assisted in setting up and staffing shifts during the run. Buz Barstow, Elizabeth Landrum, and Chae Un Kim deserve thanks for providing and preparing the crystals used.

During these runs, the vacuum camera housing, discussed in chapter 5, was not fully operational. As a result, a temporary housing was used in which a flow of nitrogen gas was used to supply a dry environment in the detector housing. This allowed limited the temperature at which the Mixed–Mode PAD was operated to $\sim$8 deg. C. In addition, the camera electronics used with this housing incorporated a version of the analog readout chain that introduced a substantial amount of noise (100 to 200 mV, varying between analog output channel) into the measurement of the analog residual due to a problem decoupling the grounds between the camera and the remote analog-to-digital converters. Consequently, the noise performance of the detector in these measurements was substantially less than the sensitivity

[9]Particularly Mark Tate and Lucas Koerner, but also Hugh Phillips and Marianne Hromalik.

which it is capable of.

7.4.3.1 CHESS F2 Beamline

The F2 station is part of the east wing of the CHESS facility and is primarily dedicated to macromolecular crystallography experiments. The station uses a double-bounce monochromator, two vertically diffracting Si(111) crystals, that monochromate the synchrotron radiation beam produced by positron bunches as they pass through the CHESS East 24-pole wiggler. This system is capable of providing monochromatic beams with a better than 0.1% bandpass over an energy range from 7.9 keV to 14 keV [20].

7.4.3.2 Full-Sized Detector Mosaic Diffraction Image

The first protein crystallography measurement performed at F2 with the Mixed–Mode PAD was the acquisition of a full detector sized, $\sim$512$\times$512 pixel, protein crystal diffraction pattern. The goal of this measurement was to provide an example image to illustrate what could be achieved with a full-size, Mixed–Mode PAD imager. The crystal used was comprised of the protein Thaumatin, a common and very robust molecule often used for developing crystallographic methods because of its well understood structure, ease of crystal growth, and diffraction durability. For this measurement, sixteen independent images were taken of the same 1 deg. in 10 s crystal oscillation, between which the Mixed–Mode PAD was translated relative to the diffraction field. In this way, the image was built up as a series of mosaic tiles with the resulting diffraction pattern shown in figure 7.10.

7.4.3.3 Spot Comparison

Since the task in an x-ray crystallography diffraction experiment is to measure the intensity of a set of diffraction spots, it is important to consider how the proper-

ties of Mixed–Mode PAD will change this measurement relative to the phosphor-coupled CCD imagers currently used at most beamlines.

An advantage of the PAD method of directly detecting x-rays over the indirect detection of phosphor-coupled CCD imagers is the tighter diffraction spots images that they are able to produce. This is a straightforward consequence of the more limited spatial response of charge spreading within a diode in comparison to the spreading of optical light within the phosphor and optical fiber taper of a phosphor-coupled CCD system. This is illustrated in figure 7.12, which shows the same diffraction series, from a Thaumatin Protein Crystal, measured with the Mixed–Mode PAD and the standard commercial phosphor-coupled CCD system[10] located at the beamline.

While the tighter impulse response improves the resolving power of the Mixed–Mode PAD over phosphor-coupled CCD imagers, and, hence, is expected to improve the crystallographic data set quality, this alone will not fundamentally change the way in which crystallographic data is taken. There are, however, other capabilities of PAD detectors that have the potential to do this.

7.4.3.4 Continuous Crystal Rotation: ϕ-Profiling

As mentioned previously, one of the expected advantages of the Mixed–Mode PAD is to divide a single frame of canonically taken data into a sequence of frames acquired under continuous rotation of the crystal, providing a detailed picture of the diffraction spot profile and eliminating excess background. This is illustrated in figure 7.13, which shows diffraction data from a Thaumatin protein crystal. Panel (a) of this figure shows the results of a 1 deg., 10 s oscillation taken by the Mixed–Mode PAD with the canonical, long-exposure, method outlined above. Panel (b) focuses in on one diffraction spot in this image, detailing the additional

[10]Quantum-210 (ADSC–Poway, CA).

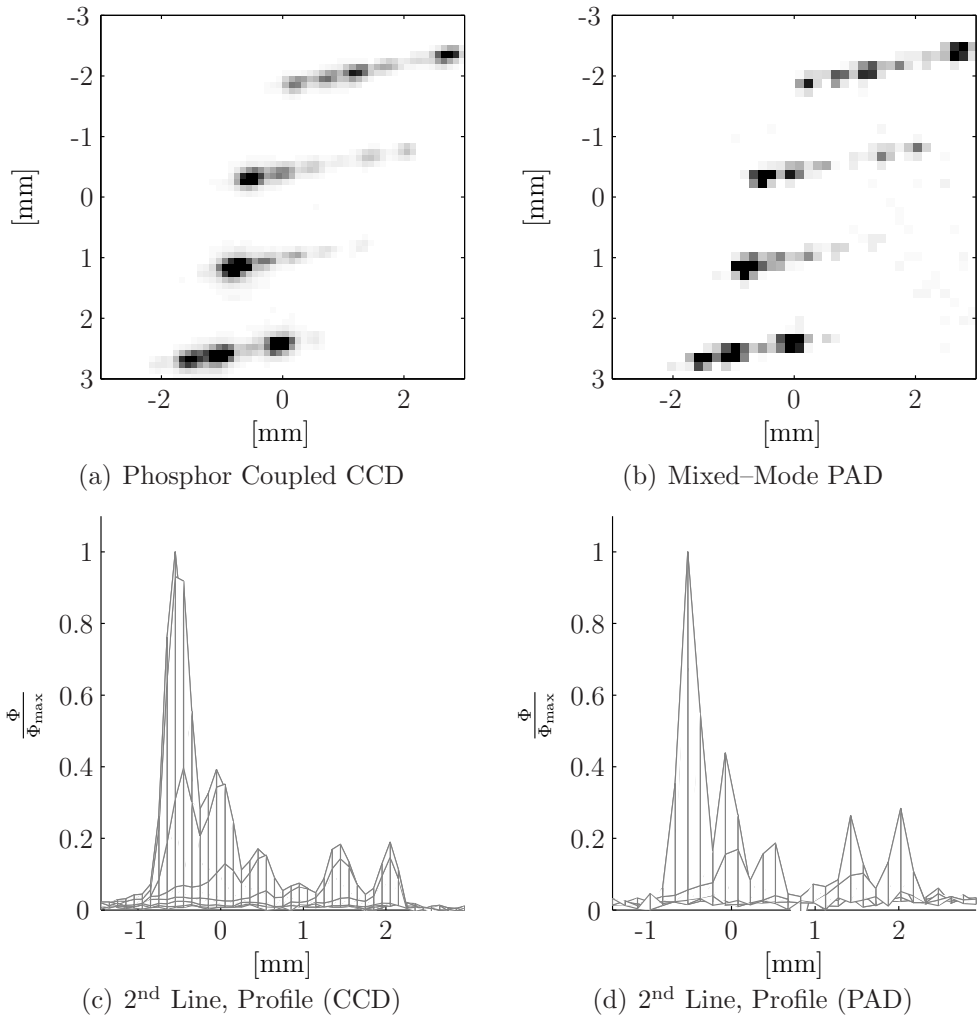

(a) Phosphor Coupled CCD

(b) Mixed–Mode PAD

(c) 2nd Line, Profile (CCD)

(d) 2nd Line, Profile (PAD)

Figure 7.12: Comparison of identical regions of a Thaumatin diffraction pattern, taken over a 1 deg., 10 s crystal oscillation, with a phosphor-coupled CCD system and the Mixed–Mode PAD. Panels (a) and (b) display a series of lines of diffraction spots taken with a phosphor-coupled CCD system and the Mixed–Mode PAD (resp.). Panels (c) and (d) show background subtracted contour profiles of the second line, indexed from the top of the respective image, normalized to the peak height of the brightest spot. The point of view for these profiles is taken to be along the vertical axis of panels (a) and (b) (resp.), in the positive direction. The missing 4th peak, indexed from the left, in the Mixed–Mode PAD line is due to a bad pixel.

information that becomes available through fine ϕ-slicing. In it a one deg., 10 s crystal oscillation is divided into fifty successive frames taken as the crystal was continuously rotated for a $\Delta\phi$ resolution of $\sim$0.02 deg. per frame.[11] The integrated intensity of a single spot is plotted as a function of the crystal oscillation angle. Dependant on the width of the $\Delta\phi$ resolution, more information may be resolved from individual diffraction spots.

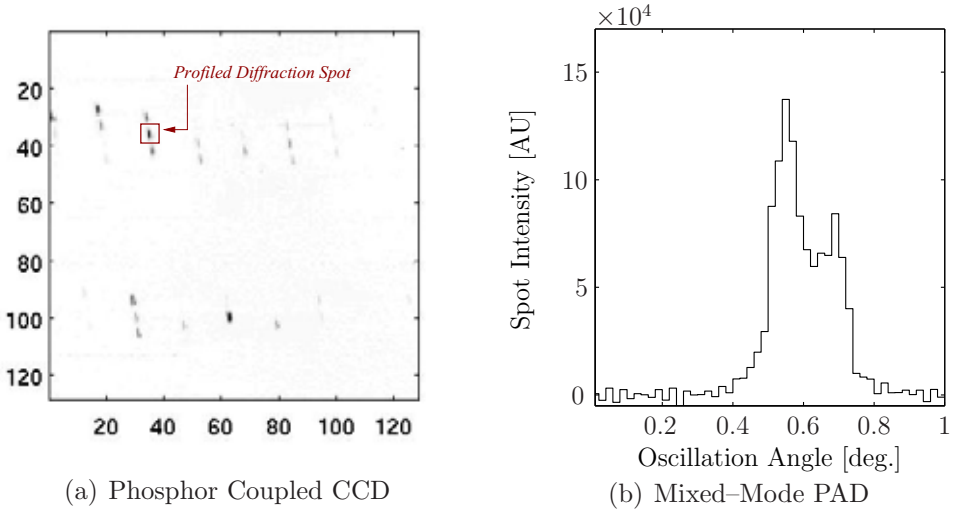

(a) Phosphor Coupled CCD

(b) Mixed–Mode PAD

Figure 7.13: Comparison of an element from a canonical macromolecular data set to the additional information revealed by fine ϕ-slicing a continuous crystal oscillation. Panel (a) shows a frame taken with the Mixed–Mode PAD containing a 1 deg., 10 s crystal oscillation. In panel (b), the same oscillation is divided into 50 frames and the integrated intensity of one diffraction spot, as indicated in panel (a), is profiled.

Beyond the apparent improvement one receives from limiting the integral of the spot intensity to regions of ϕ where the spot is actually present, thereby reducing the background integrated into the spot, increasing the $\Delta\phi$ resolution reveals new information about the underlying structure of the crystal sample. As figure 7.14 illustrates, as we increase our resolution, we go from blurred measurements of the

[11]There was a 5 ms delay between frames due to the readout time of the prototype camera. During this time the reset switch of the integrator was used to electronically shutter the detector.

