ABSTRACT

Title of dissertation: BIOLOGICALLY-INSPIRED LOW-LIGHT VISION SYSTEMS FOR MICRO-AIR VEHICLE APPLICATIONS

Andrew Berkovich, Doctor of Philosophy, 2017

Dissertation directed by: Professor Pamela Abshire
Department of Electrical Engineering

Various insect species such as the *Megalopta genalis* are able to visually stabilize and navigate at light levels in which individual photo-receptors may receive fewer than ten photons per second. They do so in cluttered forest environments with astonishing success while relying heavily on optic flow estimation. Such capabilities are nowhere near being met with current technology, in large part due to limitations of low-light vision systems.

This dissertation presents a body of work that enhances the capabilities of visual sensing in photon-limited environments with an emphasis on low-light optic flow detection. We discuss the design and characterization of two optical sensors fabricated using complementary metal-oxide-semiconductor (CMOS) very large scale integration (VLSI) technology. The first is a frame-based, low-light, photon-counting camera module with which we demonstrate 1-D non-directional optic flow detection with fewer than 100 photons/pixel/frame. The second utilizes adaptive analog circuits to improve room-temperature short-wave infrared sensing capabilities. This work demonstrates a reduction in dark current of nearly two orders of magnitude and an improvement in signal-to-noise ratio of nearly 40dB when compared to similar, non-adaptive circuits. This dissertation also presents a novel simulation-based framework that enables benchmarking of optic flow algorithms in photon-limited environments. Using this framework we

compare the performance of traditional optic flow processing algorithms to biologically-inspired algorithms thought to be used by flying insects such as the *Megalopta genalis*. This work serves to provide an understanding of what may be ultimately possible with optic flow sensors in low-light environments and informs the design of future low-light optic flow hardware.

BIOLOGICALLY-INSPIRED LOW-LIGHT VISION SYSTEMS FOR MICRO-AIR VEHICLE APPLICATIONS

By

Andrew Samuel Berkovich

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2017

Advisory Committee:
Professor Pamela Abshire, Chair
Professor Timothy K. Horiuchi
Professor Martin Peckerar
Dr. Geoffrey L. Barrows
Professor Yiannis Aloimonos

Dedication

I dedicate this dissertation to my family. Few people are able to enjoy the type of endless love and support that I have received throughout my life. I owe everything to them and am truly fortunate to have them in my corner.

Acknowledgments

First, I would like to thank my adviser Dr. Pamela Abshire for her continuing support during my Ph.D. research. She set high expectations and challenged me to evolve and improve as a researcher throughout my PhD studies. I thank her for caring enough about me to do so and for giving me the chance to take part in such exciting research. I certainly would not be where I am today without her guidance. She is a kind-hearted individual and I owe her my endless gratitude.

I would like to thank all of my committee members, Dr. Timothy Horiuchi and Dr. Martin Peckerar, Dr. Geoffrey L. Barrows, and Dr. Yiannis Aloimonos for agreeing to serve on my examination committee and taking time out of their busy schedules to meet with me and engage in valuable discussions.

A special thanks goes out to Dr. Geoffrey Barrows who has served me both as collaborator and as wonderful mentor. He taught me a great deal about how to effectively manage time and resources, and helped me grow as a researcher. I could not have asked for a better collaborator throughout my PhD research and his seemingly endless supply of new demos and projects helped inspire and excite me to get back into the lab and work on my research.

Thanks to all the members of the Integrated Biomorphic Information Systems (IBIS) Laboratory, both past and present, for helping cultivate a collaborative, open, and exciting atmosphere in the lab. Special thanks to Timir Datta, Tsung-Hsueh Lee, Marc Dandin, and Babak Nouri for laying a foundation for me to built upon, for so many thought provoking and informative discussions, and for all of their help during my time in the IBIS Laboratory. Thanks to Alex Castro for helping with some of the CTIA data collection, to Sheung Lu for help setting up optical experiments, and Bathiya Senevirathna for help with Cadence and general troubleshooting.

I also must thank Dr. Agis Iliadis, my first graduate mentor. In large part, my journey into graduate school at UMD occurred thanks to his support and encouragement. I still remember taking ENEE313 with him back in 2010 and getting my first taste of research in his lab. I can not thank him enough for his help, support, and kindness over the years.

Contents

List of Tables

List of Figures

List of Abbreviations

ADC	Analog-to-Digital Converter
APD	Avalanche Photodiode
APS	Active Pixel Sensor
AQC	Active Quench Circuit
AR	Anti-reflection
BJT	Bipolar Junction Transistor
BM	Block Matching
CCD	Charge Couple Device
CMOS	Complementary Metal-Oxide Semiconductor
CTIA	Capacitive Trans-Impedance Amplifier
DAC	Digital-to-Analog Converter
DCR	Dark Count Rate
EEPROM	Electrically Erasable Programmable Read-Only Memory
EHP	Electron-Hole Pair
EMCCD	Electron Multiplying Charge-Couple Device
EMD	Elementary Motion Detector
ES	Exhaustive Search
FG	Floating Gate
FOV	Field-Of-View
FPAA	Field Programmable Analog Array
GPS	Global Positioning System
HPF	High-Pass Filter
HR	Hassenstein-Reichardt
IIA	Image Interpolation Algorithm
IM	Intensity Measurement
IMU	Inertial Measurement Unit
LED	Light Emitting Diode
LPF	Low-Pass Filter
MAV	Micro Air Vehicle
MOS	Metal-Oxide Semiconductor
NDM	Non-Directional Multiplication
NDS	Non-Directional Summation
OF	Optic Flow
PDE	Photon Detection Efficiency
PEB	Premature Edge Breakdown
PMT	Photomultiplier Tube
PQC	Passive Quench Circuit
PWM	Pulse-Width Modulated
QD	Quantum Dot
ROIC	Readout Integrated Circuit
SF	Source Follower
SiPM	Silicon Photomultiplier
SLAM	Simultaneous Localization and Mapping
SNR	Signal-to-Noise Ratio
SPAD	Single-Photon Avalanche Diode
SWAP	Size, Weight, and Power

SWIR	Short-Wave Infrared
UAV	Unmanned Aerial Vehicles
UV	Ultra-Violet
WFI	Wide Field Integration

Chapter 1

Introduction

Over the past decade, vast improvements in the cost, size, and functionality of unmanned aerial vehicles (UAVs) have enabled them to become low-cost, commercially available platforms. As UAVs continue to shrink to the micro and nano air vehicle (MAV/NAV) regime [1], new applications in indoor, urban, cluttered, low altitude, and/or confined environments become accessible for this technology. Such platforms are well suited for intelligence, surveillance, and reconnaissance applications that require rapid sensing in environments that are inaccessible and/or dangerous to humans or ground-based robots. This shift towards near-earth flight has created a new set of technological challenges and opportunities, particularly in flight control, visual perception, and low-power systems.

Major challenges are the size, weight, and power (SWaP) restrictions placed on micro-scale vehicles. Popular MAV platforms such as the Crazyflie have a maximum payload weight of 10-20 grams and a wingspan of 10cm. These restrictions severely limit the size of batteries they carry and place power consumption at a premium. This also limits the size of sensing and computing systems placed on-board MAVs and demands highly integrated, compact systems.

Shifting from high-altitude flight to low-altitude flight increases the complexity of navigation and path planning. High-altitude flight can rely on *a priori* path planning using GPS guidance and simple IMU-based flight controllers. Because high-altitude environments are fairly static and obstacles are fairly sparse, challenges like obstacle avoidance and real-time path planning can largely be ignored.

Near-earth flight is quite different. In many low-altitude environments such as urban cities, GPS is susceptible to jamming and is generally unreliable. Typical GPS update rates vary between 1 Hz and 10 Hz, which is slow compared to the flight dynamics of MAVs. Environments are complex and dynamic which makes navigation difficult and real-time path planning a requirement. Even low-level tasks such as hover-in-place and obstacle avoidance are difficult in these conditions. Ultimately, these platforms demand on-board sensing so they are able to perceive and interact with their environment.

Optical sensing offers one of the most promising technologies to solve this challenge. For decades, CCD and CMOS imaging arrays have offered sufficiently high frame rates and spatial resolution for UAV applications. They can be paired with mature computer vision techniques to extract visual information and perform tasks such as optic flow estimation and SLAM.

However, a major limitation of current vision systems is their performance in photon-limited environments. As a result, autonomous air vehicles are not able to effectively operate outdoors at night, indoors in unlit environments, or when flying at very high speeds where image exposure times must be very short to prevent image blurring. These applications can all be classified as photon limited environments in which thousands of photons, or fewer, are incident on a photodetector each second. To date, few technologies exist that enable these capabilities. The focus of this work is to develop image sensor technology for this application.

1.1 Biological Inspiration

The rich diversity of flying organisms on this planet are able to solve many of the same problems that exist in MAV technologies. Insects, for example, are constrained to visual, chemical, and mechanical sensory inputs, extremely low power budgets, and are not equipped with radio communication, radar, or GPS. Using these sensory inputs, they are able to operate autonomously while finding food and mates, avoiding predation, navigating to and from nests, and doing all of this in the cluttered, low-light environments of dense forests at night. In addition, many of these airborne insects accomplish these tasks with just hundreds or thousands of photons each second [2] using visual sensory systems that are well-adapted for photon-limited environments.

A great example of such nocturnal flying insects is the Central American sweat bee *Megalopta genalis*. It emerges from its nest at late twilight and early dusk to forage for food. At these hours, light levels are on the order of 1 *mlux* and one of its photoreceptors may receive as few as ten photons per second [3]. Regardless, they have been shown to reliably navigate through dense forests using visual landmarks for navigation. They visually stabilize themselves, avoid obstacles in dynamic and cluttered environments, and are able to identify and land on their nest (fallen or hollowed out logs) in dense forests.

Other nocturnal and diurnal insects like the house fly *Musca domestica* and hawk moth *Deilephila elpenor* have demonstrated similar abilities to operate in photon-limited environments. The *Musca domestica* can respond to optical flow stimuli at light levels of two to three photons per photoreceptor per second while the *Deilephila elpenor* can distinguish between colors at similar light levels [4].

Inevitably, low-light vision is limited by photon shot noise which has $SNR = \sqrt{N}$, where N is the sample size (photon detection events) of a given photodetector in a given exposure time. The

only way to increase the SNR is to increase the sample size of a photodetector. In spite of this challenge, neural recordings of photoreceptors in nocturnal insects show a strong to response to individual photons [5] with very good noise characteristics. Research by Dacke, Warrant, Baird, and others has focused on understanding the visual adaptations and image processing techniques utilized by both diurnal and nocturnal insects to achieve such high performance [6–14]. Their findings suggest that the higher optical sensitivity achieved by nocturnal insects is not a result of intrinsic photoreceptor adaptations but rather how the information is processed [5]. These processing techniques include the use of neural filtering [5] as well as spatial and temporal pooling [15–21].

Spatial pooling, also called spatial summing, involves the accumulation of photon detection events from a region of photoreceptors. This technique creates virtual photoreceptors with higher overall photon rates and sample size than individual photoreceptors but comes at the cost of spatial resolution. Temporal summing refers to an integration of photon detection events over some period of time. The longer the integration period, the larger the sample size of a given photoreceptor which helps overcome photon shot noise. While the existence of temporal summation is well known in individual photoreceptors, it comes at the price of temporal resolution–naturally a trade-off exists between spatial summing and allowable flight speeds before image blurring occurs. Simultaneous spatial and temporal summing has a multiplicative effect on the effective photon rate and sample size of the resulting "virtual photoreceptors". The *Deilephila elpenor* illustrates this principle beautifully to extend its visual sensing capabilities by two orders of magnitude [13]. If implemented effectively, these biological image processing techniques offer great promise for addressing the needs of MAV vision systems in photon limited environments.

Beyond fundamental detection and low-level image processing comes the question of how the information is used and perceived. In comparison to SLAM and other high-fidelity computer vision systems, insects demonstrate a "minimalist's approach" to autonomous navigation using computationally simple optic flow algorithms, in conjunction with low-resolution imagers. Optic flow is a well-known imaging concept that is used to describe visual stimuli based on their apparent motion relative to an observer or imager. Both nocturnal and diurnal insects use this concept to control groundspeed, stabilize themselves, perform obstacle avoidance, navigate to and from nests, and forage for food [3, 6, 7]. This allows insects to simplify a visual scene to its bare essentials and avoid complex vision problems such as object recognition and high-level perception. For next-generation MAVs that are striving for payloads <1g in weight, wingspans of several centimeters, and operating capabilities in GPS-denied environments and low-light environments, these biological systems offer a proof of concept and design inspiration.

1.2 Specific Aims

The goal of this work is to develop image sensors that enhance the ability of MAVs to perform visually-based tasks in low-light, cluttered, near-Earth environments without relying upon GPS. These tasks include obstacle avoidance, hover-in-place, and self-stabilization.

We develop two types of vision systems. The first operates primarily in the visible spectrum and uses custom, silicon-based, single photon avalanche diodes (SPADs) that we fabricate in standard CMOS processes. The second operates in the short wave infrared (SWIR) spectrum, primarily between 1.2μm and 1.8μm, using InGaAs photodiode arrays. This system enables vision at night using ambient airglow illumination or active IR illumination. This system is fabricated by hybridizing an InGaAs photodiode array with a custom read out integrated circuit

(ROIC) utilizing novel circuits to reduce dark current.