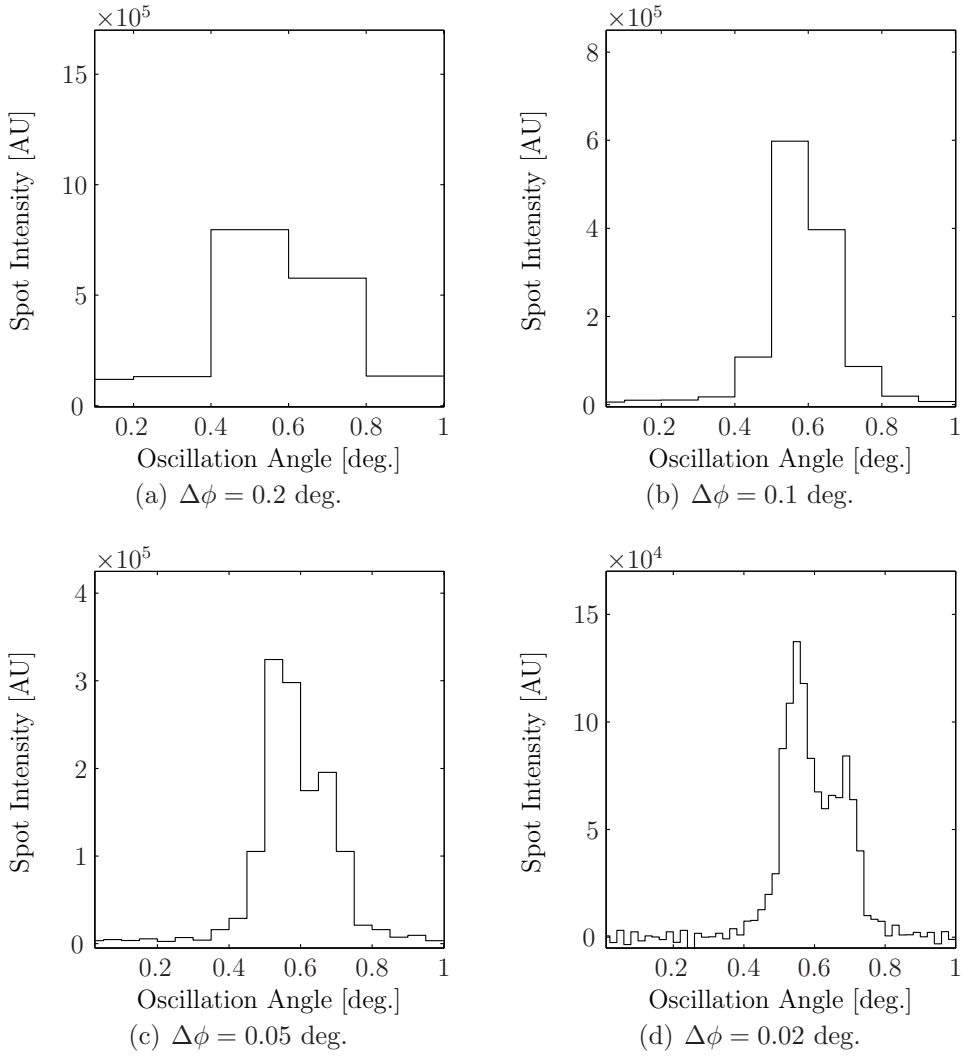

Figure 7.14: ϕ-slicing on the same diffraction spot at differing levels of $\Delta\phi$ resolution: 0.2 deg., panel (a); 0.1 deg., panel (b); 0.05 deg., panel (c); and 0.02 deg., panel (d). This spot was produced by a Thaumatin crystal, with the intensity spread over 15 pixels on the detector, during a 1 deg., 10 s continuous exposure.

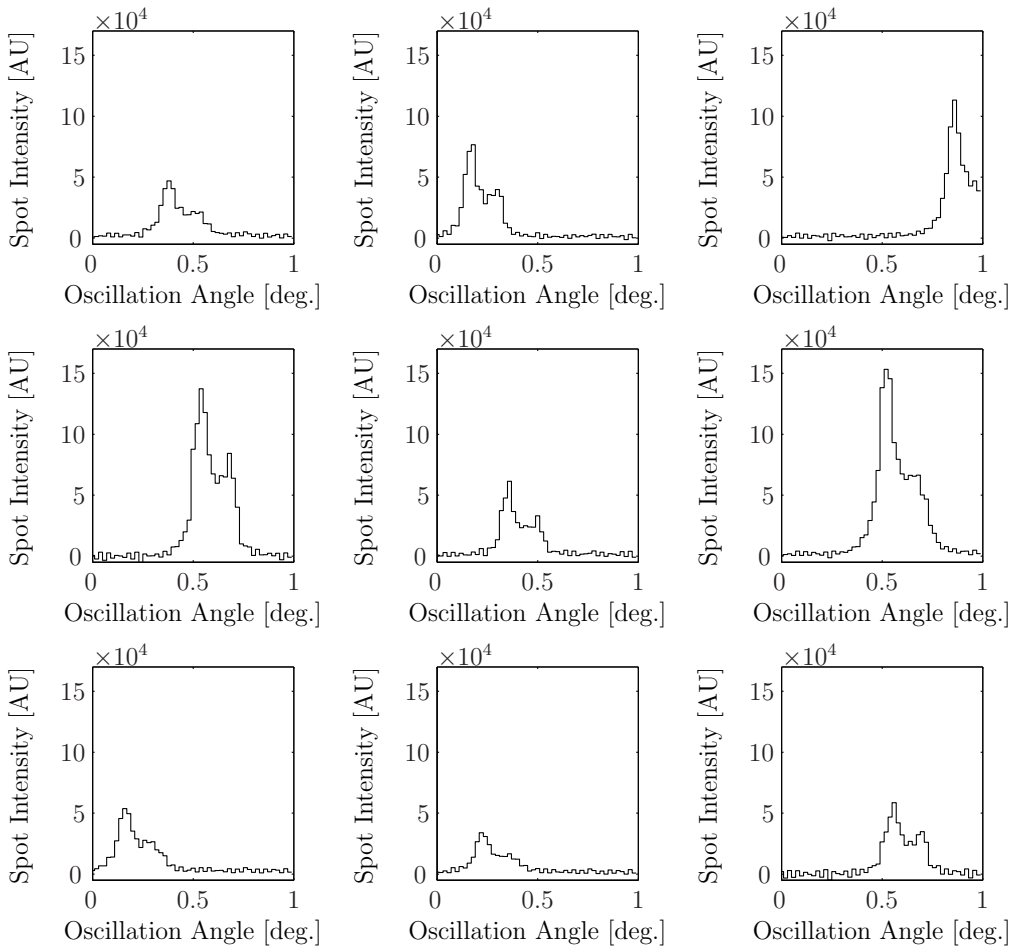

Figure 7.15: ϕ-sliced diffraction spot profiles, taken in $\Delta\phi = 0.02$ deg. steps, for a series of different spots taken from the same Thaumatin crystal, in the same frame set. The form of the diffraction profile is echoed in each spot, as one expects since the profile reflects the underlying structure of the crystal.

total diffracted intensity to being able to resolve the fine profile of this diffraction spot. In the case of this particular crystal, the information revealed by finely slicing the $\Delta\phi$ suggests that we have at least two separate crystal domains, evident in the two peaks in the diffraction profile of the spot.

Another possibility presented by finely slicing the $\Delta\phi$ steps comes from the repetition of the basic spot profile over all spots within the crystal. As discussed earlier, the profile of a diffraction spot is determined by the size and mosaic spread

of crystallites within the crystal. Hence, it should be the same for every diffraction spot from a given crystal. This is illustrated in figure 7.15, which shows nine different diffraction spots from the same crystal, taken in the same 1 deg. in 10 second sequence of frames. This fact can be used to improve the measurement of the spot intensity and centroid, particularly in weak and noisy spots, by fitting spot profiles derived from an ensemble of more intense spots.

As a final advantage of this data collection method, it can be used to observe changes within the crystal induced by external factors. To illustrate this, fine ϕ–slicing data sets were taken of a Thaumatin crystal before and after warming the crystal from 100 deg. K to 170 deg. K. A profile from one diffraction spot in this data set is shown in figure 7.16. Looking at the normalized overlay, one sees that warming the crystal, in addition to altering the lattice parameters of the crystallites as is evident by the shift in the peak position, effects the mosaic spread of crystallites within the crystal. For example, the sharp shoulder evident in near an oscillation angle of 2.5 deg. in the 100 deg. K data set has disappeared at 170 deg. K and the full angular spread of the peak has broadened by roughly a third again upon its original angular spread.

7.4.4 Reflections on Protein Crystallography

The short dead time and high frame rates which the Mixed–Mode PAD is capable of facilitate new methods of data collection that were previously not possible in phosphor-coupled CCDs and earlier area x-ray imaging technologies. The acquisition of data during continuous crystal rotation is one of the most important practical advances of the Mixed–Mode PAD over the current generation of synchrotron x-ray imagers. This acquisition mode reduces the total time spent acquiring a crystallographic data set, removes sources of systematic error, minimizes the back-

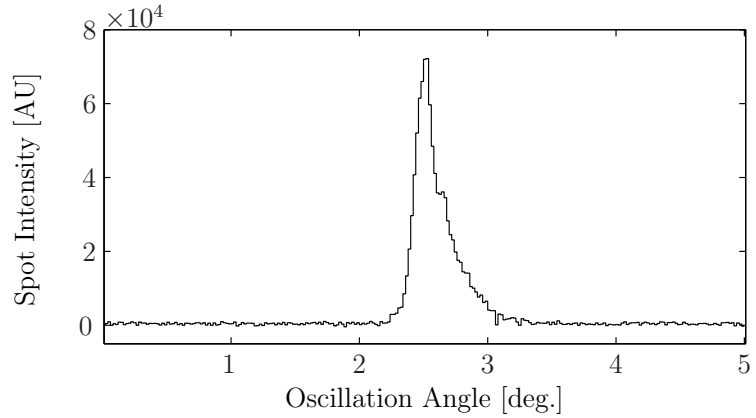

(a) 100 deg. K

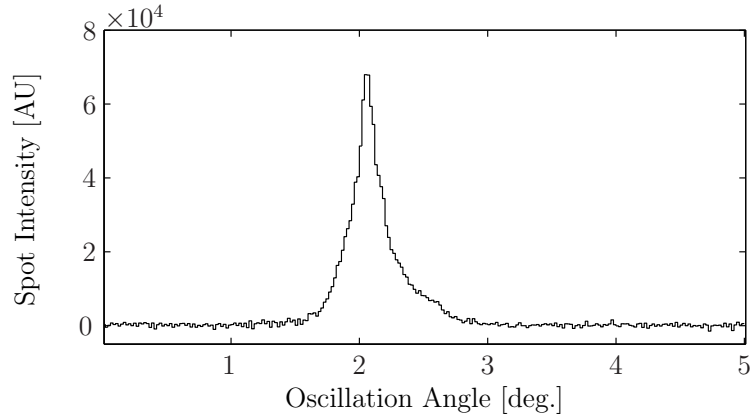

(b) 170 deg. K

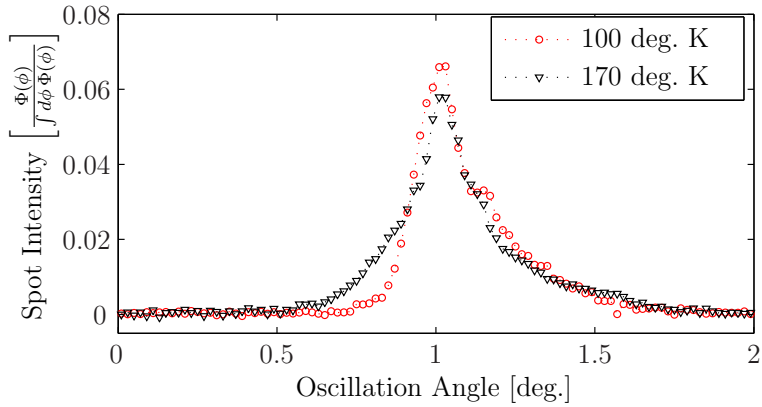

(c) Comparison of Normalized Intensity Profiles

Figure 7.16: Change in the ϕ-sliced profile of a diffraction spot from a Thaumatin crystal before, panel (a), and after, panel (b), the crystal was warmed from 100 deg. K to 170 deg. K at a warming rate of 6 deg. K per minute.

ground integrated into each diffraction spot, and gives the experimenter previously unavailable information by revealing the profile of the diffraction spot.

7.5 Time-Evolving Systems

One area where it is anticipated that Pixel Array Detectors, in particular the Mixed–Mode PAD, will make significant contributions is in dynamic studies of time-evolving systems. As mentioned earlier, the current generation of synchrotron area x-ray imagers require substantial time, typically > 1 s, to read out the detector. Consequently, most time-resolved synchrotron experiments requiring continuous data acquisition either utilize point or strip detectors [32, 52]. Imaging experiments are limited to time scales of seconds or longer [90] or utilize techniques like high-speed mechanical or electrical shutter systems to gate x-rays onto the detector for brief periods of time, building up a temporal mosaic of the system response [65]. Yet, even with these limitations time-resolved studies of dynamic systems is a very active field of research.