We also develop a novel framework to benchmark the performance of optic flow algorithms in photon-limited environments. In this work we compare "conventional" optic flow methods with biologically-inspired methods and explore the effects of spatio-temporal pooling and wide-field integration of optic flow. We present initial benchmarking results that identify what photon rates are necessary to obtain "useful" optic flow measurements and enable comparisons between different optic flow algorithms.

1.3 Thesis Organization

Chapter 2 discusses solid-state photodetectors with an emphasis on single-photon detection. We make a case for the use of single-photon avalanche diodes (SPADs) in ultra-low light vision systems. This chapter covers the design of single-photon avalanche diodes in standard CMOS processes and the integration of SPADs into image sensors. We demonstrate low-light optic flow estimation with a SPAD image sensor and characterize sensor performance. Technical contributions in this chapter include:

- Characterization of perimeter-gated SPADs including dark count rate, photon detection efficiency, and signal-to-noise ratio.
- Design and fabrication of SPAD image sensors and SPAD arrays including both frame-based and event-based architectures.
- Experimental characterization of SPAD arrays including pixel mismatch and off-chip optic flow estimation under photon-limited conditions.

Chapter 3 discusses the challenges of room-temperature short-wave infrared sensing (SWIR).

This chapter features experimental characterization of an adaptive analog circuit that improves room-temperature sensing capabilities for low-power, SWaP constrained applications. Technical contributions in this chapter include:

- Design and fabrication of both adaptive and non-adaptive CTIA structures.
- Experimental characterization of both adaptive and non-adaptive CTIA structures including limitations of dark current suppression and the effects of dark current suppression on detection limits and signal-to-noise ratio. Characterization is performed with both Si and InGaAs photodiode arrays.

This chapter contains material included in a manuscript that will be submitted for publication in the IEEE Sensors Journal and features excerpts from an article that demonstrates dark current suppression through adaptive CTIA structures. The article, which has been accepted for publication in the Proceedings of IEEE International Symposium on Circuits and Systems (2017), was co-authored by Alexander Castro, Mohammad Islam, Dr. Fow-Sen Choa, Dr. Geoffrey Barrows, and Dr. Pamela Abshire.

Chapter 4 provides a brief overview of optic flow estimation and optic flow algorithms. This chapter presents a novel simulation framework that enables benchmarking of various optic flow algorithms at low-light levels. We discuss the effects of spatial and temporal pooling and explore the limitations of photon-limited optic flow processing. Technical contributions in this chapter include:

- Development of a simulation-based framework that enables benchmarking of optic flow algorithms under varying illumination intensities.
- Analysis of optic flow performance for both biologically-inspired and traditional algorithms

using both natural and synthetic images. This includes quantifying the effects of dark count rate, optic flow velocity, and image characteristics on overall performance.

This chapter contains material included in a manuscript that will be submitted for publication in the Biological Cybernetics Journal and features excerpts from an article that benchmarks several optic flow algorithms at low-light levels using synthetic 2-D sine grating visual stimuli. The article, which has been published in the Proceedings of IEEE International Symposium on Circuits and Systems (2016), was co-authored by Dr. Geoffrey Barrows and Dr. Pamela Abshire.

Chapter 2

Low-Light Image Sensors

2.1 Fundamentals of Imaging

The field of low-light image sensors has long been dominated by photo-multiplier tubes (PMTs), charge-coupled devices (CCDs), and electron multiplying CCDs. These technologies are fairly mature and offer low-noise performance and high sensitivity. However, these sensors are not well-suited for integration with CMOS technologies and offer limited promise for next-generation low-light imaging systems. Increasingly, research efforts have focused on CMOS active-pixel sensors (APSs) and hybrid detectors for such applications. The ultimate goal of such research is to develop sensor arrays with single photon resolution, high frame rates, increased system integration, and decreased power consumption. In this chapter we discuss some of these technologies and suggest the use of single-photon avalanche diodes (SPADs) for low-light imaging applications.

2.1.1 Photodetectors

Detectors in standard CMOS processes include phototransistors, transmission gates, and photodiodes. While variations in device implementation exist, the physical mechanism of detection in each device is unique. While there is a wealth of non-CMOS devices (not fabricated in standard CMOS processes) including quantum dots, thin films (graphene, black phosphorus, etc.), and superconducting devices, in this chapter we focus on CMOS-compatible detectors and PMTs (as they have long been the gold-standard for detection of weak optical signals). Our goal is to develop low-light detectors that can be integrated into CMOS circuits for high density arrays and room-temperature operation.

Photodiode

One of the simplest and most common photo-transduction devices is the photodiode. In its simplest form, it consists of a P-N or P-I-N junction. Photon absorption in the diode generates an electron-hole pair (EHP) which can either recombine or separate into mobile charge carriers. EHPs generated in or near the depletion region have a high probability of separation due to the built-in field of the junction. While photocurrent generation in a diode can occur at both positive and negative bias voltages, typical applications reverse bias photodiodes. This is done to maximize the depletion region width and field strength. Additionally, saturation current in a diode is fairly insensitive to fluctuations in bias voltage while forward bias currents are exponentially dependent on fluctuations in bias voltage (see Eq. 2.1).

$$I_{diode} = I_{sat}\big(e^{qV/kT} - 1\big) - I_{op} \tag{2.1}$$

Linear-Mode Avalanche Photodiode

Linear-mode avalanche photodiodes (APDs) are essentially p-n junction photodiodes biased just below their breakdown voltage. The strong reverse bias produces a large internal electric field near the diode junction. Optically and thermally generated carriers near the depletion region are rapidly accelerated by the large field and are likely to undergo impact ionization. As a result, APDs have an internal multiplication gain (M), shown in Eq. 2.2, where $I_{ph,M}$ is the multiplied photocurrent, $I_{d,M}$ is the multiplied dark current, I_{ph} is the unmultiplied photocurrent, and I_d is the unmultiplied dark current. The regime in which APDs are biased is characterized by a high probability of impact ionization but no avalanche breakdown— impact ionization is not self-sustaining.

$$M = \frac{I_{ph,M} - I_{d,M}}{I_{ph} - I_d} \tag{2.2}$$

The random nature of impact ionization acts as a noise source during photodetection. A measure of this noise is the excess noise factor (F). It is related to the ratio of the standard deviation of the multiplication σ_M and multiplication factor (M) as shown in Eq. 2.3. The origin of excess noise is generally explained by the McIntyre formula [22,23] which relates multiplication noise to the relative impact ionization rates of electrons and holes.

$$F = 1 + \frac{\sigma_m^2}{M^2} \tag{2.3}$$

APDs operate as analog devices, similarly to photodiodes. Device gain is heavily dependent on bias voltage which enables APDs to be operated as gain modulated photodetectors [24]. Internal gain is generally on the order of $10^0 - 10^3$ which enables APDs to be used in low-light applications.

Geiger-Mode Avalanche Photodiode

Geiger-mode APDs, or single photon avalanche diodes, are essentially APDs biased above their breakdown voltage. In this operating regime, depletion fields are above the critical electric field ($3 \cdot 10^5 V/cm$ in Si) and impact ionization is a self-sustaining process.

SPADs operate like Geiger counters for photons and are intrinsically event-based detectors. A single optically or thermally generated electron hole pair can trigger avalanche breakdown in the device due to its high gain ($>10^6$). Avalanche current densities can reach upwards of 1 A/mm^2 and will destroy the device if not quenched. Therefore, SPADs must be quenched immediately after an avalanche process begins, and reset back to their "ON" state to enable subsequent detection. This "digital" or spiking behavior is contrary to typical photodetectors which are analog in nature.

Avalanche events triggered by thermally generated carriers are referred to as dark counts and are typically associated with trap and defect states. Reductions in DCRs have been achieved through optimization of SPAD geometries. Such optimization takes advantage of gettering inherent to CMOS fabrication processes and has been shown to effectively reduce thermal noise to around $<$1Hz/μm^2 even for large area SPADs [25]. In planar fabrication processes, electric fields tend to be enhanced around the diode periphery which causes the diode edges to be biased higher than the bulk region and results in premature edge breakdown. This results in non-uniform breakdown within the photosensitive diode area and degraded SPAD performance. Perimeter gating utilizes a poly-silicon field gate to repel carriers away from the diode edges and has been shown to effectively suppress premature edge breakdown [25]. In newer fabrication processes, retrograde wells can be used in place of poly gates to achieve the same effect.

Phototransistor

Phototransistors operate much like bipolar junction transistors (BJTs). When operating as a phototransistor, the BJT base is typically left floating and the base current is generated by a photocurrent. The "output" of the device is the collector-emitter current and, much like a BJT, is equal to $\beta \cdot I_b$. Thus, the phototransistor intrinsically has a current gain of β. Phototransistors typically have a slower temporal response than photodiodes. Recent developments in active pixel sensors have generally moved away from phototransistors in favor of p-n junction diodes.

Transmission Gate and Charge-Coupled Devices

Charge-coupled device (CCD) technologies are based on the photogate. Unlike photodiodes, they accumulate charge in metal-oxide-semiconductor (MOS) capacitors. Applying a voltage to the polysilicon gate of a MOS capacitor induces a local depletion region in the silicon substrate underneath the gate. This potential well defines an area in which photogenerated carriers can be accumulated. In CCD arrays, each photogate serves as a single pixel.

Readout requires that charges be transferred from one photogate to the next across a row/column, much like an analog shift register. Charge transfer between photogates must be highly efficient to ensure signal integrity and requires custom CCD processes to enable large CCD arrays. As a result, CCD technologies don't offer the integration of pixel-level or array-level transistors. While CCDs are common commercial cameras, they require multi-chip solutions for many vision applications.

In 2001, electron multiplying CCD (EMCCD) arrays were first introduced [26, 27]. They improve readout SNR at high frame rates with an internal gain mechanism. The gain is provided by impact ionization in a gain register which is part of the read-out mechanism. While readout

SNR is improved, the dark current is also amplified which ultimately limits device performance. They are widely used for low light applications [28] (10^{-4} lux) and offer higher data rates than CCD arrays.

While CCDs demonstrate low noise performance[1] (<10 electrons/μm^2/sec) and high fill factors (>50%), they do not enable focal plane array processing and generally have high power consumption (>100 mW). As a result, they are severely limited by the readout bottleneck and low-power CMOS techniques cannot be adopted to reduce power consumption. Thus for low power compact vision systems, such as MAVs, CCD technologies are not ideal.

Photomultiplier Tube

PMTs are high-voltage, high-gain vacuum devices for photo-detection of weak optical signals. Photon absorption occurs in a photocathode. Photons with energy greater than the work function of the photocathode produce photoelectric emission of electrons from the photocathode into the vacuum. This signal is then amplified through a series of dynodes (electron multipliers) which use large electric fields to rapidly accelerate the particles and achieve signal gain through secondary electron emission. After several multiplication stages the electrons are collected at the PMT anode.

Electron focusing and multiplication demand large fields and often high voltages, typically around 1000V. PMTs are not compatible with standard CMOS technology and require large areas relative to their CMOS and CCD counterparts. This severely limits their applications and makes PMTs unsuitable for low-power, compact, and versatile low-light imagers.

Various PMTs are commercially available through companies such as Hamatsu [29] and

[1]Data compiled from DigiKey CCD Image Sensors (www.digikey.com) and Hamamatsu CCD Cameras (http://www.hamamatsu.com)

Photonis [30]. The current market offers micro-PMTs with reduced package sizes ($<5mm^2$), high voltage modules($>1000V$), and low voltage modules ($<10V$). Such commercially available PMTs offer high gains ($>10^6$), dark currents in the nA range, and weigh upwards of 10g. Even with the miniaturization of PMTs, they are still far too large to be implemented in imaging arrays used for many computer vision applications and cannot offer satisfactory spatial resolution for most imaging applications [31].

2.1.2 Photon Shot Noise

Photon shot noise arises from uncertainty in photon arrival time. It is widely accepted that photon arrival can accurately be modeled as a Poisson process. The physical limitation of low-light imaging is shot noise. While all other noise sources (1/f noise, thermal noise, reset noise, dark noise) can be minimized through circuit design and system-level approaches, shot noise is unavoidable.

Poisson processes have a probability density function (PDF) given by Eq. 2.4 where k and λ are the measured number and expected number of events in a given time window. Poisson processes are described by their Poisson variable λ and have mean and variance equal to λ.

$$P(k) = e^{-\lambda}\frac{\lambda^k}{k!} \tag{2.4}$$

2.2 SPAD Image Sensors

SPADs are widely used in time-correlated imaging. These applications take advantage of the high temporal resolution and digital behavior of the device. It is noteworthy that there is nearly no work discussing the use of SPADs in low-light image sensors that do not focus on

time-correlated applications. In this section we argue that SPADs are useful for general purpose low-light imaging and describe the unique characteristics of SPADs that make them appealing for such applications.

2.3 Why SPADs?

Unlike CCD and PMT technologies, high performance SPADs can be implemented in standard CMOS processes. This allows for the integration of photodetection and image processing on a single chip. The benefits of this can be enormous for SWaP-constrained applications and is described in more detail in later sections. Unlike quantum devices, SPADs can operate at room temperature and do not require ultra-low or cryogenic temperatures to operate and achieve single photon resolution.

Many of the noise sources that plague APS are overcome in SPAD pixels due to their high gain and digital detection capabilities. This allows SPAD pixels to be much simpler, yet often larger, than their CMOS APS counterparts. However, this high gain introduces new noise sources and thus the selection of one technology over another requires consideration of the trade-offs between power, speed, and noise.

Fig. 2.1 illustrates the operating range for three different pixel architectures: SPAD-based pixels, linear active pixel sensors, and logarithmic active pixel sensors—this illustration offers a first-order approximation for operating regimes of each architecture.

The noise floor for SPAD sensors is set by dark count rate which can be as low as 1-100Hz depending on technology, pixel size, and device architecture. Typical SPAD operating speeds range from 1MHz to 100MHz with the fastest reported SPAD pixel operating at 185MHz [32]. Tradeoffs between speed and noise in SPADs are discussed later in this chapter. As is illustrated

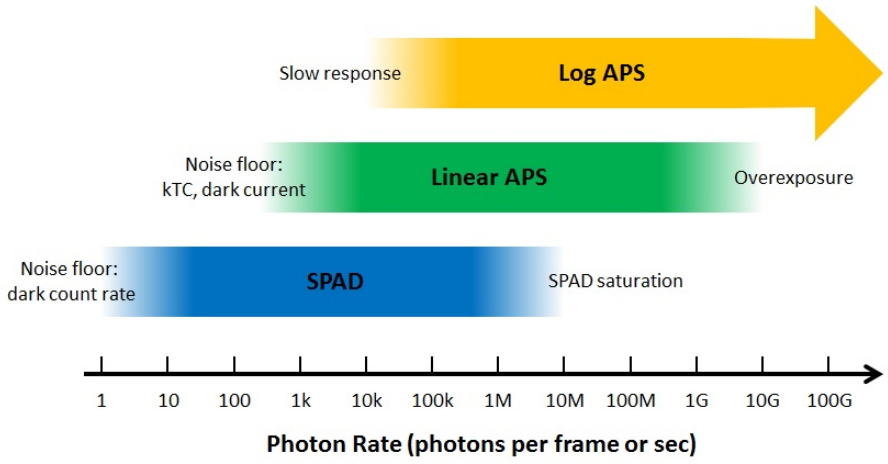

Figure 2.1: Illustration of operating ranges for SPAD-based pixels, linear active pixel sensors, and logarithmic active pixel sensors. The figure illustrates what is typical for each technology rather than peak performance for state-of-the-art sensors.

in Fig. 2.1, the high gain of SPADs makes them more resilient to kTC noise and thermal noise than their APS counterparts due to their intrinsically higher signal levels and digital processing which provides immunity to low level sources of analog noise.

Linear APS are susceptible to dark current as well as kTC noise. While precise values for these noise parameters vary greatly, kTC noise for a 100fF capacitor is >100 electrons/sec, and dark currents are similar to SPAD sensors and generally range from 10-1000 electrons/sec. The upper limit for linear APS is set by the minimum allowable exposure time for pixels. At photon rates in excess of 1G photons per exposure, discharge rates across a 100fF capacitor would require frame rates in excess of 1k frames per second. It is important to note that sensor architectures that employ correlated double sampling and gating can effectively reduce kTC noise and dark current by an order of magnitude or more. Similarly, sensor architectures that implement auto-exposure, over-sampling, and other high-dynamic range techniques can effectively prevent overexposure of pixels at much higher illumination intensities. For our first-order approximation

we ignore these approaches.

Logarithmic APS are typically limited by slow temporal response at low photon rates and non-uniformity in photo-response at extremely high photon rates due to operation near the sub-threshold limit. While the upper limit is unavoidable and rarely reached with traditional image sensors, the lower limit is proportional to to the ratio of capacitance at the sensing node and photocurrent. The well-known adaptive logarithmic photoreceptor developed by Delbruck [33] improves temporal response at low-light levels through the use of a high-gain feedback loop. Even with such approaches and optimal sensor designs that minimize capacitance at the sensing node, logarithmic APS are unable to image with time constants comparable to typical frame rates below photocurrents of around 1fA, or 10k photons/sec.

It is important to mention recent work by Fossum *et al.* [34–36] who demonstrated ultra low-noise CMOS APS referred to as quanta image sensors. These sensors achieve extremely high transduction gains (on the order of mV/e^-) without the use of avalanche breakdown. Advantages of SPAD-based imagers over such approaches are that SPAD sensors do not rely upon oversampling and high frame rates to acquire images in near-dark environments. SPAD image sensors can integrate local spike-based processing instead of relying on off-chip image processing. This makes SPAD well-suited for computational imaging applications and image sensors with embedded focal plane image processing.

Ultimately, SPADs are an attractive imaging technology because:

1. They demonstrate comparable noise performance to ultra low-noise CMOS APSs with higher timing resolution.

2. They can be integrated into standard CMOS fabrication processes.

3. Existing research on SPAD imagers has focused on time-of-flight imaging and primitive

frame-based sensors. To date, no SPAD imagers have been demonstrated with embedded image processing or neuromorphic design inspiration. Thus, this promising area of work remains untapped and offers a novel path towards improved low-light imaging capabilities.

While SPADs are not appropriate for daytime imaging or extremely high dynamic range imaging (>120dB), they are uniquely well-suited for high-sensitivity, low-light imaging applications. In this chapter we demonstrate these types of SPAD-based image sensors.

2.3.1 SPADs in CMOS

Developing SPADs in CMOS fabrication processes requires intimate knowledge and understanding of the process layers and characteristics. One of the biggest challenges with designing SPADs in CMOS processes is preventing premature edge breakdown (PEB). Ideally, a SPAD should breakdown uniformly across the entire active region which requires a uniform field. However, since CMOS processes are planar and devices have finite size, electric fields tend to be enhanced around the diode periphery. This results in the edges effectively being biased higher than the bulk region and breaking down first. Since the device periphery is optically small and has negligible contribution to detection efficiency, edge enhancement of electric field will substantially increase dark count rates and decrease detection efficiency.

Several techniques exist to prevent premature edge breakdown. All of them are designed to either reduce the electric field or to increase the breakdown voltage at the device periphery. In doing so, avalanche events have an increased probability of occurring in the center/bulk region of the device. Some of these techniques employ multiple high-voltage wells, buried diffusion layers [37], diffusion-based guard rings [38], retrograde wells [39], and polysilicon-based guard rings [25]. Our SPADs are fabricated in older CMOS processes (ON 0.5μm and XFab XC06) which do not offer many of these process layers. Therefore, we are limited to using a poly-silicon

field gate to repel carriers away from the diode edges.

Another problem that plagues CMOS SPADs is high DCRs. Thermally generated EHPs are often associated with low energy trap and defect states. Techniques for controlling DCRs are limited since: 1) it is difficult to control defect/trap densities, and 2) defect densities are controlled by the foundry and technology node. Both temperature and bias voltage have a very strong effect on DCRs. Thermal carrier generation has an exponential relationship with temperature. Cooling SPADs by 30-40^{o}C is possible with Peltier coolers and can reduce DCRs by 1-2 orders of magnitude. Operating SPADs at a low excess bias (voltage over breakdown) also reduces DCRs but comes at the expense of detection efficiency and other key SPAD performance metrics. Work by Nouri and Abshire [40] have shown that narrow or "cigar-like" device geometries can be utilized to take advantage of gettering processes inherent to all CMOS fabrication processes. This strategy helps mitigate the dark current rates that scale non-linearly with SPAD area and plague large area SPADs (>1000μm^2).

Much of the existing research on SPAD-based sensors has focused on time-correlated imaging. The goal of these efforts is the development of imagers with ultra-high temporal resolution for 3D imaging applications [41], positron emission tomography (PET) applications [42], fluorescence lifetime imaging [43], and other time-correlated imaging applications [44]. These applications have largely driven improvements in SPAD technology over the past two decades. Most recently, SPADs have been implemented in deep submicron technology nodes including 90nm [45], 130nm [46], and 180nm [47]. At the device level, SPAD research has worked to suppress dark counts and premature edge breakdown in advanced technology nodes through the use of retrograde well, perimeter gating, and enhancement implant layers [48–54].

2.3.2 SPAD Noise

Cross-Talk and After-pulsing

Optical and electrical crosstalk are a result of one SPAD event inducing avalanche breakdown in a neighboring SPAD. Electrical crosstalk is caused by electrical interference through the substrate or power supply lines. Optical crosstalk is due to photon emission of a SPAD during avalanche breakdown, triggering breakdown in another SPAD. It is heavily dependent on the number of carriers involved in an avalanche event and on pixel pitch. Recent work has explored and quantified the effects of SPAD density and substrate thickness on optical crosstalk [55–57].

After-pulsing is similar to crosstalk except that instead of a SPAD generating a "false" event in a neighboring pixel, it generates a "false" event in itself—crosstalk and after-pulsing are the only temporally-correlated noise sources in SPADs. This false triggering is a result of carriers that are trapped during avalanche breakdown and released after the avalanche event is quenched. It has been demonstrated that after-pulsing rates can be substantially suppressed in actively-quenched SPADs through programmed hold-off times following an avalanche event. While effective, this technique reduces the maximum operating speed and dynamic range of a SPAD. We also see a strong dependence between after-pulsing probability and SPAD excess bias voltage. These dependencies introduce significant trade-offs between speed, detection efficiency, and noise in SPADs. Moreover, these trade-offs vary drastically across device architectures and fabrication technologies. Recent work by Bronzi *et al.*, Pancheri *et al.*, and others has discussed the effects of excess bias and dead-time on after-pulsing probability [58–61].

SPAD Gating

A technique often utilized with noisy SPADs, particularly IR SPADs, is SPAD gating [62,63]. As opposed to free-running mode, in gated mode SPADs are photosensitive only during a defined, periodic gate-on window—during this time they are reverse biased above breakdown. During the gate-off window, the SPAD reverse bias is lowered below the breakdown voltage such that no avalanche events can occurs and the SPADs are not photosensitive. This operation prevents false detection of photons outside of some defined timing window. Sufficiently small timing windows can substantially reduce the effects of all temporally-correlated noise sources as well as thermal noise and other uncorrelated noise sources. This technique is similar to the use of optical choppers and lock-in amplifiers in highly sensitive optics experiments.

Timing Jitter

Much like after-pulsing and cross-talk, timing jitter is a noise source specific to SPADs [64, 65]. It refers to an uncertainty in the temporal dynamics of a given avalanche event. Each avalanche process begins near the location of photon absorption. The avalanche process builds up locally and spreads through multiplication-assisted diffusion and photon-assisted propagation (photon emission and re-absorption). These processes are stochastic in nature and introduce timing jitter. Avalanche build up is also strongly dependent on electric field intensity with varies throughout a device and introduces additional jitter to the temporal dynamics of avalanche events.

Dark Count Rate

Dark count rate (DCR) refers to the generation of photo-generated carriers when a SPAD is in a dark environment. In general, two process dominate the generation of dark counts: 1) tunneling, and 2) thermal carrier generation. These processes cannot be differentiated from optical carrier

generation and DCRs often set the noise floor for low-light sensing.

Tunneling refers to the penetration of electrons through a potential barrier. This process is highly dependent on electric field strength and is often observed in SPADs due to their large reverse bias voltage. Tunneling typically occurs between the valence and conduction band of the SPAD anode and cathode and is more prevalent in deep sub-micron processes that have higher doping concentrations and narrower depletion regions. Tunneling can be assisted by trap states and is often the dominant source of dark counts at large excess bias voltages. Tunneling is not a prevalent noise source in standard p-n junction diodes and is unique to their high-voltage counterparts

Thermal carrier generation, also known as Shockley-Read-Hall generation, refers to the process by which electrons or holes may be promoted from the valence band to the conduction (or vice versa) through thermal processes. Direct promotion is unlikely in Si SPADs because the bandgap of Si (1.1eV) is substantially larger than thermal energy at room temperature (kT = 25.9meV), and because Si is an indirect bandgap semiconductor. Thus thermal generation is often assisted by trap and defect states. These mid-band energy states provide intermediaries by which carriers can increase their energy. Shockley-Read-Hall generation is highly sensitive to temperature and is the dominant source of dark counts at high temperatures and small excess bias voltages. Thermal generation is ubiquitous in CMOS photodetectors and similar silicon processes can reasonably expect comparable rates of thermal generation.

State-of-the-art low noise performance for SPADs is DCR<1Hz / μm^2 [66]. A summary of SPAD performance across several technology nodes is provided in Table 2.1.

2.3.3 Photon Detection Efficiency

Photon detection efficiency (PDE), or photon detection probability (PDP), is a measure of how many photons are detected in a time window at a known photon rate. The ratio of photons detected to photons incident, or expected detection events, within a given time window provides PDE. PDE increases with bias voltage. However, at high voltages PDE increases slower than DCR. Common techniques for boosting PDE include the use of anti-reflection (AR) coatings and micro-lenses. These have the effect of increasing the number of photons collected in the Si detector.

State-of-the-art CMOS and CCD image sensors exhibit peak detection efficiencies >70% [67, 68] with the use of AR coatings and micro-lenses. More primitive CCD sensors (front-side illuminated and poly-Si gates) and CMOS sensors without micro-lenses have PDEs closer to 30-50% [67, 69] depending on technology. Typical PDEs for SPADs range from 1% to around 30%. Table 2.1 provides a comparison of state-of-the-art SPAD devices across several technology nodes. The highest reported PDE for CMOS SPADs is 72% [70] in a 130nm node. It is important to note that this is achieved at a 12V excess bias where the normalized DCR is in excess of $20\text{Hz}/\mu m^2$. The importance of fabrication processes with image sensor friendly processing is demonstrated by Niclass *et al.* [71]. SPADs in the same 0.35μm process were fabricated with and without a nitride passivation layer. The nitride layer reduced peak PDE by a factor of 7 for one value of excess bias while also shifting the location of the maximum PDP by >100nm. At wavelengths near 450nm, a decrease in PDP by a factor of approximately 10 is observed.