The Mixed–Mode PAD promises to extend the reach of continuous framing imaging experiments into the millisecond regime, effectively adding a new dimension to a domain currently limited to point and strip detectors. It is well suited to this work, not only because of its fast framing capabilities, but also because of the large dynamic range it offers for short exposures. To understand this claim consider that the standard definition of dynamic range,

$$\text{Dynamic Range} = \frac{\text{Well Depth}}{\sigma_{\text{read}}}, \tag{7.5}$$

where σ_{read} is the read noise of the detector—canonically taken to be the smallest signal one can detect above the detector noise[12]—does not extend well to short

[12]For a reliably operating photon counting devices this denominator is simply 1, because this is the smallest level at which one may detect signal.

exposures. As an example, consider a photon-counting detector with a well depth of 2^{20} counts (x-rays) which may be reliably operated up to a flux limit of 10^6 x-rays/pixel/s. In a 10 ms exposure, this detector can record at most 10^4 ($\sim 2^{13}$) x-rays/pixel, substantially less then the detector total well depth. Consequently, for short exposures an alternative definition of the detector dynamic range proves more useful,

$$\text{Effective Dynamic Range} = \frac{\Phi_{\max} \cdot t_{\exp}}{\sigma_{\text{read}}}, \tag{7.6}$$

where $\Phi_{\max}$ is the per-pixel flux limit of the detector and $t_{\exp}$ is the exposure duration. This definition shows the advantage that a detector with a high flux tolerance, like the 10^8 x-rays/pixel/s of the Mixed–Mode PAD, has over devices with lower flux limits in high-speed imaging.

As a first experiment using the Mixed–Mode PAD to study the dynamics of an evolving system, the imager was employed in the study of homoepitaxial $SrTiO_3$ thin film growth through Pulsed Laser Deposition (PLD). This particular experiment was chosen because it is a system known to exhibit dynamics on millisecond time scales and is the subject of active research by another group at Cornell.

In the remainder of this section, we first present an overview of the PLD thin-film growth technique followed by a discussion of the ongoing PLD research being conducted by our collaborators in this experiment, the Brock Group in the Applied and Engineering Physics Department at Cornell University. Results form the first Mixed–Mode PAD PLD experiment are then presented, along with a discussion of the Mixed–Mode PADs strengths and weaknesses in this experiment. Finally, we conclude with some comments on the unique possibilities a fast-framing, wide-effective-dynamic-range, imaging detector offers for time-resolved studies.

7.5.1 PLD Overview

Pulsed Laser Deposition (PLD) is a thin film growth technique that has enjoyed much attention in the research community since the late 1980s when it emerged as a promising technique for the growth of high–T_C superconducting films.[13] Today, it is an active area of research that has expanded significantly from its original focus on superconducting thin films. Researchers hope to one day be able to grow materials that would be thermodynamically impossible with other techniques, such as multilayer films with layers an arbitrary number of monolayers deep or arrays of quantum dots that could be integrated into optoelectronic or microelectronic devices [21].

The experimental setup required for PLD thin film growth is relatively straightforward. Within a deposition chamber, the beam from a high-powered laser is used as an external energy source to ablate material from a target. This ablated material forms a plume which deposits a thin film of particles on a nearby substrate. The composition of the target, the ambient environment, the substrate material, temperature and orientation, as well as the laser pulse duration, intensity, and wavelength may all be varied to control film growth.

The laser–target and plume–substrate interaction are very complex physical phenomena. Theoretical descriptions are multidisciplinary combining both equilibrium and non-equilibrium processes. For this reason, significant amounts of experimental information are needed to develop accurate descriptions of thin-film growth. Historically, PLD researchers have relied on post growth microscopy, such as Atomic Force Microscopy (AFM), Scanning Tunneling Microscopy (STM), Scanning Electron Microscopy (SEM), Transmission Electron Microscopy (TEM), etc.,

[13]The history of PLD actually began long before this. The first paper on the topic was published in 1965 [91], a few years after the first high-powered ruby lasers became available. However, the field remained somewhat stagnant until the first successful growth of high–T_C superconducting films was reported in 1987.

to analyze the effects of different growth parameters [21]. This approach suffers from a number of limitations, foremost being that it is *ex-situ*, so only information about the film's final state is available. Furthermore, in order to be examined, the substrate must be removed from the deposition environment and cooled, often hundreds of deg. C, during which changes in the film structure are expected.

In the early part of this decade, work with Reflective High–Energy Electron Diffraction (RHEED) and x-ray surface studies, the latter made possible by the availability of two PLD at synchrotron sources (one at the European Synchrotron Radiation Facility (ESRF) in Grenoble, France and one at CHESS in Ithaca, New York, USA), yielded *in-situ* information about film growth that challenged many assumptions about the processes governing the monolayer growth [21, 32]. The bulk of the work at millisecond time scales, however, was limited to a general characterization of film surface roughness, with post growth microscopy relied upon to characterize surface structures. Very recently, as we will discuss further in the next section, this work was extended through *in-situ* measurements in the scattering plane to reveal information about the correlation length of structures in the plane of the film, as they evolve during the growth.

7.5.2 PLD Studies by the Brock Group at CHESS

Dr. Joel Brock's group in the Applied and Engineering Physics Department of Cornell University has been using the PLD chamber in the CHESS's G3 hutch to study thin-film growth using a point detector with the sample in an anti-Bragg reflection geometry. In this mode, the intensity observed by the detector is directly related to the surface roughness of the thin film [33]. By depositing less material than is necessary to complete an atomic layer one can observe oscillations in the reflected intensity—typically called RHEED oscillations in deference to their original

discovery in electron scattering experiments. This occurs because, under certain kinetic conditions, deposition of material on a smooth surface serves to roughen it while deposition on a rough surface serves to smooth it, resulting in periodic roughening and smoothing of the sample, thereby allowing one to monitor single molecular layer growth.

While PLD has been shown capable of growing high-quality crystalline films, a complete model of the PLD process at the atomic level does not exist. The conventional picture has held that molecules and atoms in the ablation plume strike the substrate or film surface at randomly distributed positions. These particles may then evaporate, bond to the surface at existing step edges, or collide with other particles to nucleate new islands [21]. This treatment assumes that the behavior of the particles on the surface is predominantly a thermal process, neglecting the kinetic energy of the incident beam. By using *in-situ* measurements using the point detector techniques discussed above, members of Dr. Brock's group were able to show, in 2005, that this picture cannot fully explain PLD film growth [32].

More recently, this group has employed a phosphor-coupled CCD operated in strip detector mode[14] in their PLD growth studies. In this operating mode, the CCD acts as a large 1D detector, accomplished by shifting rows within the CCD array into the readout register of the CCD, yet waiting to read this register until it has collected the charge from a full column of exposed pixels. The operation of shifting charge from pixel rows into the readout register effectively sums the charge, allowing an entire column of pixels, or sections thereof, to be combined with very little noise contribution. In addition, the frame rate of the detector is substantially increased, because the next frame's worth of data may be acquired

[14]This mode is sometimes referred to as 'streak camera mode'; however, this terminology is confusing as detectors known as streak cameras are often employed in particle accelerators and storage rings to characterize particle beam properties [88]. As this operation is quite different from how the Brock group has been using their CCD, we adopt the 'strip detector' terminology for clarity.

while the current frame is being read from the readout register.

By operating the CCD in strip detector mode, the members of the Brock group have been able to study the diffuse scattering that accompanies the specular reflection in the off-axis field on either side of the specular reflection. The profile of this scattering, as a function of scattering angle, gives information about the length scales of structures formed as intermediaries of a complete monolayer. Because the Brock group is, currently, studying the growth of homogeneous monolayers on a substrate of the same material (homoepitaxial growth) and the system has only one variable dimension, the size and separation of the monolayer islands, the scattering only varies in the plane of the sample. Thus, for this experiment, there is little a 2D detector can offer beyond the capabilities of an ideal 1D detector.

However, the CCD in strip detector mode is far from an ideal 1D detector. A significant drawback of the strip detector method is that it substantially diminishes the integrated flux the detector is able to acquire per unit area of each pixel column. This is because the well depth of the readout register does not change when in strip mode although the effective area of each pixel increases dramatically. As a consequence, moderately intense features on the camera in 2D mode may easily cause it to saturate when operated in strip mode. This is problematic for the PLD experiment, because it means that it is not possible to measure both the strong specular reflection and the weak diffuse scattering, in the same image, without substantial attenuation of the specular intensity.

In contrast, the Mixed–Mode PAD has the ability to measure the entire pattern, both the specular and diffuse, directly and simultaneously. In addition, its high frame rate makes imaging the 2D specular and diffuse patterns, on the same time scales as the Brock Group's previous 1D measurements, quite straightforward.

7.5.3 Synchrotron Studies of Monolayer Growth

The PLD measurements were performed at the G3 station at CHESS. The measurements shown here came from one run that lasted for the majority of October 2007.[15] This work was done in collaboration with members of Dr. Brock's group from the Applied and Engineering Physics Department of Cornell University; in particular, John Ferguson and Gokhan Arikan, graduate students from this group, were essential in providing samples, operating the laser and deposition chamber, as well as assisting in data collection. The G-Line staff scientist, Arthur Woll, was very helpful by providing example scripts illustrating how we could interface the Mixed–Mode PAD single hybrid prototype camera to the beamline control software as well as offering very useful advice on the experiment, as the run progressed. Members of the Cornell PAD group[16] also assisted and much thanks is especially due to Lucas Koerner for taking on an equal share of the operating time throughout the duration of the run.

7.5.3.1 CHESS G3 Beamline

The G3 Beamline is a part of the G-line complex built as an addition to the Cornell High Energy Synchrotron Source (CHESS). It operates off of positron bunches circulating through the 49 pole wiggler that also feeds CHESS West, from the storage ring's electron current. G3, in particular, is a hutch designed for time–resolved studies requiring a high flux, delivering upwards of 10^{13} x-rays/s/mm^2. A series of two synthetic multilayer monochromators are used supply a monochromatic beam over an energy range of 8 keV to 12 keV with a 1% bandpass [20, 52].

[15] PLD measurements are extremely time intensive, with pre- and post-sample requiring many hours. As a result, when everything is working well, testing four or five samples in a 24 hour period is an achievement.

[16] Mark Tate, Marianne Hromalik, and Hugh Philipp assisted in setup and covered some data taking shifts.

7.5.3.2 Homoepitaxial $SrTiO_3$ Growth

As this data was taken collaboratively with a group that focuses on PLD growth, and, as such, constitutes a portion of a body of work they are developing, we will not draw any conclusions regarding PLD growth in this thesis. Instead, we will focus on the performance of the Mixed–Mode PAD in these experiments. To this end, we will present data from one homoepitaxial $SrTiO_3$ growth observed with the Mixed–Mode PAD during this CHESS run. As with most of the other growths, a series of 800 images were taken,[17] where the duration of each exposure was 95 ms with 5 ms spent reading out the imager for a 10 Hz frame rate. The repetition rate of the laser was set at 0.2 Hz so that 50 frames would be acquired between each laser shot with the initial firing occurring after the first 50 frames were taken.

The laser used to create the PLD plume was a 348 nm KrF Excimer Laser whose fluence on the single-crystal $SrTiO_3$ target was $\sim$2 J-cm^2. The separation between the target and the growth substrate was 6 cm. These substrates were prepared by members of the Brock group using standard procedures for forming well-ordered TiO_2 terminated (0 0 1) surfaces of $SrTiO_3$ [52]. Typically, an additional annealing step was performed within an oxygen-rich environment in the growth chamber, just prior to deposition, to further diminish surface roughness. Details regarding the growth chamber as well as sample preparation steps may be found in [33].

A typical image captured from a 95 ms integration with the Mixed–Mode PAD is shown in figure 7.17 where panel (a) shows the detector's full range while panel (b) depicts the same image scaled from 0 to $\sim$10 x-rays. These two panels detail the two regions of interest, contrasting the intense specular reflection near the center of the image with its substantially weaker wings of diffuse scattering. The absence of x-rays outside of $\pm$5 mm of the central axis of the image was due to slitting

[17] A limitation of the prototype camera that was discussed in chapter 5.

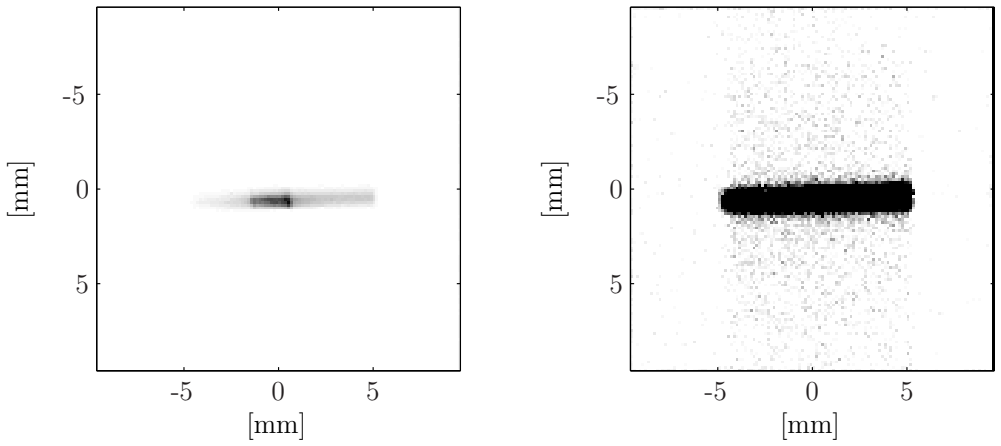

(a) Full Scale (b) Scaled from 0 to 10, 10 keV X-Rays

Figure 7.17: 95 ms exposure of $SrTiO_3$ at the $\left(0\,0\,\frac{1}{4}\right)$ after a growth series, shown if figure 7.18, near maximum diffuse scattering oscillation. Panel (a) is scaled to cover the entire range of the image, from 0 x-rays per pixel to 2,218 x-rays per pixel, while panel (b) is limited to show the diffuse scattering, whose intensity is at most a few x-rays per pixel. Both images are shown in the negative and the intensity floor is set at twice the read noise ($2\sigma_{\text{read}}$) so that spots in the images actually represent 1 or more x-rays.

of the beam. Within the un-occulted region, there is interest in the variation of integrated intensity of the specular reflection and diffuse scattering, as well as the profile of the diffuse scattering, with time from the start of the growth and each successive material deposition. Our results for this growth are summarized in figure 7.18. Panel (a) of this figure depicts the evolution of the integrated specular intensity with material deposition and surface relaxation. Within the specular reflection growth oscillations, each maxima represents the completion of a single $SrTiO_3$ monolayer. Accompanying out of phase oscillations in the integrated diffuse scattering are shown in panel (b), along with the profile of this scattering in panel (c).

Before going further, some remarks are warranted to explain how the data in figure 7.18 was derived from images akin to our example from figure 7.17. The straightforward method of measuring the specular and diffuse signal with the

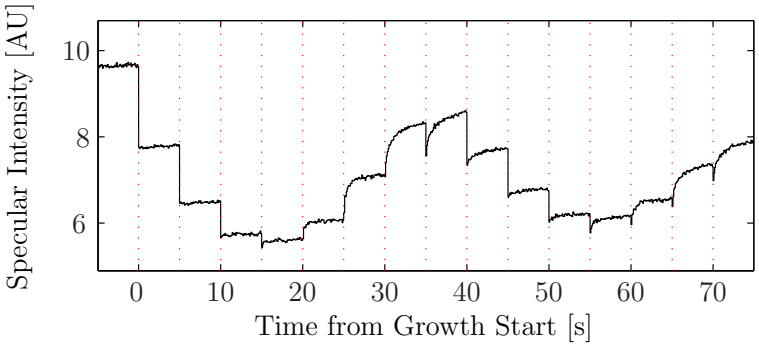

(a) Integrated Reflected Specular Intensity

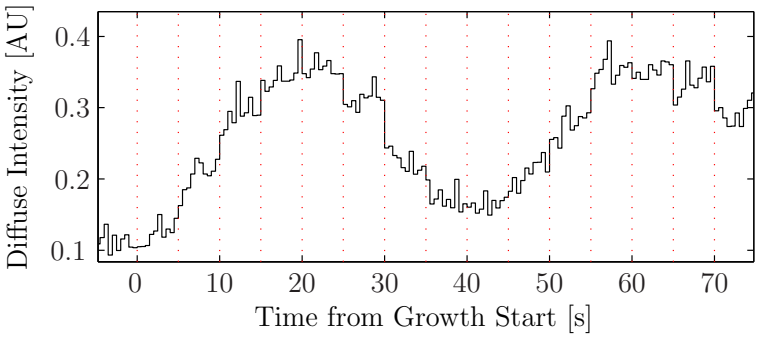

(b) Integrated Diffuse Scattering Intensity

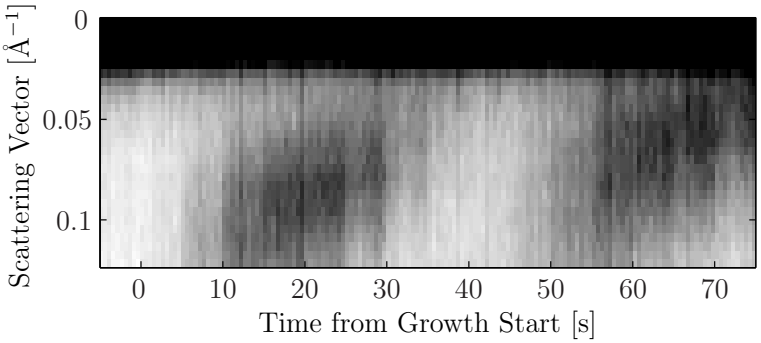

(c) Diffuse Scattering Profile (neg.)

Figure 7.18: Homoepitaxial growth of a $SrTiO_3$ thin film, as observed with the Mixed–Mode PAD. Each peak in the reflected specular beam, panel (a), represents the completion of a single monolayer growth. The accompanying oscillations in the diffuse scattering are shown in integral form in panel (b), while the time evolution of the diffuse scattering profile are shown in panel (c). This last panel is plotted in the negative with the dark strip at the top of the image denoting the location and extent of the specular reflection. Dashed vertical lines are included in panels (a) and (b) to denote new material was deposited.

Mixed–Mode PAD by directly summing the equivalent analog voltage observed in all pixels at the same scattering vector is problematic for the diffuse scattering, because it effectively multiplies any coherent noise, in the image by the number of pixels used in the sum. As the diffuse scattering yield is frequently less than 1 x-ray per pixel per frame, collected over a large area (∼50–100 pixels), even noise correlations smaller than our digitizer resolution may combine to become significant.[18] To minimize the systematic error from correlated noise one may restrict the pixels included in the integration to only those whose measurement contains signal from one or more x-rays.

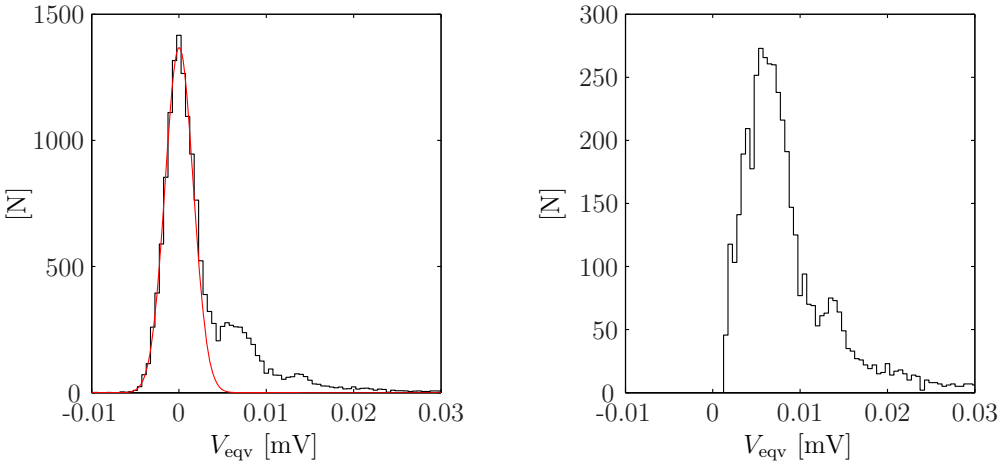

(a) All Pixels (b) Pixels with X-Ray Signal

Figure 7.19: Histogram of the non-specular measurements measurements from a single image within the $SrTiO_3$ homoepitaxial growth series. Panel (a) shows the complete data set along with a fit to the zero x-ray distribution. Panel (b) shows the remaining data following a cut against pixels with no x-ray signal.

[18]Part of the reason why the correlated noise is such an issue for this particular measurement, and a global zero correction as was discussed in section 6.2 does not remove the problem, has to do with the structure of the Mixed–Mode PAD analog readout. This device is laid out in eight banks of sixteen pixels, where the analog residual voltage from eight pixels, one in each bank, are sampled in parallel. Sixteen such samplings reads an entire row of pixels in microseconds, resulting in stronger noise correlations along rows in the imager than columns. Because of the physical orientation of the Mixed–Mode PAD hybrid in these measurements most pixels in the same row will contain data at the same scattering vector, thereby introducing systematic noise as a function of scattering angle. In hindsight, this problem could have been avoided by mounting the detector on its side.

What makes this ultimately possible is the precision with which the Mixed–Mode PAD is able to measure the x-ray charge yield, a consequence both of the quality of the electronics and the high conversion efficiency of direct x-ray detection. Because of this, a histogram of the signal distribution over the entire image, an example of which is shown in figure 7.19, reveals the form of the zero x-ray distribution, *i.e.* the distribution of measurements from pixels containing no x-ray signal. Fitting this distribution with a Gaussian, with care to avoid introducing bias through partial x-ray signal from charge sharing between neighboring pixels, gives a global offset that can be used to reduce the effect of the noise correction. In addition, by comparing the amount of signal observed in each bin of the histogram with the predictions of this fit, it is possible to assign a probability, for each pixel, that the measurement it reports contains signal from an x-ray as opposed to a random noise fluctuation. It is then straightforward to remove the zero x-ray measurements by comparing, on a pixel by pixel basis, a random number drawn from a flat distribution to the probability that the measurement from the pixel contains signal from an x-ray.[19] The data remaining after this cut was then used to determine the scattered intensity as a function of in plane scattering vector.

For the determination of the diffuse scattering profile, the orientation of the x-ray image as well as the center of the diffraction ($\vec{q}_{||} = 0$) were determined through a linear fit to the weighted mean of each detector column, *i.e.* the vertical element in figure 7.17, pixels within the diffuse scattering region were then binned based on the minimal separation of their centers and this fitted line. As the number of elements in each bin could vary, the final parameter reported was the average intensity observed, calculated by integrating the measurements from all

[19]This technique could be improved by applying spatial constraints, such as the presence of an identifiable x-ray signal in a neighboring pixel; however, this is non-trivial to implement and, as the overlap of the x-ray signal and the zero x-ray distribution are sufficiently small and the statistics sufficiently good, this was generally not deemed necessary.