Table 2.1: State-of-the-Art SPAD Comparison

	Niclass [72]	Niclass [71]	Mandai [73]	Gersbach [74]	Gersbach [74]	Richardson [75]	Webster [70]
Tech. Node	0.8μm	0.35μm	180nm	130nm	130nm	130nm	130nm
DCR (Hz) @ RT	350 $V_{ex} = 5$	750 $V_{ex} = 3.3$	<5k $V_{ex} = 4$	11k $V_{ex} = 5$	220 $V_{ex} = 2$	25 $V_{ex} = 1.4$	>1k $V_{ex} = 12$
Active Area (μm^2)	38	78	78	58	58	50	50
Norm. DCR (Hz/μm^2)	9.2	9.6	<64	190	3.8	0.5	>20
Peak PDP (%) @ λ (nm)	12 635 $V_{ex} = 5$	36 465 $V_{ex} = 3.3$	36 600 $V_{ex} = 4$	36 480 $V_{ex} = 5$	26 480 $V_{ex} = 2$	28 500 $V_{ex} = 1.4$	72 560 $V_{ex} = 12$

2.3.4 SPAD Front-End

A key function of SPAD front-end circuitry is to detect and quench avalanche events. If an avalanche event is not quenched, high current densities during avalanche breakdown can result in thermal destruction of the device. Following quenching, front-end circuitry must reset the SPAD to its "active" state where it is ready to detect additional avalanche events. The time between quenching and reset is referred to as dead-time because SPADs are in an off state and are not sensitive to thermally- or optically-generated carriers. SPAD front-ends are generally classified into one of three categories: passive quench circuit (PQC), active quench circuit (AQC), or hybrid quench circuit. The primary differences between each of these circuits is the mechanism that controls the timing of the quench phase and reset phase.

Passive Quenching Circuits

Passive quenching circuits (PQCs) consist of a ballast resistor connected in series with a SPAD. When avalanche breakdown occurs, a large current passes through the ballast resistor, increasing the voltage drop across the resistor, and decreasing the voltage across the SPAD. When the SPAD bias decreases below the breakdown voltage, the avalanche process is stopped. SPADs self-quench in a PQC and eliminate the need for additional sensing circuitry. Once the avalanche event is quenched, the SPAD current will drop and the SPAD bias voltage will increase back to its "active" state. PQC benefit from a fast quench time but suffer from long recharge times that are set by the RC time constant of the circuit. The resistance in this RC constant is equal to the value of the ballast resistor and the capacitance is equal to the parasitic capacitance of the diode. As a result, PQC typically operate at speeds 10-100 time slower than AQC. PQC do not employ timing circuits that enable a programmable "dead-time". As a result they often suffer from higher after-pulsing probabilities than AQCs, can be triggered while recharging, and are referred to as paralyzable detectors.

Active and Hybrid Quenching Circuits

Active quenching circuits (AQCs) and hybrid quenching circuits (HQCs) operate by detecting avalanche events and adjusting the SPAD's bias conditions with active circuit elements. Compared to PQCs, they are able to achieve higher operating speeds and more precise timing control at the expense of size and power consumption.

Active quenching lowers the SPAD bias voltage below breakdown once an avalanche event is detected. This ensures that quenching time is not effected by statistical variations in avalanche multiplication. Following quenching, the bias voltage is kept low for a pre-determined length of

time. This enables a well-defined dead-time or hold-off time to prevent after-pulsing. After a programmed dead-time, the AQC increases the bias voltage above breakdown. Reset times are much faster for AQCs than PQCs since they are driven by active components and can range from nanoseconds to tens of nanoseconds. Generally, AQCs are non-paralyzable since an active reset prevents the SPAD from firing until the reset phase is complete. AQCs can operate at rates >150MHz [32].

Hybrid circuits combine active and passive techniques. Typically, a large ballast resistor is used to implement passive quenching and active components are incorporated to add a hold-off time and fast reset [76]. In this configuration large ballast resistors can be used without compromising high after-pulsing probabilities.

Dynamic Range

SPAD dynamic range is heavily dependent upon front-end circuitry. As previously mentioned, PQCs are characterized as paralyzable detectors. At high photon rates a nonlinear decrease in count rate is observed because passively-quenched SPADs can be triggered while recharging. This has the effect of extending the dead time of the detector and prevents the SPAD output from crossing a spike threshold voltage. If overexposed, passively-quenched SPADs are "paralyzed", detect no photons, and effectively have in infinite dead-time. AQCs are generally non-paralyzable and use active circuitry to guarantee a well-defined dead-time regardless of illumination intensity. Photons absorbed within this dead-time are not detected and cannot trigger an avalanche event.

The relationship between input photon rate (λ_{in}) and measured photon rate or count rate (λ_{out}) for both paralyzable and non-paralyzable detector models are given by Eq. 2.5 and Eq. 2.6, respectively. The upper-limit of count rate for paralyzable detectors is given by $1/(e \cdot \tau)$ and $1/\tau$ for non-paralyzable detectors, where τ is the dead-time of the detector. Thus AQCs offer

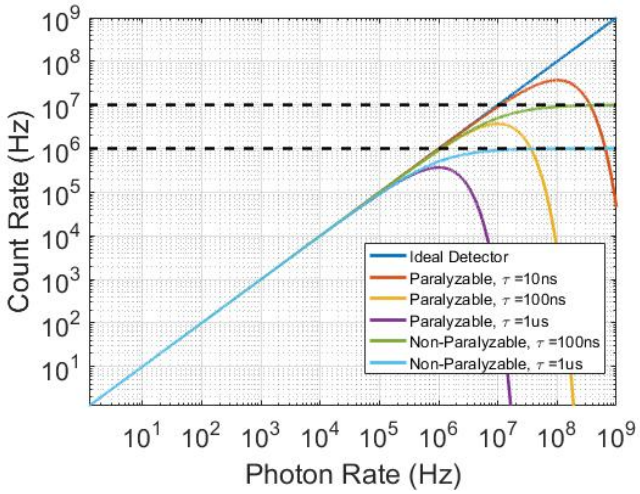

Figure 2.2: Count rate versus photon rate for ideal, paralyzable, and non-paralyzable photon counting modules. Dead time, τ, is varied over three orders of magnitude to illustrate its effect on maximum count rate.

higher dynamic range and detection rates than their passive counterparts. Ultimately, the dynamic range of each model is limited by the dark count rate of the SPAD and the dead-time of the detector. Assuming a modest DCR of 100Hz and a dead-time of 100ns, we can achieve a maximum dynamic range of <100dB. It is important to note that linearity in count rate decreases near peak count rates. This is illustrated in Fig. 2.2 for multiple values of dead-time.

$$\lambda_{out} = \lambda_{in} \cdot^{-\lambda_{in} \cdot \tau} \tag{2.5}$$

$$\lambda_{out} = \frac{\lambda_{in}}{1 + \lambda_{in} \cdot \tau} \tag{2.6}$$

2.3.5 SPAD Spike Processing

Time-correlated SPAD imagers rely upon the high temporal resolution of SPADs and are sensitive to precise spike timing. For intensity measurements, individual spikes carry virtually no information. Instead spikes must be integrated over some exposure time to provide meaningful intensity estimates.

Digital Counter Pixel

Digital counters are well suited for integrating SPAD pulses [77,78]. They offer low-noise performance, low leakage currents, and low pixel mismatch. However, high resolution digital counters (>8 bits) can occupy a large area and can be difficult to implement in compact pixel designs.

Digital counters require an input signal that clearly crosses logic levels. This requires that SPADs have an AQC with rail-to-rail output or PQC with a digitizing block (comparator, Schmitt trigger, etc.). Digital counters are not well-suited for local image processing tasks in the focal plane as digital values would likely have to be converted to an analog voltage and would require a digital-to-analog converter (DAC) in each pixel. This further limits pixel size and fill factor. Overall, digital counters offer an attractive solution for frame-based SPAD imagers with off-chip processing.

Analog Counter Pixel

Analog counters offer an attractive alternative to their digital counterparts. They are more compact and convert spikes directly to analog voltages, making them well-suited for embedded image processing tasks. Ultimately, counter resolution is limited by pixel and readout noise. However, even a modest step-size of 5mV corresponds to roughly 9-bit resolution with a 3.3V supply voltage. The biggest challenges with analog counters are: achieving high linearity, keeping

leakage currents small, and minimizing pixel mismatch.

Pixel-level analog counters for SPAD-based photon counting applications have recently been demonstrated by Pancheri *et al.* [79], Chitnis *et al.* [80] or Dierickx *et al.* [81]. High fill-factors of 20% have been demonstrated with photo-response non-uniformities of 1%. However, existing counters suffer from limited dynamic range, large area implementation, and/or non-linear outputs which require corrections for linearization.

2.4 SPADCAM

We implemented a 16x16 array of passively quenched SPADs with analog counters in each pixel. The chip was fabricated in a ON0.5μm process on a 1.5mmx1.5mm TinyChip through MOSIS. A screenshot of the chip layout is shown in Fig. 2.3.

2.4.1 Pixel Architecture

The pixel implements a perimeter-gated SPAD with passive quenching. We chose passive quenching because it enables smaller pixels and a simpler pixel design. This means we have paralyzable SPADs and the image sensor can be "photo-bleached" if sufficiently bright stimuli are provided. Passive quenching is notoriously plagued by slow and highly variable recovery times in excess of 1μsec. To mitigate these effects we implement a pulse shaping circuit following the SPAD. This consists of a comparator and monostable pulse generator, shown in Fig. 2.4, to improve pulse-width uniformity of SPAD pulses. The monostable pulse generator was designed to produce a pulse width of 4.3ns. We verified this through Monte Carlo simulations and showed a standard deviation in pulse width of 93ps, or approximately 2% of the mean pulse width. These digital pulses propagate to an analog counter that integrates photon counts. The analog counter value

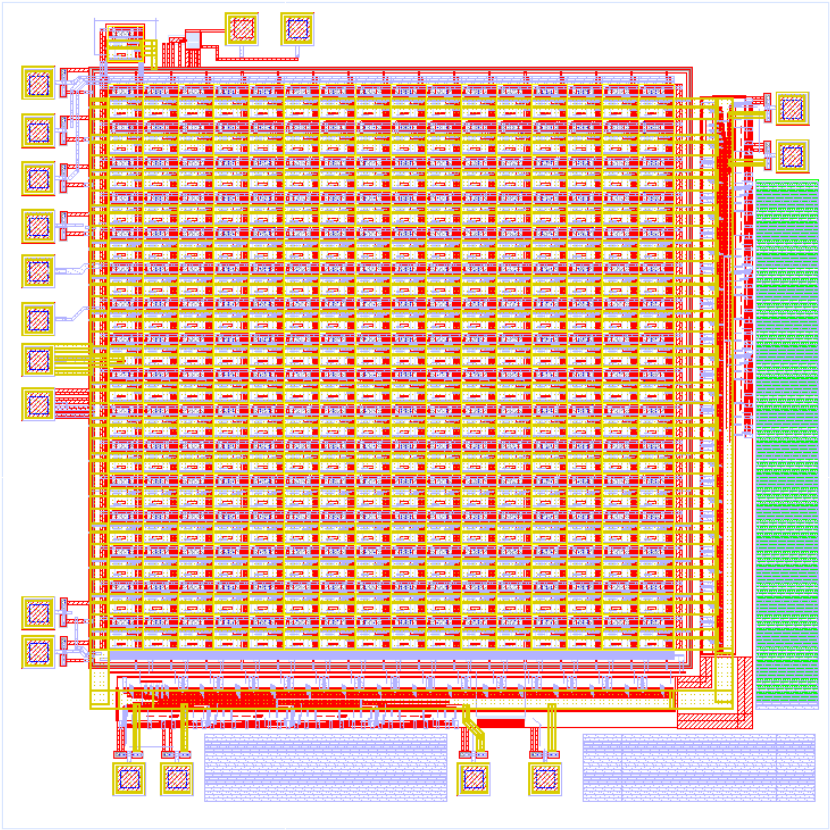

Figure 2.3: Layout picture of SPADCAM. Array is comprised of a 16x16 pixel array with frame-based, analog readout.

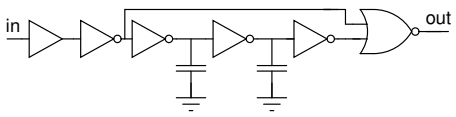

Figure 2.4: Schematic of pulse shaping circuit. The circuit consists of a comparator followed by a monostable pulse generator and is designed to generate a pulse width of 4.3ns.

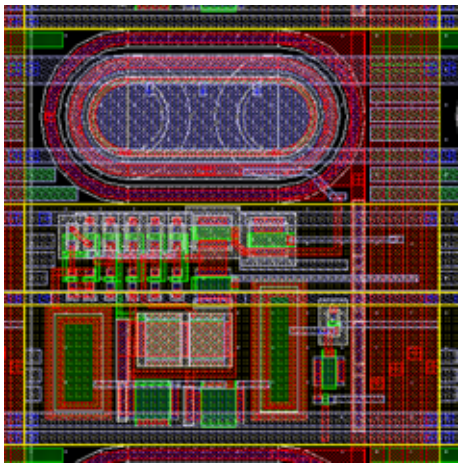

Figure 2.5: Layout picture of the SPAD pixel. It contains a 360μm^2 perimeter-gated SPAD in an area of 60μm x 60μm.

is readout through a simple source follower circuit.