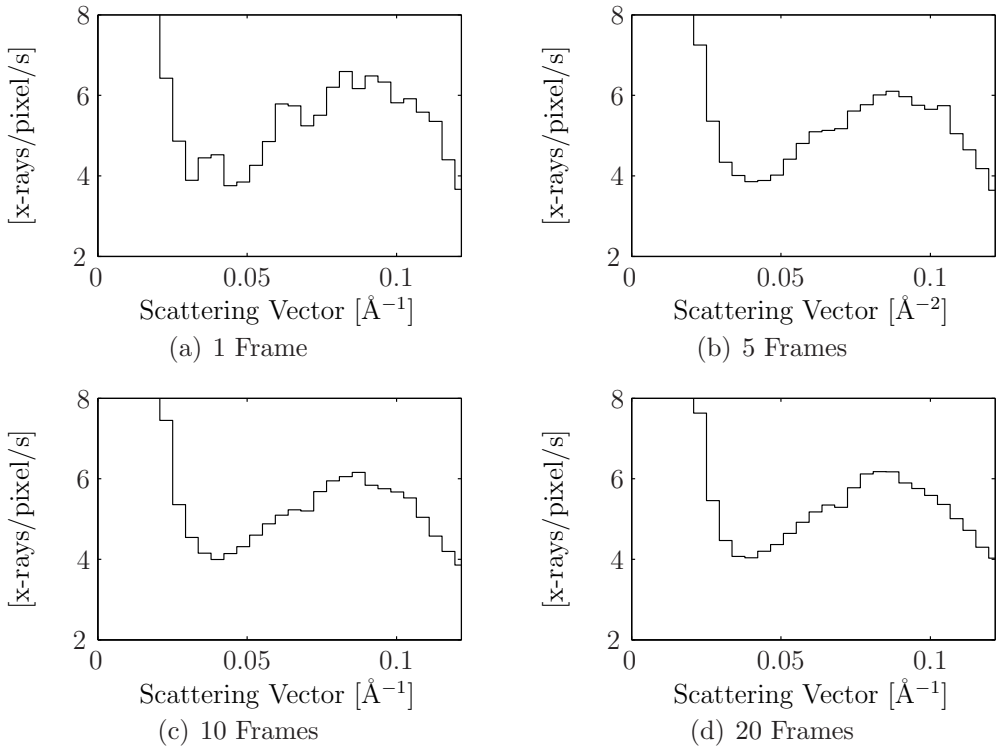

Figure 7.20: Diffuse scattering intensity profile near the first specular intensity minima. Panels (a) through (d) show how the profile improves through merging frames, a combined effect of improved statistics and cancellation of correlated noise effects within each frame.

pixels determined by the method above to hold x-ray signal and dividing this result by the total number of pixels within the range of scattering vectors of the bin. Example results at a portion of the growth are shown in figure 7.20 while results from the full data set are given in panel (c) of figure 7.18. Merging the data in this fashion sufficiently reduces the systematic noise to bring out the diffuse scattering profile, however effects of systematic noise are still evident as it is present in the pixel measurements with x-ray signal. To mitigate these residual effects of coherent noise, it is possible to combine data from multiple frames as exemplified in figure 7.20. Not only does this improve measurement statistics, but coherent noise fluctuations in the independent images will tend to cancel out, leading to notable improvement in the the measurement quality.

The scientific implication of the diffuse diffraction profiles, shown individually in figure 7.20 and as a time series in panel (c) of figure 7.18, is a measurement of the correlation length of structures on the film surface. *A priori*, one cannot tell if these structures represent voids or islands; however, this may be discerned post-growth through the *ex-situ* analysis methods mentioned previously. What this information can directly tell us, though, is if and when intermediate growth structures form on the surface, their scale and proportional distribution, information that is critical for the evaluation of theoretical models of PLD thin film growth.

7.5.4 Mixed–Mode PAD Performance Critique

The PLD studies of $SrTiO_3$ homoepitaxial growth reported here demonstrate both the power and limitations of using the Mixed–Mode PAD for dynamic measurements. As one expects, the Mixed–Mode PAD excelled in measuring the strong, unattenuated, specular reflection of the synchrotron beam. More difficulty, however, is found in measuring the diffuse scattering from the growth due to the very

weak signal levels involved and the need to combine data from a large number of pixels. Because the Mixed–Mode PAD lacks the pre-sampling summation capabilities of a CCD, data from each pixel must be sampled independently and combined offline—a process that makes the system susceptible to even very low levels, smaller than the quantization threshold of the detector, of correlated noise. Still, for a 2D detector making a fundamentally 1D measurement the Mixed–Mode PAD performs quite well.

It is reasonable to ask what would happen if this measurement were performed with a photon-counting PAD, particularly as such a device would eliminate the systematic noise seen in the Mixed–Mode PAD. This is, unarguably, quite an advantage, but it should be noted that it comes at a price in terms of the active area, as the quantum efficiency of photon-counting PADs is suppressed in the charge sharing regions between pixels. Given the very weak diffuse background, this can prove problematic, because correcting for this effectively magnifies the statistical variations in the signal. Because of this, a more accurate result should be possible, in cases where a weak flux is combined with a short integration, from a detector with higher total quantum efficiency. For this to work, though, the accuracy of the analog measurement must be improved.

As discussed in the conclusion to chapter 4, one way to improve this measurement is to "free-wheel" the analog portion of the Mixed–Mode PAD readout. The notion behind this is to use the sample and hold circuit in each pixel to sample the analog data at a much faster than the frame rate. This oversampling offers improved statistics for the weak analog signal and removes short time scale noise correlation effects. As our discussion in the conclusion to section 4.4 argues, development of this capability is non-trivial, as it requires moderately sophisticated digital signal processing by the FPGA supporting Mixed–Mode PAD hybrid—yet

it has the potential to let an analog imager like the Mixed–Mode PAD outperform a photon-counting PAD in weak-flux dynamic measurements.

7.5.5 Prospects for 2D Growth

The interest in thin film growth is not limited to homogeneous monolayer structures, though understanding these systems is, arguably, a prerequisite to understanding the growth of more complex structures. The devices envisioned as applications of this technology are multi-layer, potentially with substructure built into these layers. The continuous dynamics of the growth of such structures cannot be studied by point or strip detectors because it will require both in-plane and out-of-plane degrees of freedom. Techniques for these studies have been developed and active research is currently underway study a variety systems on time scales longer than 1 s [90].

Detailed description of the Grazing Incidence X-Ray Scattering (GISAXS) technique can be found in references [84, 89, 60]. Briefly though, GISAXS is a technique that combines x-ray reflectivity with small angle scattering to provide ensemble structural information about in-plane and out-of-plane ordering of a thin film. It has been used extensively to study the growth of block–copolymer systems [36, 90]. Because these systems have evolution time scales on the order of seconds to hours, dynamical studies can be conducted with the current generation of x-ray detectors. GISAXS has also been used to study PLD films growth [82]. However, these experiments have not been dynamical studies into the surface science of PLD material deposition. Instead, researchers sought to study the evolution of the film as a function of the number laser pulses, allowing the film to equilibrate after each laser pulse and before acquiring each GISAXS image. However, as we have discussed, there are important and poorly understood surface dynamics that occur on short

time scales following each laser pulse that equilibrium studies do not address.

Finally, in addition to the possibility of 2D dynamic information that Mixed–Mode PAD offers, its improved dynamic range and resolving power over the phosphor CCDs that have been used in most GISAXS experiments should yield new information, even on long time scales. As discussed earlier, the images produced in surface scattering experiments often exhibit strong specular reflection and weak diffuse scattering signals. These specular reflections are often orders of magnitude more intense than the diffuse scattering, which becomes a challenge both in terms of detector dynamic range and Point Spread Function (PSF) that is beyond the capabilites of most phosphor-coupled CCD systems. Consequently, these CCD systems rely on beam stops or attenuators to see the weak diffuse signal, a limitation that is not necessary with the Mixed–Mode PAD.

7.6 Conclusion

The five experiments presented in this chapter serve to illustrate how the capabilities of the Mixed–Mode PAD extend beyond that of the current generation of x-ray imagers. Combining this detector's fast framing capabilites and high flux tolerance with its single x-ray sensitivity, pixelation limited spatial resolving power, and substantial dynamic range allows it to acquire data sets that are not possible with current x-ray imaging technologies. As a result, this detector promises not only to improve the quality of data sets gathered in currently possible experiments, but to enable new experiments that were previously impossible.

CHAPTER 8

CONCLUSIONS

The importance of detector development to the advancement of modern science was recognized in the press release announcing the 1992 Nobel Prize in Physics, wherein the Nobel Committee wrote, "The development of detectors very often goes hand in hand with progress in fundamental research."[1] This observation is at the heart of the motivation for developing the Mixed–Mode Pixel Array Detector. Years of investment into synchrotron light sources has yielded sources capable of producing intense x-ray fluxes, setting the stage for experiments that would have been impossible with more conventional sources. The full potential for x-ray science using the very intense x-ray flux from synchrotrons is not, however, currently being realized because of an absence of suitable detectors. Therefore, new x-ray detectors that address the limitations of the current generation of x-ray imagers would enable a broad range of new science.

In the early 1990s, research and development began on a new class of x-ray detectors, Pixel Array Detectors, intended to meet this need. PAD detectors combine direct x-ray detection with custom, in-pixel signal processing electronics. These imagers were initially developed along two paths: Digital Pixel Array Detectors, that use photon discriminating signal processing electronics to count the number of individual photons observed; and Analog Pixel Array Detectors, that accumulate the current generated by the x-ray photons into an analog storage element for post exposure digitization. Digital PADs offer a large dynamic range and very fast readout, yet suffer from a dead time associated with processing each photon that, effectively, imposes a flux limits on these detectors that is low in comparison to the x-ray intensities attainable in many synchrotron experiments. On the other hand,

[1] In 1992 the Nobel Prize in Physics was awarded to Professor Georges Charpak for his invention and development of particle detectors, in particular the multi-wire proportional chamber.

Analog PADs can tolerate the extremely intense x-ray fluxes produced at modern synchrotron light sources but are limited in terms of their dynamic range and, in designs that do not incorporate a high level of parallel digitization and read out, frame rate. The Mixed–Mode PAD merges the integrating front end of an Analog PAD with the digital storage of a Digital PAD to achieve a device that is more than the sum of its components.

In this final chapter, we will briefly review the characteristics that distinguish the Mixed–Mode PAD, both from the current generation of x-ray imagers employed at synchrotron light sources and from other contemporary PAD projects, and the opportunities these characteristics present for science at synchrotron light sources.

8.0.1 Performance Highlights

What enables the Mixed–Mode PAD to extend science at synchrotron light sources is a combination of performance characteristics that is unique to this imager and well suited to brightness of the modern synchrotron. Within the melange of imager properties discussed over the course of this thesis, two classes of distinguishing characteristics can be identified: those that distinguish the PAD methodology from current generations of non-PAD x-ray detectors, and therefore are held in common with all PADs; and those that distinguish the Mixed–Mode PAD from its Analog and Digital PAD contemporaries.

The characteristics of the Mixed–Mode PAD that distinguish it from the current generation of synchrotron imagers, but that it holds in common with all PAD detectors, are benefits it receives from direct x-ray detection. These benefits include a comparatively large signal yielded per x-ray with intrinsically small variation (*e.g.* for a 10 keV x-ray one expects 445 ± 3 aC of charge yield from a PAD as opposed to 2 to 5 ± 2 to 3 aC for a phosphor coupled CCD) that allows the Mixed–Mode PAD to observe quanta of x-rays (as was demonstrated in section

7.1) and Digital PADs to reliably count single x-rays. In addition, direct detection results in an analog impulse response that is constrained to a small spatial region, typically less than the size of a PAD pixel, thus making it possible to resolve very faint x-ray signals in close proximity to intense x-ray signals. Finally, though not a prerequisite, the signal processing electronics of the Mixed–Mode PAD and all contemporary PADs offer an electronic shutter that controls the gating of x-ray signal into the PAD signal processing chain; this, in-turn, frees PAD detectors from the experimental limitations and timing uncertainty introduced by the mechanical shutters required by the current generation of synchrotron imagers.

The characteristics that distinguish the Mixed–Mode PAD from other PADs stem from the signal processing electronics built into each pixel. Foremost among these characteristics are the large well depth (2.6×10^7 x-rays/pixel[2]) and single x-ray sensitivity (through a FWHM read noise of 0.4 x-ray) that give the Mixed Mode PAD a dynamic range of 156 dB; in comparison with 120 dB for the largest dynamic range Digital PAD or 76 dB for the largest dynamic range Analog PAD. Similarly important, particularly for imaging with brief exposures, is the 10^8 x-rays/pixel/s flux tolerance of the Mixed–Mode PAD, two orders of magnitude higher than what can be expected from a contemporary Digital PAD. This high flux tolerance makes more of the imager dynamic range available in brief exposures, as per the discussion of section 7.5. Finally, the capacity of the Mixed–Mode PAD detector hybrids to accurately readout in < 1 ms offers continuous frame rates up to 1 kHz, which is roughly equal to the continuous frame rates that may be obtained from Digital PADs but much faster than all but the most highly parallel of Analog PADs.

[2]All x-ray referenced parameters in this chapter assume an x-ray energy of 10 keV.

8.0.2 Science Opportunities

While the individual performance characteristics noted above all represent significant advancements beyond the capabilities of the current generation of synchrotron imagers, what ultimately distinguishes the Mixed–Mode PAD from contemporary PAD projects is the combination of these characteristics that it offers. This combination places the Mixed–Mode PAD in a unique position to effectively make use of the intense x-ray fluxes available at modern synchrotron light sources to extend x-ray science to new areas.

One very important property that the Mixed–Mode PAD brings to x-ray scattering experiments is its combination of large dynamic range and sub-pixel point spread. In earlier imagers, like phosphor-coupled CCDs, indirect x-ray detection would yield non-negligible signal on mm length scales. This effectively imposes a dynamic range limit on these imagers by making it very difficult to resolve weak signals in the presence of the presence of strong ones. In a direct detection imager, like the Mixed–Mode PAD, the analog response of the detector diode to each x-ray is concentrated to within a sufficiently small spatial region that, so long as there is a pixel separating the weak and intense signals, it is possible to resolve them both. However, to measure both signals requires, in addition to resolving power, sufficient dynamic range to observe both with good statistics. As scattering systems can produce peaks or rings of scattered intensity that differ by many orders of magnitude and synchrotron sources provide sufficient flux to observe these on reasonable time scales, the availability of a detector that offers the dynamic range of the Mixed–Mode PAD with the resolving power of direct detection will reveal new information in currently studied systems while enabling the investigation of new systems are beyond the resolution and dynamic range limits of current detectors.

The potential the Mixed–Mode PAD offers for studying wide-dynamic-range

systems is, however, not limited to static measurements. One of the most exciting application area for the Mixed–Mode PAD is in the imaging of systems that are continuously evolving on ms time scales. Currently, this time scale is almost exclusively the domain of point or 1D detectors, as imagers capable of framing on this time scale are simply not available.

The Digital PADs, currently in development, offer electronic shuttering capabilities and small read out dead time required to frame at the high rate needed for work on ms time scales; however, because of their flux limitations they are limited in the effective dynamic range that they offer for short exposures. The Mixed–Mode PAD also possess electronic shuttering capabilities and small read out dead time but combines these performance characteristics with a flux tolerance two orders of magnitude higher than what one can expect from a well calibrated Digital PAD. Consequently, it is a device that is uniquely well suited to imaging on ms time scales and, therefore, is poised to enable investigation of previously inaccessible dynamic systems with x-rays.

8.0.3 Work Ahead

The work presented in this thesis demonstrated the functionality and capabilities of the Mixed–Mode PAD while illustrating it potential impact on science at synchrotron light sources through a series of demonstration experiments. This work was performed using a single hybrid prototype camera, as was discussed in chapter 5. While this prototype served as a good characterization and demonstration platform, its practical utility is limited by its small active area.

However, as was also discussed in chapter 5, while this characterization work and the inaugural synchrotron experiments using the Mixed–Mode PAD were undertaken at Cornell, our industrial collaborators at ADSC have been working to-

wards developing custom support electronics and a cryostat housing suitable for a large-area (512×512 pixel/4×4 detector hybrid) camera. As this phase of the project is, fundamentally, a commercialization effort, the involvement of the Cornell Detector Development Group is limited to an advisory role.

Tangentially, an effort is underway by the Cornell Detector Development Group to upgrade the support electronics of the single hybrid prototype camera to support four detector hybrids in a 1×4 configuration (128×512 pixel). This imager would have limited utility for application requiring large 2D area coverage, such as the crystallography market that our collaborators at ADSC are targeting. However, it should be very useful when employed at the G-line facility of CHESS, where there is much interest in studying the dynamics of thin film growth processes as well as solution scattering in the small angle regime—systems that require substantial active area in one dimension but much less in the other.

8.0.4 Closing Remarks

It was argued in the introduction to this thesis that, for the development of new instrumentation to be of scientific merit, it should be designed with, and demonstrate an ability to, enable a broad class of new scientific investigation. The Mixed–Mode PAD demonstrably meets these objectives by bridging a portion of the gulf between the capacity of synchrotron light sources to produce intense x-ray fluxes and the capabilities of modern x-ray detectors to measure the resulting signals. In the process, it offers the first opportunities for wide dynamic range, continuous imaging of dynamic phenomena on ms times scales using x-rays.

APPENDIX A

LINEAR FEEDBACK SHIFT REGISTER THEORY

The theory underlying the operation of the pseudorandom counter is the theory of finite fields, which was originally built on a foundation created by the French mathematician Évariste Galois.[1] The circuit itself springs from considerations of the reducibility of polynomials over finite fields; in particular, the binary field $\mathbb{Z}/2\mathbb{Z} = \{0, 1\}$. As this is a topic that has been worthy of many texts, we will only present an overview oriented discussion, sketching out the details as they relate to our particular application. Readers interested in a deeper explanation of finite fields are referenced to [24] as well as [37] for further details on their application to the theory of Linear Feedback Shift Registers (LFSR).

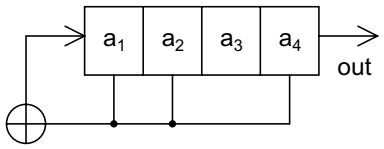

Figure A.1: Example linear feedback shift register.

To illustrate the problem, we begin by looking at an example of an LFSR that would be unsuitable for a counter, because it cannot reach all possible register states and analyzing the properties that indicate this. Consider the three-tap Fibonacci mode LFSR shown in figure A.1. We may write the state of this register

[1] Évariste Galois, 1811-1832. Quite a unique figure in the history of mathematics. While still in his teens, he determined a necessary and sufficient condition for a polynomial to be solvable by radicals, laying the foundation for the branch of Finite Field Theory known as *Galois Theory*. He died from wounds suffered in a duel at the age of twenty; whether the cause of the duel was politics or a matter of the heart is a point of debate.

as a vector

$$\vec{a} = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{pmatrix}, \tag{A.1}$$

where each component (a_1, a_2, a_3, and a_4) is a member of the binary field ($\mathbb{Z}/2\mathbb{Z}$)). Given a state vector $\vec{a}$, to generate the register's next state we apply the generation operator ($\mathcal{G}$),

$$\mathcal{G} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 1 \end{pmatrix}, \tag{A.2}$$

or more generally if $\vec{a}_0$ is the initial state of the system then the k^{th} subsequent state is given by $\vec{a}_k = \mathcal{G}^k \cdot \vec{a}_0$. Since we are working with a nonsingular generator on a finite vector space, for every state there exists an integer p such that $\mathcal{G}^p\vec{a} = \vec{a}$ or equivalently $\mathcal{G}^p = \mathcal{I}$, where $\mathcal{I}$ is the identity matrix. The integer p is said to be the *period* of the generation operator for that vector and the period is termed *maximal* if $p = b^d - 1$, where b base of the finite field comprising the vector space components (2 in the case of the binary field) and d is the dimensionality of the space. Inspection shows that this is one less than the total number of possible states and the largest period achievable, as the null vector will always generate a separate, singular subspace. The splitting of this vector space by the generator defined above is shown in figure A.2. As this figure illustrates, this generator splits the 16 possible states into four subspaces and as such does not yield a maximal period.

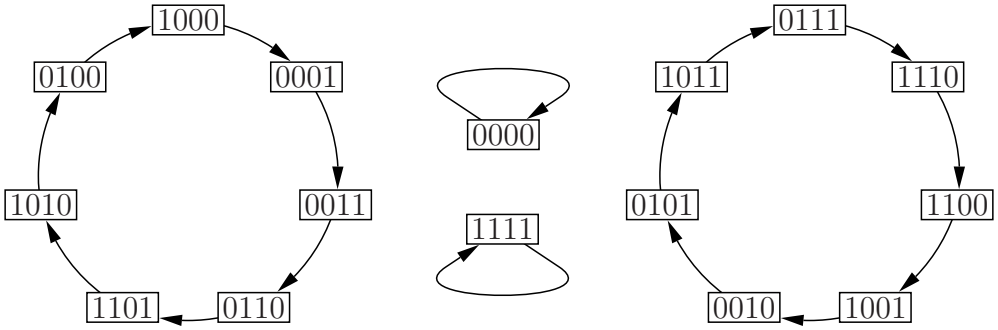

Figure A.2: Graphical descriptions how the generator Ω, as defined in equation A.2, splits vector space of 4-tuples with binary components, $(\mathbb{Z}/2\mathbb{Z}) \otimes (\mathbb{Z}/2\mathbb{Z}) \otimes (\mathbb{Z}/2\mathbb{Z}) \otimes (\mathbb{Z}/2\mathbb{Z})$.

A more direct, though much less illustrative, path to this result comes from considering the characteristic polynomial of the generation operator, $\mathcal{G}$,

$$g(x) = \det(x \cdot \mathcal{I} - \mathcal{G}) \tag{A.3}$$

$$= x^4 - x^3 - x^2 - 1$$

$$= x^4 + x^3 + x^2 + 1$$

$$= (x+1)(x^3 + x + 1), \tag{A.4}$$

where we have used the fact that for $x \in \mathbb{Z}/2\mathbb{Z}$, $x = -x$ to arrive at our final simplification. As in traditional linear algebra, the characteristic polynomial of the generator can reveal much about its behavior. For our purposes, of principle importance is the polynomial's *reducibility*; that is, whether the polynomial may be written as the product of two or more distinct polynomials without remainder. If the characteristic polynomial can be reduced, then the vector space will split under the generator's operation, as our example illustrated, and thus the generator cannot be maximal. If, however, the characteristic polynomial ($g(x)$) is irreducible then Finite Fields Theory assures us that it will divide $f(x) = x^{b^{D(g)}-1} - 1$, where b is again the base of our finite field (2 in this case) and $D(g)$ is the degree of the characteristic polynomial. Now, if g divides f, then any root of g is also a root

of f and since the operator $\mathcal{G}$ is a root of its characteristic polynomial, g, $\mathcal{G}$ must also be a root of f. This is equivalent to saying

$$\mathcal{G}^{b^{D(g)}-1} = \mathcal{I}. \tag{A.5}$$

Unfortunately, we are not guaranteed that $b^{D(g)}-1$ is the smallest integer exponent for which $\mathcal{G}$ raised to that power becomes the identity, only that the smallest integer exponent will divide $b^{D(g)}-1$. Polynomials which are both irreducible and maximal are termed *primitive*. To verify if a particular operator is primitive what remains is to check if g divides $x^k - 1$ for any integer k that divides $b^{D(g)} - 1$. While computationally laborious if done by hand algorithms exist to allow this to be efficiently checked by computer. In this way, analysis of the characteristic polynomial provides a direct and efficient method for determining if a generator has a maximal period.

Although this operator approach offers a pedagogically straightforward means to determine if a LFSR is maximal, most treatments of the Fibonacci mode LFSR architecture do not invoke it as there is a much simpler way of deriving the systems characteristic polynomial. More frequently referred to in the literature as the *connection polynomial*, it is defined as

$$q(x) = \sum_{i=0}^{m} q_i x^{(m-i)}, \tag{A.6}$$

where m is the number of bits in the register, $q_i = 1$ if there is a tap on the i^{th} bit (using the indexing from figures 4.20 and A.1) or 0 otherwise, and $q_0 = 1$ by definition. The connection polynomial is equivalent to the generator operator's characteristic polynomial, and, thus, it may be analyzed analogously to determine if a LFSR has a maximal period.

APPENDIX B

FREQUENCY ANALYSIS OF INTEGRATOR WITH INJECTED CURRENT

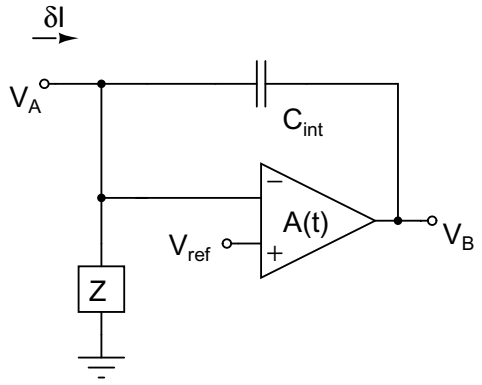

Figure B.1: Model used in current injection analysis.