The pixel is implemented in an area of 60μm x 60μm with a fill factor of 10%. The SPAD is cigar-shaped and perimeter-gated with an active area of 360μm^2. A screenshot of the pixel layout is provided in Fig. 2.5.

2.4.2 Counter Architecture

The analog counter is adapted from a design by Gasparini *et al.* [82]. A schematic of the counter circuit is provided in Fig. 2.6. A digital spike is input to the circuit at C1 and passes through a high-pass filter (HPF) formed by C1 and M1. When the gate of M2 goes high, capacitor C2 partially discharges through M2 and the counter output voltage, stored on C2, decreases.

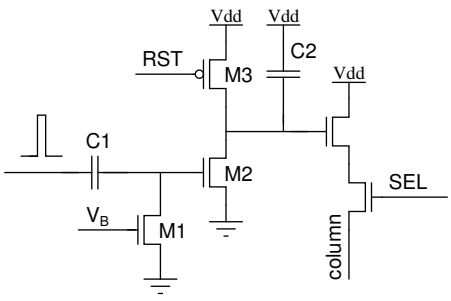

Figure 2.6: Circuit schematic of analog counter. Input spikes are capacitively coupled to the gate of M2 through C1. The output voltage is stored across C2 and readout through a source follower.

By removing charge packets in this manner, the circuit achieves high linearity over its entire dynamic range.

The step size of the counter is tuned by adjusting V_B, the gate voltage of M1. The limiting factor of the step size is the noise floor. One noise source is leakage at the output node. We experimentally measured the average leakage rate across the array as 7.2 mV/sec with a standard deviation of 0.69 mV/sec. This corresponds to a leakage current of about 1fA. For moderate exposure times of 10-100ms this corresponds to a discharge of <1 mV per frame. This leakage rate is well below the step size of the counter even when tuned to 9-bit resolution (6.4mV step size). By adjusting V_B, we demonstrate an average step size of approximately 10mV (see Fig. 2.7) and 30mV (see Fig. 2.8). In both cases, counters display high linearity with R^2>0.99 over the entire dynamic range of the counter.

2.4.3 System Architecture

This imager uses a standard frame-based readout with a simple source follower for each column. The output of the chip is an analog voltage and is read by an off-chip ADC. Maximum readout

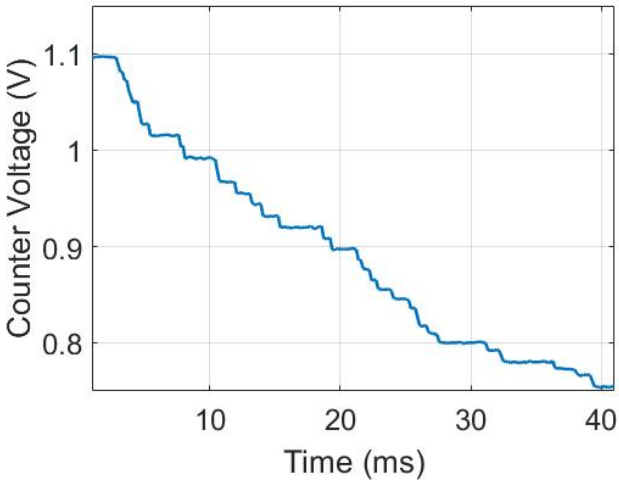

Figure 2.7: Transient response of analog counter in low-light conditions. Average step-size is approximately 10mV.

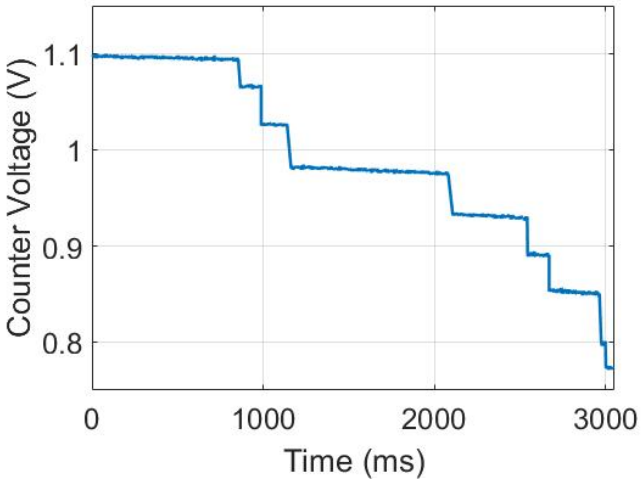

Figure 2.8: Transient response of analog counter in low-light conditions. Average step-size is approximately 30mV.

time per pixel is conservatively 50ns. The entire array can be readout in 12.8μs which has an equivalent frame rate of about 78kHz. The operation of the imager readout is as follows: First, a global reset signal is sent to all pixels to recharge counter outputs to VDD. Second, the chip integrates counts during an exposure time. Third, each element in the array is sequentially read.

2.4.4 Dark Count Rate

Dark count rate was measured by averaging the total number of SPAD pulses over at least 500 consecutive frames in dark conditions. Each frame had an exposure time of >50ms. During all measurements the anode and gate connection are shorted. DCR is measured for anode voltages between 13.40V to 13.80V. This represents a rail-to-rail voltage of about 16.7V to 17.1V and an excess bias of approximately 1V-1.4V. Fig. 2.9 shows DCR for each of the 16 pixels in a single row of the image sensor as a function of anode voltage. DCR is normalized to units of Hz/μm^2 for a 360μm^2 SPAD. We note saturation in the DCR curves which corresponds to saturation in the analog counter, not DCR. In Fig. 2.10 we show the cumulative distribution of dark count rates in the same row of the array. At an anode voltage of 13.68V, nearly all pixels are fully saturated by dark counts.

2.4.5 Photon Detection Efficiency

Optical characterization is performed with a monochromator (Newport Cornerstone 260 1/4M, model 74100) in an optically isolated Faraday cage. The monochromator illuminates the sensor with a monochromatic light beam generated by a 150W Hg-Xe light source. A diffraction grating provides a narrow-band stimulus with 2nm resolution and enables wavelength to be swept from 200nm to 1100nm.

Optical power is measured with a calibrated photometer (Newport 818-UV Low-Power UV

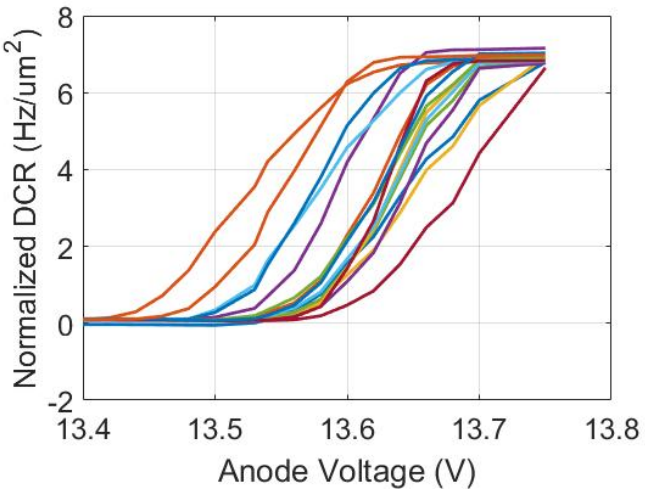

Figure 2.9: Normalized dark count rate versus anode voltage across a single row of SPADCAM.

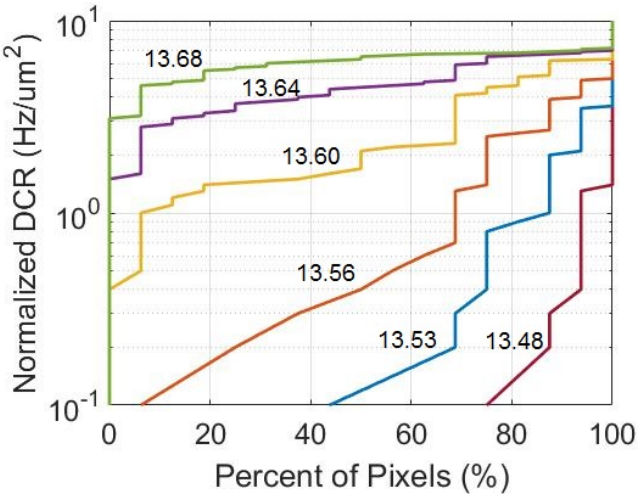

Figure 2.10: Cumulative distribution of dark count rate across a single row of SPADCAM. Each curve corresponds to a unique anode voltage ranging from 13.48V to 13.68V. Anode voltage is indicated next to each curve.

Enhanced Silicon Photodetector) and an optical power meter (Newport 1830-C). The detector has a sensor area of 100mm^2 and a spectral range of 200nm to 1100nm. A test beam with power density of 17.5nW (measured on the optical power meter) is configured at a wavelength of 450nm which corresponds to a mean photon rate of 396 photons/μm^2/sec.

Photon detection efficiency (PDE) is measured by averaging the total number of SPAD pulses over at least 500 consecutive frames, first in dark conditions, then in light conditions. The dark measurement provides DCR at a specific bias while the light measurement provides the total number of dark counts and the total number of photons detected per frame. We estimate the number of detected photons as the difference between the bright and dark measurement. The PDE is approximately equal to the number of photons detected per frame (N_d) divided by the number of incident photons expected per frame (N_{ph}).

Fig. 2.11 shows PDE for each of the 16 pixels in a single row of the image sensor as a function of anode voltage. PDE is computed for the area of the SPAD, not pixel, so fill-factor is not included in the measurement. PDE measurements are limited by saturation of the pixel-level analog counters. At high anode voltages, where peak PDE is expected, high DCRs quickly saturate the counter and prevent accurate measurements of PDE. For this reason, PDE are limited to anode voltages before saturation occurs. Thus, we cannot characterize PDE performance for the SPADs at larger excess bias voltages.

We note a low peak PDE of <1% at a wavelength of 450nm. This occurs at a modest excess bias voltage around 1V. This stands in contrast to typical SPAD PDEs ranging from 1% to 30%. One explanation for this result is the fabrication process. The SPADs described in this work have been fabricated in an ONC5 process which not offer optical windows removed from passivation layers, deposition of dielectric materials that enhance photon absorption, or micro-lenses that

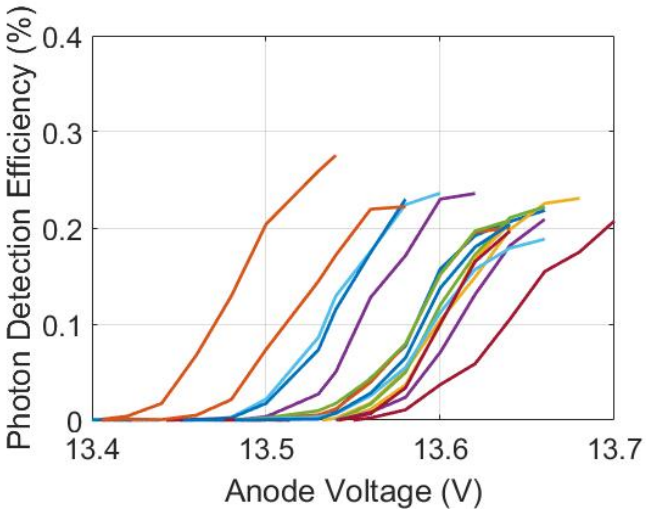

Figure 2.11: Photon detection efficiency versus anode voltage across a single row of SPADCAM.

provide a concentration factor and increase photon flux for each pixel. As previously mentioned, Niclass *et al.* [71] demonstrate that passivation layers can reduce PDE by approximately one order of magnitude. We expect that this is the dominant factor in the abnormally low peak PDE measurements. Increasing frame-rates during pixel-level characterization would have the effect of increasing the capacity of pixel-level counters and would enable PDE characterization at larger excess bias voltages which should also increase peak PDE.

2.4.6 Signal-to-Noise Ratio

Signal-to-noise ratio (SNR) is the ratio of photon detection rate over dark count rate. We calculate it as the product of quantum efficiency and photon flux over dark count rate. We estimate SNR at a photon flux of 1M photons/SPAD/sec across a range of bias voltages. Fig. 2.12 shows SNR estimates for select pixels across a single row of the SPAD image sensor. SNR is estimated using PDE and DCR measured in the previous sections.

A notable result is that peak SNR does not occur at the same bias voltage as peak PDE. This

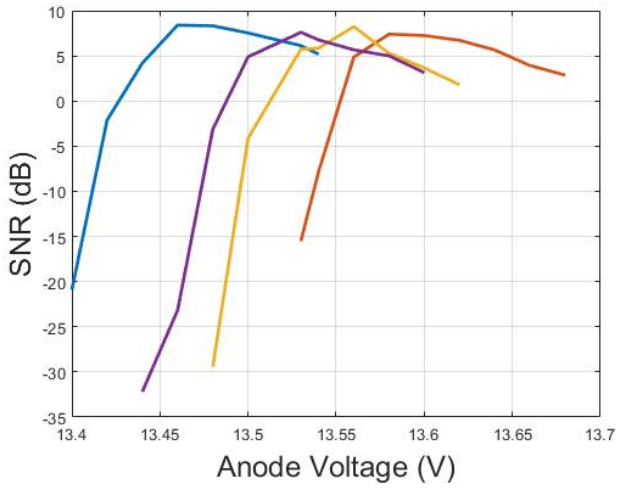

Figure 2.12: Estimated SNR of SPAD photodetection for a photon flux of 1M photons per SPAD per sec.

is because as bias voltage increases, gains in PDE come at the expense of increased DCR. The bias voltage at which peak SNR is achieved is insensitive to photon flux. An increase or decrease in photon flux simply shifts the SNR curves vertically without changing the shape of the SNR curve or the location of its maximum. To the best of our knowledge, this result has been reported in literature. Typically, optimization of SPADs focuses on either minimizing DCR or PDE, not SNR. This contribution is particularly significant for intensity-based measurements that do not rely upon temporal correlation techniques (such as gating) as a way to reduce effective DCR.