From figure B.1, the impedance seen looking into node V_{A} is given by

$$\left(Z \middle|\middle| \frac{1}{i\omega C_{\mathrm{int}}(1+A)}\right) = \frac{Z}{1+i\omega C_{\mathrm{int}} Z(1+A)}, \tag{B.1}$$

where A is the amplifier gain at the frequency ω. Then, if a current (δI) is injected into this node at the frequency ω,

$$\delta V_{\mathrm{A}} = \frac{Z}{1+i\omega C_{\mathrm{int}} Z(1+A)} \cdot \delta I, \tag{B.2}$$

and

$$\delta V_{\mathrm{B}} = -A\delta V_{\mathrm{A}} \tag{B.3}$$

$$= -\frac{ZA}{1+i\omega C_{\mathrm{int}} Z(1+A)} \cdot \delta I. \tag{B.4}$$

For the special case of $Z = \frac{1}{i\omega C_{\mathrm{pix}}}$, this result simplifies to

$$\delta V_{\mathrm{B}} = -\frac{A}{i\omega(C_{\mathrm{pix}} + C_{\mathrm{int}}(1+A))} \cdot \delta I. \tag{B.5}$$

APPENDIX C

ASIC SUBMISSION HISTORY

All submissions were fabricated using the TSMC 0.25 μm process, contracted through the Metal Oxide Semiconductor Implementation Service (MOSIS) service. Each submission is accompanied by release documentation, prepared by Skip Augustine, containing full schematics a list of design changes. This documentation is on record with the Cornell X-Ray Detector Development Group.

Table C.1: Mixed–Mode PAD prototyping submission history.

Submission ID	Date	Comment
Cornell.A	Nov. 2003	Test structures
AE176	Jan. 2004	16×128 pixel array
AE180	Mar. 2004	16×16 pixel array
Cornell.B	May 2004	Test structures
AE184	Aug. 2004	16×16 pixel array
AE190	Apr. 2005	16×128 pixel array
AE196	Aug. 2005	16×16 pixel array
AE203	Nov. 2005	16×16 pixel array
AE207	Feb. 2006	Final 128×128 pixel array

BIBLIOGRAPHY

[1] Jens Als-Nielsen and Des McMorrow, *Elements of modern x-ray physics*, John Wiley & Sons, New York, 2001.

[2] G. Anelli, M. Campbell, M. Delmastro, F. Faccio, S. Floria, A. Giraldo, E. Heijne, P. Jarron, K. Kloukinas, A. Marchioro, P. Moreira, and W. Snoeys, *Radiation tolerant vlsi circuits in standard deep submicron cmos technologies for the lhc experiments: Practical design aspects*, IEEE Transactions on Nuclear Science **46** (1999), no. 6, 1690–6.

[3] Giovanni M. Anelli, *Conception et caracterisation de circuits integres resistants aux radiations pour les detecteurs de particules du lhc en technologies cmos submicroniques profondes*, Ph.D. thesis, Institut National Polytrchnique de Grenoble, 2000.

[4] A.G. Angello, F. Augustine, A. Ercan, S. Gruner, R. Hamlin, T. Hontz, M. Renzi, D. Schuette, M. Tate, and W. Vernon, *Development of a mixed-mode pixel array detector for macromolecular crystallography*, IEEE Nuclear Science Symposium Conference Record (2004).

[5] George Arfken, *Mathematical methods for physicists*, Academic Press, 1966.

[6] N. Aschcroft and N. Mermin, *Solid state physics*, Harcourt, New York, 1976.

[7] Skip Augustine, *Release notes: Ae2007 – 128x128 mixed–mode pixel x-ray detector asic*, Tech. Report 1, Augustine Engineering, 2006.

[8] B. Balland and G. Barbottin, *Trapping and detrapping kinetics impact on c(v) and i(v) curves*, Instabilities in Silicon Devices: New Insulators, Devices and Radiation Effects, Vol. II (G. Barbottin and A. Vapaille, eds.), Elsevier, 1989, pp. 7–81.

[9] Gérard Barbottin and André Vapaille (eds.), *Instabilities in silicon devices: Silicon passivation and related instabilities*, vol. 2, Elsevier, New York, 1989.

[10] Gérard Barbottin and André Vapaille (eds.), *Instabilities in silicon devices: New insulators, devices and radiation effects*, vol. 3, Elsevier, Amsterdam, 1999.

[11] S. L. Barna, M. W. Tate, S. M. Gruner, and E. F. Eikenberry, *Calibration proceedures for charge-coupled device x-ray detectors*, Review of Scientific Instruments **70** (1999), no. 7, 2927–34.

[12] S. Basolo, J.-F. Berar, N. Boudet, P. Breugnon, B. Caillot, J.-C. Clemens, P. Delpierre, B. Dinkespiler, I. Koudobine, Ch. Meessen, M. Menouni, C. Mouget, P. Pangaud, R. Potheau, and E. Vigeolas, *Xpad : Pixel detector for material sciences*, IEEE Transactions on Nuclear Science **52** (2005), no. 5, 1994–8.

[13] M. W. Bautz, G. E. Berman, J. P. Doty, and G. R. Ricker, *Charge-coupled-device x-ray detector performance model*, Optical Engineering **26** (1987), no. 8.

[14] G. Bertolini and A. Coche (eds.), *Semiconductor detectors*, North-Holland, Amsterdam, 1968.

[15] Carl Branden and John Tooze, *Introduction to protien structure*, 2nd. ed., Garland Publishing, Inc., New York, 1998.

[16] D. Bräunig and F. Wulf, *Radiation effects in electronic components*, Instabilities in Silicon Devices: New Insulators, Devices and Radiation Effects, Vol. III (G. Barbottin and A. Vapaille, eds.), Elsevier, 1999, pp. 639–722.

[17] Ch. Broennimann, E. F. Eikenberry, B. Henrich, R. Horisberger, G. Huelsen, E. Pohl, B. Schmitt, C. Schulze-Briese, M. Suzuki, T. Tomizaki, H. Toyokawa, and A. Wagner, *The pilatus 1m detector*, Journal of Synchrotron Radiation **13** (2006), 120–30.

[18] Barry E. Burke, John Tonry, Michael Cooper, Gerard Luppino, George Jacoby, Richard Bredthauer, Kasey Boggs, Michael Lesser, Peter Onaka, Douglas Young, Peter Doherty, and David Craig, *The orthogonal-transfer array: a new ccd architecture for astronomy*, Optical and Infrared Detectors for Astronomy, Proceedings of the SPIE, vol. 5499, September 2004, pp. 185–92.

[19] M. Campbell, G. Anelli, E. Cantatore, F. Faccio, E. H. M. Heijne, P. Jarron, J.-C. Santiard, W. Snoeys, and K. Wyllie, *An introduction to deep submicron cmos for vertex applications*, Nuclear Instruments and Methods in Physical Research A **473** (2001), no. 1–2, 140–5.

[20] CHESS, *Online cornell high energy synchrotron source (chess) facility description*, http://www.chess.cornell.edu, 2008.

[21] D. Chrisey and G.Hubler (eds.), *Pulsed laser deposition of thin films*, John Wiley & Sons, New York, 1994.

[22] Lilián del Risco Norrlid, *Modeling the performance of an x-ray imaging detector*, Ph.D. thesis, Uppsala University, 2004.

[23] C. M. Dozier, D. M. Fleetwood, D. B. Brown, and P. S. Winokur, *An evaluation of low-energy x-ray and cobolt-60 irradiation of mos transistors*, IEEE Transactions on Nuclear Science **35** (1987), no. 6, 1535–9.

[24] D. S. Dummit and R. M. Foote, *Abstract algebra*, 2nd. ed., Prentice–Hall, New Jersey, 1999.

[25] A. Ercan, M. W. Tate, and S. M. Gruner, *Analog pixel array detectors*, Journal of Synchrotron Radiation (2006).

[26] Alper Ercan, *Development of a high frame rate x-ray imager*, Ph.D. thesis, Cornell University, Ithaca, NY, 2007.

[27] Shyamsunder Erramilli, Frederik Osterberg, Sol M. Gruner, Mark W. Tate, and Manfred Kriechbaum, *Time-resolved x-ray studies of pressure-jump-induced topological transitions in biological membranes*, vol. 2521, SPIE, 1995, pp. 188–96.

[28] Federico Faccio and Giovanni Cervelli, *Radiation–induced edge effects in deep submicron cmos transistors*, IEEE Transactions on Nuclear Science **52** (2005), no. 6, 2413–20.

[29] U. Fano, *Ionization yield of radiations. ii. the fluctuations of the number of ions*, Physical Review **72** (1947), no. 1, 26–9.

[30] A. Fant, M.J. French, P. Murray, P. Seller, S. Manolopoulos, and S.L. Thomas, *Silicon-based photon counting x-ray detector for synchrotron applications*, Nuclear Instruments and Methods in Physical Research A **535** (2004), 442–7.

[31] John C. Feltz, *Development of the modulation transfer function and contrast transfer function for discrete systems, particularly charge–coupled devices*, Optical Engineering **29** (1990), no. 8, 893–904.

[32] A. Fleet, D. Dale, Y. Suzuki, and J.D. Brock, *Observed effects of a changing step-edge density on thin-film growth dynamics*, Physical Review Letters **94** (2005), p. 036102.1–4.

[33] Aaron Fleet, *Time resolved x-ray scattering during strontium titanate homoepitaxy by pulsed laser deposition*, Ph.D. thesis, Cornell University, 2005.

[34] Boyd Fowler, *Cmos area image sensors with pixel level a/d conversion*, Ph.D. thesis, Stanford University, 1995.

[35] Elspeth F. Garman and Thomas R. Schneider, *Macromolecular cryocrystallography*, Journal of Applied Crystallography **30** (1997), 211–37.

[36] A. Gibaud, S. Dourdain, O. Gang, and B. Ocko, *In-situ grazing incidence small-angle x-ray scattering real-time monitoring of the role of humidity during the structural formation of templated silica thin films*, Physical Review B (Condensed Matter and Materials Physics) **70** (2004), no. 16.

[37] Solomon W. Golomb, *Shift register sequences*, revised ed., Aegean Park Press, Laguna Hills, CA, 1982.

[38] Paul R. Gray, Paul J. Hurst, Stephen H. Lewis, and Robert G. Meyer, *Analysis and design of analog integrated circuits*, 4th ed., John Wiley & Sons, New York, 2001.

[39] S. M. Gruner, E. F. Eikenberry, and M. W. Tate, *Comparison of x-ray detector*, International Tables For Crystallogrpahy, Vol. F (Michael G. Rossmann and Eddy Arnold, eds.), Kluwer Academic Publishers, 2001, pp. 143–53.

[40] Sol M. Gruner, Don Bilderback, and Maury Tigner, *Synchrotron radiation sources for the future: White paper*, online, April 2001.

[41] Sol M. Gruner, J. R. Milch, and Geo. T. Reynolds, *Evaluation of area photon detectors by a method based on detective quantum efficieny (dqe)*, IEEE Transactions on Nuclear Science **NS–25** (1978), no. 1, 562–4.

[42] Sol M. Gruner, James R. Milch, and George T. Reynolds, *Survey of two-dimensional electro-optical x-ray detectors*, Neclear Instruments and Methods **195** (1982), 287–97.

[43] Jr. H. E. Boesch and J. M. McGarrity, *Charge yield and dose effects in mos capacitors at 80 k*, IEEE Transactions on Nuclear Science **NS-23** (1976), no. 6, 1520–5.

[44] Fredric J. Harris, *On the use of windows for harmonic analysis with the discrete fourier transform*, Proceedings of the IEEE **66** (1978), no. 1, 51–83.

[45] J. R. Helliwell, S. Ealick, P. Doing, T. Irving, and M. Szebenyi, *Towards the measurement of ideal data for macromolecular crystallography unsing synchrotron sources*, Acta Crystallographica D **D49** (1993), 120–8.

[46] B. L. Henke, E. M. Gullikson, and J. C. Davis, *X-ray interactions: photoabsorption, scattering, transmission, and reflection at e=50-30000 ev, z=1-92*, Atomic Data and Nuclear Data Tables **54** (1993), no. 2, 181–342.

[47] Paul Horowitz and Winfield Hill, *The art of electronics*, 2nd. ed., Cambridge University Press, Cambridge, MA, 1989.

[48] Gregor Hülsen-Bollier, *The pilatus 1m detector a novel large area pixel detector*, Ph.D. thesis, University of Erlangen–Nürnberg, PSI, 2005.

[49] C. Jacoboni, C. Canali, G. Ottaviani, and A. Alberigi Quaranta, *A review of some charge transport properties of silicon*, Solid-State Electronics **20** (1977), 77–89.

[50] James Janesick, James Andrews, John Tower, Tom Elliott, John Cheng, Michael Lesser, and Jeff Pinter, *Fundamental performance differences between cmos and ccd images; part i*, Proceedings of the SPIE (2007).

[51] James R. Janesick, *Scientific charge coupled devices*, SPIE Press, 2001.