2.4.7 Pixel Mismatch

SPADs across the SPADCAM image sensor suffer from much higher levels of mismatch than traditional APS. A major source of mismatch that is unique to SPADs is variability in breakdown voltage. Breakdown voltage varies due to: 1) effectiveness of the poly-Si guard ring in suppressing premature edge breakdown, 2) doping concentrations and their effect on depletion region width

and field strength, and 3) system-level considerations such as Ohmic voltage drops in wiring. These variations are observed in PDE curves that show nearly a 150mV spread in the anode voltages required to achieve equivalent PDE (see Fig.2.11). This corresponds to a standard deviation of 42mV, or <1% of the anode voltage, from a limited sample size of 16 SPADs. We see a comparable result in work by Antolovic *et al.* [83] which reports a standard deviation in breakdown voltage of 67mV across a 512x128 pixel array.

Another major source of mismatch is variability in defect density across the image sensor. This mismatch is present in traditional CMOS image sensors and gives rise to FPN. DCR distributions across SPAD image sensors can generally be modeled with a normal distribution [83]. The high sensitivity of SPADs gives rise to a high sensitivity to DCR. Across SPAD literature we see particularly noisy SPADs referred to as "hot" pixels. These are SPADs with high DCRs that typically fall outside of the normal distribution of the SPAD array. These SPADs are often ignored or shut off for imaging applications as they can dominate readout resources and generate noisy data. Across one row of our image sensor we see a mean DCR of 1.77 Hz/μm^2, a median DCR of 0.93 Hz/μm^2, and a high standard deviation of 1.77 Hz/μm^2 for an anode voltage of 13.58V. Generalizing statistics is not possible for such a small dataset. However, we can clearly separate out noisy SPADs (DCR>3 Hz/μm^2) from "typical" SPADs (DCR<1.5 Hz/μm^2). The second population accounts for nearly 11 out of our 16 SPADs (68%) and have a mean DCR of 0.71 Hz/μm^2, a median DCR of 0.63 Hz/μm^2, and a standard deviation of 0.32 Hz/μm^2 for an anode voltage of 13.58V. Of the four noisy pixels, two have a DCR between around 3.6 Hz/μm^2 and two have a DCR in excess of 5 Hz/μm^2.

2.4.8 Low-Light Optic Flow

In spite of the wealth of research on optic flow algorithms, optic flow hardware, and low-light imaging hardware, there is little to no work discussing the integration of low-light imaging hardware with optic flow detection. SPADCAM provides a platform to pursue this work. We focus on detection of 1-D optic flow using a single row of our 16x16 array. We implement optic flow processing off-chip to enable testing of multiple algorithms and image processing techniques. The ultimate goal of this work is to lay the groundwork for an embedded SPAD-based, low-light optic flow sensor.

Algorithm Selection

In our exploration of low-light algorithms, we focus on block matching, image interpolation, and three biologically-inspired correlation-based motion detectors (elementary motion detector, non-directional multiplication, and non-directional summation). A detailed description of each algorithm is provided in Chapter 4 along with benchmarking results for each algorithm at low-light levels. We ultimately focus on off-chip optic flow detection using the two non-directional correlation-based motion detectors: non-directional multiplication (NDM), and non-directional summation (NDS). These algorithms clearly outperform the elementary motion detector (EMD) in photon-limited environments. Block matching (BM) struggles to detect motion with low-resolution images and is not well-suited for 1-D optic flow on a 16x1 array. Also, pixel mismatch and the presence of "hot" pixels introduce significant biases in block matching which degrades overall performance. Performance of the image interpolation algorithm (IIA) also degrades substantially due to pixel mismatch. Neither IIA nor BM implement techniques for mismatch compensation and both assume that image contents remain constant between sub-

sequent frames. This assumption breaks down when significant variations in PDE and DCR exist. NDM and NDS, on the other hand, initially process SPAD signals with a high-pass filter (HPF). This approach has been demonstrated throughout biological systems as a method for overcoming mismatch between neurons. Additionally, these algorithms are correlation-based and more sensitive to temporal constant than to absolute intensity measurements. As a result, they are well-suited for noisy SPAD arrays.

Experimental Setup

The optical setup is similar to the setup for characterization of PDE and DCR. All measurements are performed in an optically isolated Faraday cage. A monochromator illuminates the sensor with a monochromatic light beam at a wavelength of 700nm and optical power of 93nW. Image motion is produced by passing a metal plate above the image sensor. The plate travels between the optical fiber output of the monochromator and the image sensor. The metal plate passes back and forth, parallel to the orientation of rows in the SPAD sensor, and modulates the photon flux of the imager. The metal plate is connected to a 3D printer (Borlee Mini01 Desktop Compact 3D Printer) which is programmed to generate 1-D motion at velocities ranging from 0-16.7 pixels/frame, or 1-10mm/sec.

Image Processing

A single frame corresponds to a 16x1 image. Prior to optic flow processing, each frame is spatially-pooled from 16x1 to 12x1. This filter is applied for two reasons: 1) to reduce pixel offset, and 2) improve SNR of optical signals. The spatial filter is an adaptation of the triangular image filter with dimension 5x1. The 12x1 spatially-pooled frames are then processed by the NDM and NDS algorithms. Both algorithms first apply a high-pass filter followed by a low-pass

filter (LPF) to each pixel in the 12x1 array. The output of the LPF is a temporally delayed version of the HPF output. The final stage of each algorithm is to correlate the HPF output of each pixel to the delayed outputs of its two neighboring pixels. Since only 10 of the 12 pixels have two neighbors, the output vector of the NDM and NDS algorithms have dimension 10x1. A visual representation of these algorithms is presented in Fig. 4.4(b) and Fig. 4.4(c).

Optic Flow Response

Biologically-inspired algorithms are traditionally characterized with sinusoidal gratings [84]. As a result, analysis of their transfer functions and overall characterization is usually presented exclusively for such visual stimuli. However, this analysis is generally applied due to its mathematical simplicity and ease of description. Such algorithms are, in fact, well-suited for a wide-range of images, not exclusively sinusoidal gratings. A cartoon representing the signal processing flow of a single pixel is shown in Fig. 2.13. A visual stimulus that consists of a rising or falling edge can produce a sharp peak at the output of the NDS as well as the NDM and EMD algorithms. The image processing for optic flow detection of such stimuli is identical to that of sine gratings. However, the flow estimate is measured by the mean algorithm response for sine gratings and the peak response for stimuli such as the one illustrated in Fig. 2.13. In our experimental setup, we generate a visual stimulus that mimics a square wave with small spatial frequency (<1 cycle/frame) instead of a sine wave. Thus, in order to characterize each algorithm we look at peak response produced by the moving edge of the stimulus.

Experimental Characterization

Images are captured at a frame rate of 10fps. Frame rate is limited by the relatively slow operating speed of the micro-controller (Teensy3.2). The micro-controller is equipped with an

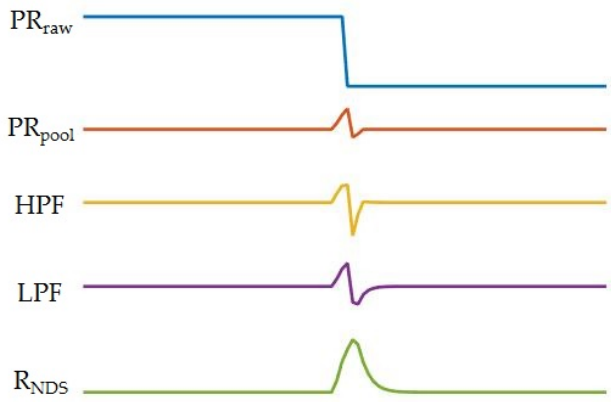

Figure 2.13: Cartoon representing signal processing flow for a single pixel in the NDS algorithm. Raw pixels (PR_{raw}) are spatially pooled (PR_{pool}) and then sequentially processed in a HPF and LPF. The final step (R_{NDS}) correlates HPF with LPF from its two neighboring pixel with addition.

integrated ADC used to read analog outputs from the image sensor and is used to generate control signals for the sensor. As previously mentioned, the image sensor is able to operate at frame rates in excess of 1kfps. Since parameters such as photon flux, frame rate, PDE, and fill-factor may vary drastically depending on environment and sensor design, a more universal description of low-light performance can be provided in absolute number of photons required for a particular operation. Thus in our analysis of algorithm performance, we quantify performance as a function of photons/frame or photons/pixel/frame.

As previously discussed, we characterize non-directional optic flow for simple square wave stimuli. Stimuli produce high contrast edges between "dark" and "bright" light intensities. We configure the experimental setup so that "dark" conditions correspond to approximately 27 photons/pixel/frame and "bright" conditions correspond to approximately 65 photons/pixel/frame.

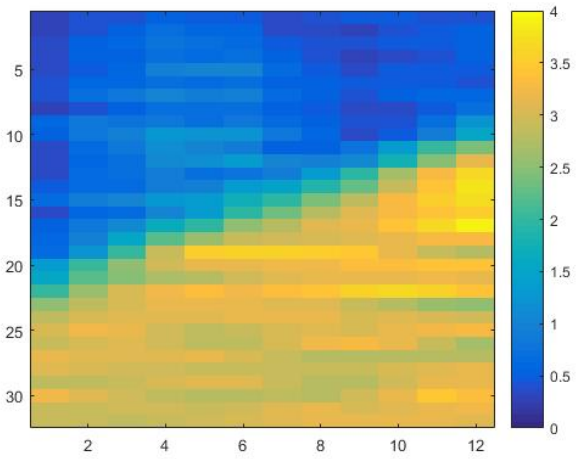

Figure 2.14: Visualization of 32 consecutive 1-D frames capturing a moving edge. The x-axis corresponds to pixel number for a 12x1 spatially-pooled frame. The y-axis corresponds to frame number of 32 consecutive frames.

This corresponds to an image contrast of 64%. Fig. 2.14 visualizes 32 consecutive frames capturing a moving edge of the visual stimulus. Blue pixels corresponds to "dark" pixels that are partially occluded while yellow pixel correspond to "bright" pixels. Total photon flux corresponds to roughly 1k photons/frame when the sensor is fully illuminated.

Optic flow processing produces frames of size 10x1—each "pixel" corresponds to the output of an optic flow detector. Similar to Fig. 2.14, we can visualize the output frames from the NDM and NDS algorithms. Fig. 2.15 and Fig. 2.16 visualizes 32 consecutive frames (same frames as shown in Fig. 2.14) corresponding to optic flow response of the NDM and NDS algorithms. Each algorithm displays a strong response (yellow) as the moving edge of the visual stimulus is detected and a weak baseline response (blue) at all other times.

Fig. 2.17 shows the normalized transient response of a single spatially-pooled pixel (PR_{pool}). It clearly displays a difference in photon flux between dark and bright scenes. The curves for

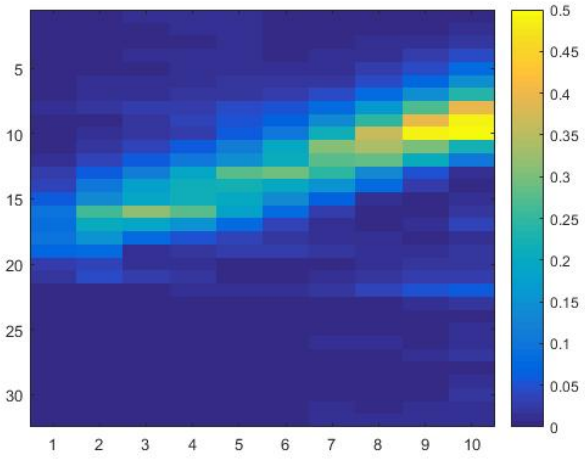

Figure 2.15: Visualization of 32 consecutive 1-D frames capturing optic flow response of the NDM algorithm during a moving edge. The x-axis corresponds to detector number for a 10x1 detector array. The y-axis corresponds to frame number of 32 consecutive frames.

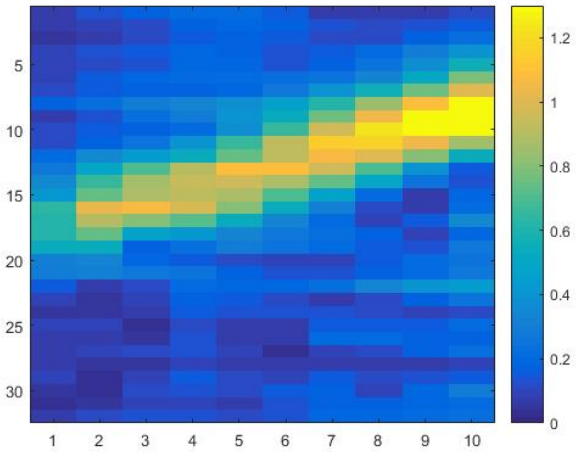

Figure 2.16: Visualization of 32 consecutive 1-D frames capturing optic flow response of the NDS algorithm during a moving edge. The x-axis corresponds to detector number for a 10x1 detector array. The y-axis corresponds to frame number of 32 consecutive frames.