[52] Jr. John D. Ferguson, Arthur R. Woll, Gokhan Arikan, Darren S. Dale, Aram Amassian, Mark W. Tate, and Joel D. Brock, *Time resolved in-situ diffuse x-ray scattering measurements of the surface morphology oh homoepitaxial srtio3 films during pulsed laser deposition*, Material Research Society Symposium Proceedings **1024** (2008).

[53] Howard Johnson and Martin Graham, *High speed signal propogation, advanced black magic*, Prentice-Hall, New Jersey, 2003.

[54] R. Clark Jones, *On the point and line spread functions of photographic images*, Journal of the Optical Society of America **48** (1958), no. 12, 934–7.

[55] Scott Kirkpatrick, *Modeling diffusion and collection of charge from ionizing radiation in silicon devices*, IEEE Transactions on Electron Devices **ED-26** (1979), no. 11, 1742–53.

[56] Glenn F. Knoll, *Radiation detection and measurement*, 2nd. ed., John Wiley & Sons, New York, 1989.

[57] Lucas Koerner, *Radiation damage testing of mmpad ae190*, Internal note, Gruner Biophysics Group, Cornell University., September 2005.

[58] R. Lacoe, J. Osborn, R. Koga, S. Brown, and D. Mayer, *Application of hardness–by–design methodology to radiation–tolerant asic technologies*, IEEE Transactions on Nuclear Science **47** (2000), no. 6.

[59] R. C. Lacoe, J. V. Osborn, D. C. Mayer, S. Brown, and J. Gambles, *Total–dose tolerance of the commercial taiwan semiconductor manufacturing company (tsmc) 0.35–μm cmos process*, 2001 IEEE Radiation Effects Data Workshop, IEEE Nuclear and Space Radiation Effects Conference, IEEE, IEEE Press, 2001.

[60] R. Lazzari, *Isgisaxs: A program for grazing-incidence small angle x-ray scattering analysis of supported islands*, Journal of Applied Crystallography **35** (2002), no. 4, 406–21.

[61] Rob Lewis, *Position sensitive detectors for synchrotron radiation studies: The tortoise and the hare?*, Nuclear Instruments and Methods in Physical Research A **513** (2003), 172–7.

[62] F. Li and A. Nathan, *Ccd image sensors on deep-ultraviolet: Degradation behavior and damage mechanisms*, Microtechnology and MEMS, Springer, New York, 2005.

[63] Bernabé Linares-Barranco and Teresa Serrano-Gotarredona, *On the design and characterization of femtoampere current–mode circuits*, IEEE Journal of Solid State Circuits **38** (2003), no. 8, 1353–1363.

[64] N.V. Loukianova, H.O. Folkerts, J.P.V. Maas, D.W.E. Verbugt, A.J. Mierop, W. Hoekstra, E. Roks, and A.J.P. Theuwissen, *Leakage current modeling of test structures for characterization of dark current in cmos image sensors*, IEEE Transactions on Electron Devices **50** (2003), no. 1, 77–83.

[65] Andrew G. MacPhee, Mark W. Tate, Christopher F. Powell, Yong Yue, Matthew J. Renzi, Alper Ercan, Suresh Narayanan, Ernest Fontes, Jochen Walther, Johannes Schaller, Sol M. Gruner, and Jin Wang, *X-ray imaging of shock waves generated by high-pressure fuel spray*, Science **295** (2002), 1261–3.

[66] Stéphane Mallat, *A wavelet tour of signal processing*, 2nd. ed., Academic Press, 1999.

[67] Franco Maloberti, *Analog design for cmos vlsi systems*, Kluwer Academic, Boston, MA, 2001.

[68] P. J. McWhorter, S. L. Miller, and W. M. Miller, *Modeling the anneal of radiation–induced trapped holes in varying thermal environment*, IEEE Transactions on Nuclear Science **37** (1990), no. 6, 1682–9.

[69] C. A. Mead and T. Delbrück, *Scanners for visualizing activity of analog vlsi circuitry*, Analog Integrated Circuits and Signal Processing **1** (1991), no. 2, 93–106.

[70] George C. Messenger and Milton S. Ash, *The effects of radiation on electronic systems*, Van Nostrand Reinhold Company, 1986.

[71] B. Mikulec, M. Campbell, E. Heijne, X. Llopart, and L. Tlustos, *X-ray imaging using single photon processing with semiconductor pixel detectors*, Nuclear Instruments and Methods in Physical Research A **511** (2003), 282–286.

[72] P. Paillet and J. L. Leray, *Defects and radiation–induced charge trapping phenomena in silica*, Instabilities in Silicon Devices: New Insulators, Devices and Radiation Effects, Vol. III (G. Barbottin and A. Vapaille, eds.), Elsevier, 1999, pp. 723–80.

[73] P. Pangaud, S. Basolo, B. Chantepie, J.C. Clemens, P. Delpierre, B. Dinkespiler, M. Menouni, A. Bonissent, F. Debarbieux, and C. Morel, *First results of xpad3, a new photon counting chip for x-ray ct-scanner with energy discrimination*, IEEE Nuclear Science Symposium Conference Record, IEEE, 2007, pp. 14–8.

[74] Stephen K. Park, Robert Schowengerdt, and Mary-Anne Kaczynski, *Modulation-transfer-function analysis for sampled image systems*, Applied Optics **23** (1984), no. 15, 2572–82.

[75] George G. Pavlov and John A. Nousek, *Charge diffusion in ccd x-ray detectors*, Nuclear Instruments and Methods in Physical Research A **428** (1999), 348–66.

[76] Hugh T. Philipp, Lucas J. Koerner, Marianne Hromalik, Mark W. Tate, and Sol M. Gruner., *Pixel array detector for the capture of femtosecond duration x-ray images*, Optics & Photonics: Photonic Devices + Applications, SPIE, SPIE, 2007.

[77] L. Pollack, M.W. Tate, A.C. Finnefrock, C. Kalidas, S. Trotter, N.C. Darnton, L. Lurio, R.H. Austin, C.A. Batt, S.M. Gruner, and S.G.J. Mochrie, *Time-resolved collapse of a folding protein*, Physical Review Letters **86** (2001), 4962–5.

[78] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery, *Numerical recipies in c, the art of scientific computing*, 2nd. ed., Cambridge University Press, 1992.

[79] T. F. J. Quinn, *X-rays, electrons, and crystalline materials*, Butterworth and Co., London, 1970.

[80] Manijeh Razeghi, *Fundamentals of solid state engineering*, Springer-Verlag, New York, 2002.

[81] Stephen E. Reichenbach, Stephen K. Park, and Ramkumar Narayanswamy, *Characterizing digital image acquisition devices*, Optical Engineering **30** (1991), no. 2, 170–7.

[82] G. Renaud, R. Lazzari, C. Revenant, A. Barbier, M. Noblet, O. Ulrich, F. Leroy, J. Jupille, Y. Borensztein, C. Henry, J. Deviile, F. Scheurer, J. Mane-Mane, and O. Fruchart, *Real-time monitoring of growing nanoparticles*, Science **300** (2003), no. 5624, 1416–9.

[83] Matthew Joseph Renzi, *Pixel array detectors for ultra-fast time-resolved x-ray imaging*, Ph.D. thesis, Cornell University, Ithaca, NY, 2003.

[84] C. Revenant, F. Leroy, R. Lazzari, G. Renaud, and C. Henry, *Quantitative analysis of grazing incidence small–angle scattering: Pd/mgo(001) growth*, Physical Review B (Condensed Matter and Materials Physics) **69** (2004), no. 3.

[85] G. Rossi, M. Renzi, E. F. Eikenberry, M. W. Tate, D. Bilderback, E. Fontes, R. Wixted, S. Barna, and S. M. Gruner, *Development of a pixel array detector for time resolved x-ray imaging*, Synchrotron Radiation Instrumentation: Eleventh US National Conference, American Institute of Physics, AIP, October 1999.

[86] ———, *Tests of a prototype pixel array detector for microsecond time-resolved x-ray diffraction*, Journal of Synchrotron Radiation **6** (1999), 1096–1105.

[87] L. Rossi, P. Fischer, T. Rohe, and N. Wermes, *Pixel detectors*, Springer, New York, 2006.

[88] K. Scheidt, *Review of streak cameras for accelerators: Features, applications and results*, Proceedings of the European Particle Accelerator Conference, 2000, pp. 182–5.

[89] S. Sihna, E. Sirota, S. Garoff, and H. Stanley, *X-ray and neutron scattering from rough surfaces*, Physical Review B (Condensed Matter and Materials Physics) **38** (1988), no. 4, p.2297–311.

[90] D. Smilgies, P. Busch, D. Posselt, and C. Papadakis, *Characterization of polymer thin films with small-angle x-ray scattering under grazing incidence (gisaxs)*, Synchrotron Radiation News **15** (2002), no. 5, 35–42.

[91] Howard M. Smith and A. F. Turner, *Vacuum deposited thin films using a ruby laser*, Applied Optics **4** (1965), 147–8.

[92] W. Snoeys, T. Gutierrez, and G. Anelli, *A new nmos layout structure for radiation tolerance*, IEEE Transactions on Nuclear Science **49** (2002), no. 4.

[93] S. M. Sze, *Semiconductor devices*, 2nd. ed., John Wiley & Sons, Inc., 2001.

[94] T-0250MM-TM-001, *Tsmc 0.25um mixed-signal 2p5m pip or 1p5m+mim salicide 2.5v/5.0v or 2.5v/3.3v design guideline*, Tech. report, TSMC Co., LTD, 2000.

[95] Mark W. Tate, Darol Chamberlain, and Sol M. Gruner, *Area x-ray detector based on a lens-coupled charged-coupled device*, Review of Scientific Instruments **76** (2005), 81301–1–10.

[96] Mark W. Tate, Sol M. Gruner, and Eric F. Eikenberry, *Coupled format variations in x-ray detectors based on charge coupled devices*, Review of Scientific Instrumentation **68** (1997), no. 1, 47–54.

[97] Yuan Taur and Tak H. Ning, *Fundamentals of modern vlsi devices*, Cambridge University Press, 1998.

[98] Barry N. Taylor, *Nist special publication 811: Guide for the use of the international system of units (si)*, Tech. report, National Institute of Standards and Technology, Washington, April 1995.

[99] A. Thompson, D. Mills, S. Naday, S. Gruner, P. Siddons, J. Arthur, R. Wehlitz, and H. Padmore, *A program in detector development for the us synchrotron radiation community*, Tech. Report LBNL–48596, Lawrence Berkeley National Lab., CA, July 2001.

[100] Lukas Tlustos, *Performance and limitations of high granularity single photon processing x-ray imaging detectors*, Ph.D. thesis, Technischen Universität Wien, March 2005.

[101] L. K. Townsley, P. S. Broos, G. Chartas, E. Moskalenko, J. A. Nousek, and G. G. Pavlov, *Simulating ccds for the* Chandra *advanced ccd imaging spectrometer*, arXiv:astro-ph/0111003v1, October 2001.

[102] Raoul Velazco, Pascal Fouillat, and Ricardo Reis (eds.), *Radiation effects on embedded systems*, Springer, Dordrecht, The Netherlands, 2007.

[103] Wilfried von Ammon and Heinz Herzer, *The production and availability of high restivity silicon for detector application*, Nuclear Instruments and Methods in Physical Research **226** (1984), 94–102.

[104] Norbert Wermes, *Pixel vertex detectors*, arXiv:physics/0611075v1, 2006.

[105] Bernard Widrow and Samuel Stearns, *Adaptive signal processing*, Prentice-Hall Signal Processing Series, Prentice-Hall, New Jersey, 1985.

[106] Herman Winick, *Synchrotron radiation sources - present capabilities and future directions*, Journal of Synchrotron Radiation **5** (1998), no. 3, 168–75.

[107] D. Yang, B. Fowler, and A. El Gamal, *A nyquist rate pixel level adc for cmos image sensors*, IEEE Journal of Solid State Circuits (1999).

[108] Mirang Yoon, S. G. J. Mochrie, M. W. Tate, S. M. Gruner, and E. F. Eikenberry, *Anisotropic coarsening of periodic grooves: Time-resolved x-ray scattering*, Phys. Rev. Lett. **80** (1998), no. 2, 337–40.