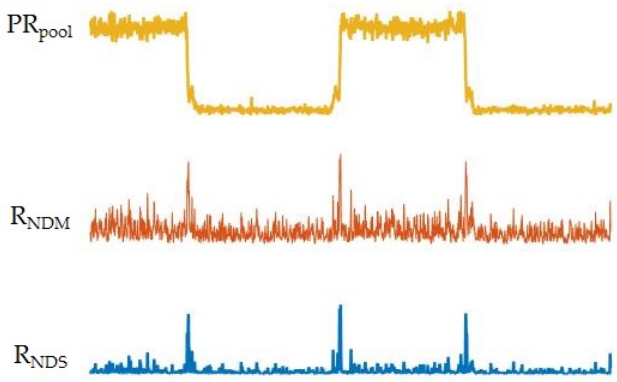

Figure 2.17: Transient response of a single spatially-pooled pixel (PR_{pool}) and its corresponding optic flow detector outputs (R_{NDM}, R_{NDS}). Amplitudes for the three signals are normalized to improve reader clarity. Image velocity corresponds to 1.67 pixels/frame.

R_{NDM} and R_{NDS} are also normalized and show a strong response to both rising and falling edges of the stimulus (motion right and left).

As we sweep image velocity from 0-16.7 pixels/frame, we observe variations in peak response of each algorithm. Both the NDS and NDM algorithms are configured such that the time constant of both the HPF and LPF are equal. For NDM, we set the time constant to 60ms. For NDS, we set the time constant to 40ms. These time constants are selected as they produce transfer functions that are able to clearly discriminate static and dynamic scenes, they produce a fairly linear response to small image velocities (<4 pixels/frame), and overall are well-suited for the visual stimuli. The transfer function for the NDM and NDS algorithms are provided in Fig. 2.18 and Fig. 2.19, respectively. As discussed in Chapter 4, the NDS algorithm offers a higher mean SNR than the NDM algorithm across the velocity range. We measure a mean SNR of 12.6dB for the NDM algorithm and a mean SNR of 15.7dB for the NDS algorithm.

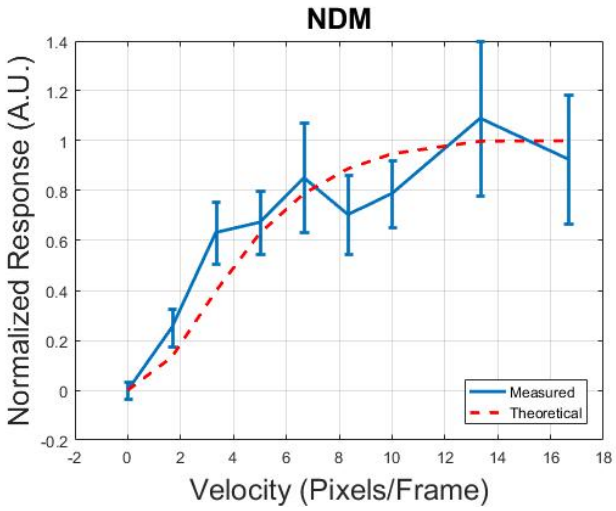

Figure 2.18: Transfer function for NDM algorithm in response to 1-D motion. Image velocity is swept from 0-16.7 pixels/frame. Error bars denote standard deviation in velocity measurement. The curve labeled theoretical corresponds to the ideal response of the NDM algorithm for a particular sinusoidal algorithm and is included to show a general agreement in shape of the measured transfer function. NDM algorithm has a time constant of 60ms for all filters.

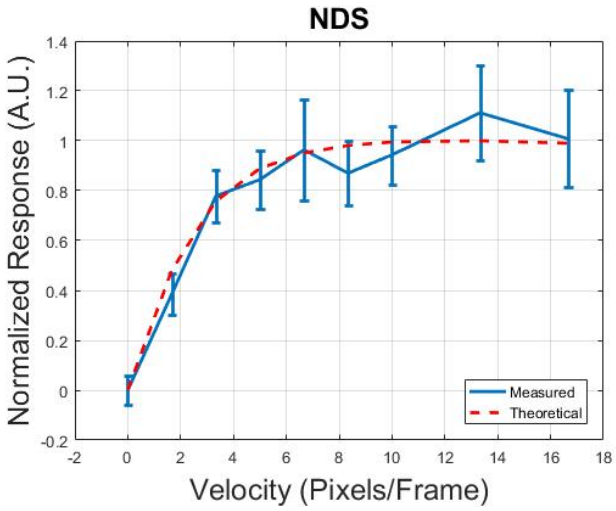

Figure 2.19: Transfer function for NDS algorithm in response to 1-D motion. Image velocity is swept from 0-16.7 pixels/frame. Error bars denote standard deviation in velocity measurement. The curve labeled theoretical corresponds to the ideal response of the NDS algorithm for a particular sinusoidal algorithm and is included to show a general agreement in shape of the measured transfer function. NDS algorithm has a time constant of 40ms for all filters.

Conclusion

We experimentally demonstrate 1-D non-directional optic flow with a SPAD-based image sensor at light levels corresponding to an average of <100 photons/pixel/frame and image contrast of 64%. This is the first known demonstration of optic flow with a SPAD image sensor as well as the first demonstration of optic flow detection at such low-light levels.

Chapter 3

Room Temperature SWIR Sensing[‡]

Short wave infrared (SWIR) imaging is widely used for low-light and night vision applications. Typically we define SWIR photons as having wavelengths between 900nm and 1700nm and energies between 1.3eV and 0.7eV. Due to the longer wavelengths of SWIR photons, compared to visible photons, they benefit from reduced Rayleigh scattering and enable "see through" imaging in fog, haze, and rain.

SWIR night vision systems benefit greatly from a phenomenon known as nightglow. This refers to the process by which molecules in the atmosphere become excited through the absorption of solar radiation during the day and then return to their low-energy states at night, radiating photons as they relax. The radiation spectrum is broadband and stretches from the

[‡]Excerpts from a manuscript in preparation for submission to *IEEE Sensors Journal* as **A. Berkovich**, G. Barrows, P. Abshire, "Improving Signal-to-Noise Ratio (SNR) for Room Temperature SWIR Sensing Through Adaptive, Floating Gate CTIA Photocircuit", 2017 and a conference paper published as **Andrew Berkovich**, Alexander Castro, Mohammad M. Islam, Fow-Sen Choa, Geoffrey Barrows, and Pamela Abshire, "Dark Current Reduction by an Adaptive CTIA Photocircuit for Room Temperature SWIR Sensing", *International Symposium for Circuits and Systems (ISCAS)*

visible spectrum across the infrared spectrum. However, the spectrum has large peaks in SWIR spectrum and throughout the IR spectrum which correspond to transition energies of airborne molecules such as CO_2, OH, H_2O, and O_2. These peaks provide significantly higher illumination intensities than starlight and visible wavelengths at night, and has contributed to the widespread use of SWIR image sensors for night vision applications. For applications where nightglow is unavailable, such as indoors, active illumination may be possible. In such applications, SWIR imagers are well suited for covert or stealth applications as SWIR photons are invisible to the human eye.

3.1 SWIR Detectors

Much like the photodetectors discussed in Chapter 2, SWIR sensors are diverse in their sensing mechanisms and physical composition. Common sensor structures include microbolometers [85], image intensifiers [86], photodiodes [87], avalanche diodes [88, 89], and quantum devices [90, 91]. Today, the large majority of commercially available SWIR cameras utilize either microbolometers or photodiodes [92–94]. Research and development of SWIR detectors is largely driven by applications in security, defense, and health care. These R&D efforts largely focus on techniques for reducing the high cost of SWIR detectors, improving device reliability, and integrating detectors into high-resolution, 2D arrays.

3.1.1 Microbolometers

The origins of microbolometers go back to the 1950's [95]. The simplest bolometer is a single, homogeneous, semiconductor used for thermometry. Infrared radiation is detected by measuring changes in resistance of the semiconductor that result from changes in temperature of the mate-

rial. This change in temperature corresponds to absorption of infrared photons. Recent work has improved the sensitivity and overall performance of these sensors through the use of composite and MEMS structures, as well as through the optimization of material fabrication [95–97].

Bolometers are used for optical sensing across the infrared and millimeter spectrum and offer a viable alternative to photodiodes in many applications. The distinguishing difference between bolometers and photodiodes is that bolometers absorb and thermalize photons. Due to rapid thermalization, bolometers detect the elevated detector temperature. Photon detectors (e.g. photodiodes) extract thermally-excited carriers before they thermalize.

3.1.2 Photodiodes and Avalanche Photodiodes

Photodiodes and avalanche photodiodes are described in more detail in Chapter 2. SWIR diodes are often fabricated with III-V semiconductors (InGaAs, InP). SWIR APDs and SPADs have been demonstrated by various groups [98–101]. Inevitably, these detectors are plagued by high levels of noise (after-pulsing, cross-talk, dark count rates) and rely upon cooling or temporal gating to improve sensing capabilities. These technologies suffer from relatively new and immature fabrication techniques which leads to high sensor costs and unreliable sensor performance. As a result, materials such as black silicon and germanium have garnered significant attention for SWIR sensing [102,103]. These material systems offer substantially more mature fabrication techniques and lower costs which has motivated efforts to extend their spectral response deeper into the IR spectrum.

3.1.3 Quantum Detectors

Thanks to rapid advances in materials science, novel SWIR photodetectors based on nanostructures, quantum devices, and emerging materials have become an increasingly popular ap-

proach [104–107] to IR sensing. These systems exploit narrow-band gap materials such as IV-VI semiconductors (PbS, PbSe) and tunable structures like HgCdTe and graphene [104–107].

The biggest challenges with such approaches are the integration of these devices with CMOS technology, and scaling from individual detectors to high-resolution 2-D image sensors. Demonstrations of quantum SWIR detectors often present results for individual detectors, not arrays. Increasingly scaling and system integration are becoming the focus of work in this area and these material systems offer a promising alternative to more mature technologies such as bolometers and photodiodes.

3.1.4 Dark Current in Photodiodes

One of the biggest hurdles in implementing SWIR sensors is the high dark current that is typically associated with narrow-band gap photodetectors. A major focus of ongoing research is the development of SWIR sensors that are capable of operating at room temperature. Substantial efforts have been made in the area of device fabrication and optimization in order to minimize dark current [108–112] at the device level. At the system level, thermal cooling is often used to further boost device performance and reduce dark currents. However, it is important to note that for mobile, low-power application, thermal cooling often exceeds size, weight, and power constraints of such systems.

At the system-level, a variety of active pixel sensors and sensing techniques, such as current skimming and signal compression, [113–118] have been demonstrated to mitigate the effects of dark noise. The goal of such work is typically to extend dynamic range or improve uniformity across the image array. However, it is important to note that such techniques simply compensate for dark current without actually reducing it. Thus, dark current still contributes shot noise and limits detection sensitivity.

In this chapter we present a novel approach to dark current reduction that precisely adjusts the measurement conditions so that dark current does not affect the measurement in the first place. We demonstrate a photocircuit that holds the applied bias voltage of a photodiode near zero volts (within 1mV). This approach has the advantage of reducing dark current so that shot noise from the dark current does not affect the measurement in the first place.

3.2 Capacitive Transimpedance Amplifier Pixel

The capacitive transimpedance amplifier (CTIA) photocircuit uses a transimpedance amplifier (TIA) with a capacitor in the feedback loop to integrate photocurrent. Typically, one of two configurations is used with this circuit. Either the photodiode's anode and cathode are placed between the inverting and non-inverting terminal of the amplifier, or one of the diode terminals is connected to the non-inverting terminal while the other is connected to an external reference voltage (see Fig. 3.1).

In Fig. 3.1a, the voltage across the diode is controlled by the external reference voltages DAC and VP. This configuration is often used when it is desirable to apply, either a large reverse bias (>1V) across the diode, or a dynamic bias (e.g. modulation). In Fig. 3.1b, the voltage across the diode is zero for an ideal amplifier, but practically is set by the input offset of the amplifier. This configuration is used when it is desirable to apply a small, static reverse bias across the diode (<1V).

3.2.1 Amplifier Gain

An important consideration for the design of the amplifier in the CTIA photocircuit is the gain. The input-output relationship is $v_{out} = A_v(V_+ - V_-) = A_v \cdot dV$. This means that as the CTIA

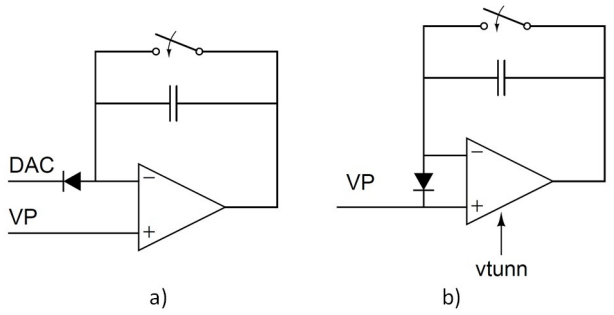

Figure 3.1: a) Non-adaptive CTIA circuit configuration, b) Adaptive CTIA circuit configuration

pixel discharges during integration, the input differential must be nonzero to allow the output to change. As the input offset grows, the diode bias voltage changes. For large reverse biases (>1V), this effect may be negligible. However, for bias voltages comparable to the turn-on voltage of the diode this can have significant consequences. This effect is particularly relevant for the CTIA configuration shown in Fig. 3.1b.

In Fig. 3.2 we simulate this effect assuming an initial amplifier offset (diode bias) of 1mV. This simulation assumes an exponential I-V curve for the diode, a saturation current of 100pA, and an integrating cap of 100fF. We plot discharge curves for different amplifier gains ranging from 40dB to 90dB. We see that for high-gain amplifiers (90dB) the discharge is highly linear, as we expect. This linearity indicates that the discharge rate is fairly constant throughout the integration period and is equal to the dark current at a bias of 1mV. For low-gain amplifiers, we clearly see a non-linear discharge curve. This indicates a substantial increase in dark current during the integration time and is a direct result of the growing differential input required to change the output. For an output swing of 3.3V, the input differential changes by 33mV with a 40dB gain, and 33μV with a 100dB gain.

When designing CTIA structures we ensure that the differential amplifier has a gain of >100dB to minimize the effects of varying input differential voltage. It is important to note

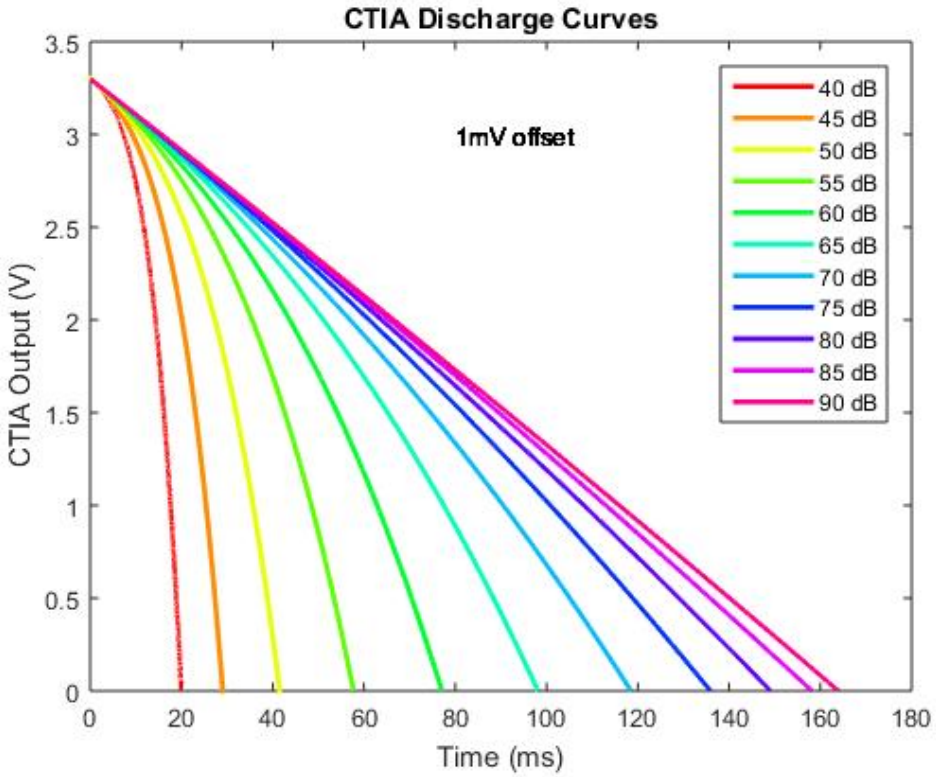

Figure 3.2: Simulation of a CTIA pixel discharging with an initial input offset of 1mV for different values of open-loop gain.

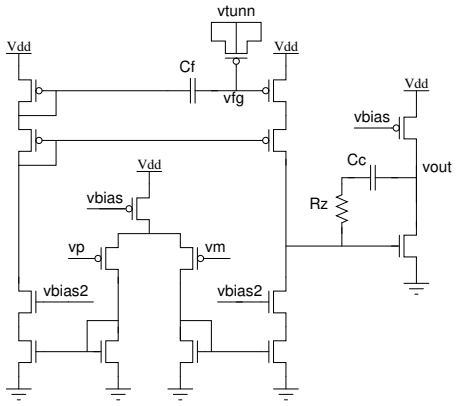

Figure 3.3: Circuit schematic of offset canceling adaptive amplifier.

that this comes at the cost of increased noise. This is an important trade-off to explore but is not discussed in this chapter.

3.2.2 Amplifier Offset

Typical amplifier input offsets are on the order of 10mV. Due to the exponential relationship between current and voltage, an input offset of 10mV versus 1mV can alter dark current by more than an order of magnitude. Thus, input offset can be a limiting factor in low light detection and motivates this work.

We modified a standard amplifier circuit to include a floating gate (FG) current mirror, as shown in Fig.3.3. This FG enables us to compensate for input offset in the amplifier by precisely adjusting currents in the current mirror under specific bias conditions. As the specific bias conditions diverge, the offset cancellation will be less accurate. In this chapter we demonstrate the feasibility of this approach in a CTIA photocircuit implemented in a 0.6μm CMOS process.

Many discrete amplifiers ICs offer offset null capabilities to reduce input offset voltages to just several μV and are often used in instrumentation applications. Jandu *et al.* demonstrate an architecture with an auto-zero amplifier in a feedback loop to limit input offset to approximately

2μV and fairly high power consumption of 5mW [119]. Also, several patents discuss techniques for minimizing input offset [120–122].

The fundamental approach underlying most offset null or input-offset minimization techniques are precisely tuned current sources that sink or source current to compensate for mismatch across the circuit. While different approaches may implement a wide range of techniques for programming these current sources, including feedback loops, adaptive circuits, and digital tuning, the underlying principles are similar. We employ this basic principle in our adaptive CTIA structures. The advantages of using a single floating gate device for input-offset compensation are: 1) low circuit complexity, 2) minimal increase in total circuit area and power consumption, and 3) scalability of circuit architecture.

3.2.3 Floating Gate Devices: Overview

Floating gate devices are commonly available in standard digital CMOS technologies and provide local nonvolatile analog storage with long storage lifetimes [123] (1-100μV voltage drop over 10 years), and high-resolution (14-bit) memory cells [124]. They are the basis for EEPROM and have been used in a wide range of adaptive circuits including imagers [125], data converters [126], field-programmable analog arrays [127–129], and amplifiers [130]. Carver Mead and Misha Mahowald can largely be credited for the first use of floating gate devices in adaptive neuromorphic sensing and computing systems [131, 131]. This line of research continued for well over two decades [132, 133] and still is pursued today [123]. The work reported in this chapter is the first known use of floating gates in readout integrated circuits for SWIR sensors.

3.2.4 CTIA Operation

We compared the adaptive CTIA circuit with a non-adaptive version, which is identical except that capacitor Cf, shown in Fig. 3.3, is shorted. Both CTIA structures include a 10pF integration capacitor as shown in Fig. 3.1. Monte Carlo simulations show that the amplifier has an average open-loop gain of 130dB with a std of 3.4dB and an average input offset of -18.7mV. This ensures a gain of $>10^6$ and means that a change in the output of 3V corresponds to a change of $<3\mu$V at the input.

3.3 Experimental Characterization: Dark Current

We performed a series of experiments to characterize the dark current of the CTIA photocircuit. The CTIA circuits were fabricated in a standard 0.6μm CMOS process; a micrograph is shown in Fig. 3.4. All circuits were placed in a box that provides electrical and optical shielding. We tested the CTIA structures with two different photodiodes, one InGaAs and one Si. The Si device is a commercially available photodiode (OSRAM #SFH229FA) with a dark current of 50pA at a reverse bias of 10V. The InGaAs device is a custom 4x4 array fabricated by Islam *et al.* [134] and has a smaller turn on voltage and much higher dark current density compared to Si diodes. We selected a sample with unusually high defect densities to represent a "worse-case" scenario with respect to dark current. This diode has a dark current of 170nA at a reverse bias of 100mV. All measurements were performed at room temperature.

3.3.1 Non-Adaptive CTIA Pixel

A schematic of the measurement setup is shown in Fig. 3.1a. We connect the photodiode anode to a 12-bit DAC. The DAC has a 3.3V rail and offers voltage resolution of <1mV. We measured

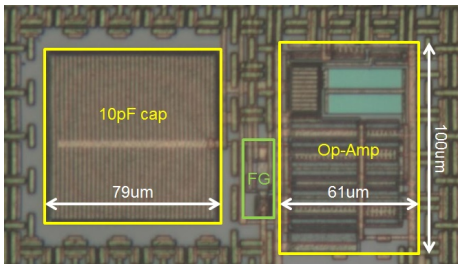

Figure 3.4: Micrograph showing an adaptive CTIA pixel. The adaptive op-amp and integration capacitor are highlighted in yellow while the floating gate (FG) devices are in green.

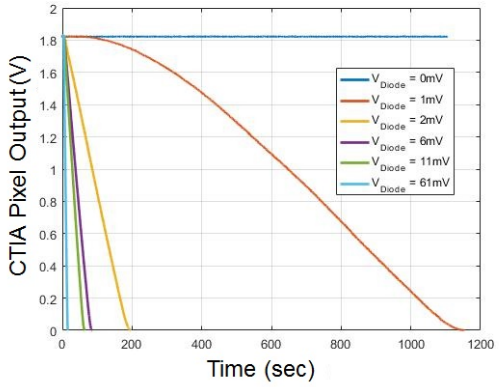

Figure 3.5: Transient response of CTIA pixel discharging with Si photodiode in the dark. V_{Diode} indicates reverse bias of the photodiode.

discharge curves of the CTIA pixel for various anode voltages. By sweeping the anode voltage we tuned the diode bias to within 1mV of its "optimal" zero bias condition.

Fig. 3.5 shows the transient response of the non-adaptive CTIA structure with the Si photodiode. The different curves are for different values of reverse bias voltages across the photodiode. As we raise the diode's anode voltage from 0V towards the amplifier reference voltage VP, the slope of the discharge curves decreases. This holds until the anode bias exceeds the cathode voltage at which point the diode becomes forward biased and can no longer discharge.

We estimate dark current from these curves by calculating the average slope ($I = C \cdot \Delta V / \Delta t$).

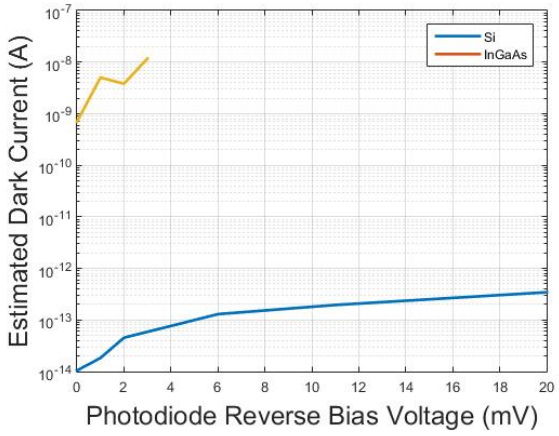

Figure 3.6: Dark current for InGaAs and Si photodiodes as a function of reverse bias voltage. Results are for the non-adaptive photocircuit.

We measure discharge time (Δt) as the time required for the CTIA output to discharge from 90% to 10% of its voltage range. As expected, by decreasing the applied bias we are able to suppress the dark current. For each diode we are able to decrease dark current by 1-2 orders of magnitude. It is worth noting that for the InGaAs diode we were not able to measure leakage currents for large input offsets since the dark currents were so large that discharge was nearly instantaneous and could not be adequately sampled in our experimental setup.

3.3.2 Adaptive CTIA Pixel

In the adaptive CTIA structure, the photodiode is placed between the inverting and non-inverting amplifier inputs, and the input offset sets the bias conditions for the diode. We tune the input offset by adjusting vfg. We cannot directly measure the input offset as any connection to the non-inverting terminal would introduce additional leakage current that would interfere with dark current measurements. The floating gate voltage is measured through a readout buffer. We repeated the same measurements as for the non-adaptive structure, except

instead of sweeping the diode anode voltage using a DAC, we adjust the diode bias by varying the floating gate voltage. A more detailed description of the floating gate operation is provided in the following section.

The discharge curves (not shown) appear similar to those of the non-adaptive structure. Again we estimate dark current by measuring the discharge time (Δt) from 90% to 10% of the voltage range. We demonstrate two states of the adaptive CTIA pixel: poorly tuned and finely tuned. For the InGaAs diode we define poorly tuned as having a floating gate voltage 15-20mV from its optimal value and finely tuned as within 1mV. We categorize our results in this way because, as previously mentioned, we cannot directly measure input offset and must estimate its value. The "optimal value" is defined as the floating gate voltage that produces an input offset of exactly 0V – which would, theoretically, produce an infinite discharge time. In these two cases we measure an average dark current of 5.8nA for the poorly tuned pixel and 2.25pA for the finely tuned pixel. The poorly tuned case performs slightly worse than the best-case scenario observed for the non-adaptive CTIA structure (0.7nA). The finely tuned case represents a reduction in dark current of 2-3 orders of magnitude over both the non-adaptive and poorly-tuned adaptive cases. Again, for large offsets produced by a floating gate >100mV from its optimal value, the dark current is so large that discharge is nearly instantaneous and cannot be measured in our experimental setup.

For the Si diode we define the poorly tuned case as having a floating gate voltage 10-20mV from its "optimal" value and finely tuned as within 5mV from its "optimal" value. We measure an average dark current of 1.6pA for the poorly tuned case and 33.8fA for the finely tuned case. The finely tuned case is close to the lowest observed dark current for the non-adaptive circuit (10.4fA).

3.3.3 Floating Gate Devices: Programming

As shown in Fig. 3.3, the floating node is connected to the gate of a MOSCAP that tunes the floating gate voltage via Fowler-Nordheim tunneling. The body of the MOSCAP is connected to an external high-voltage power supply *vtunn*. Increasing *vtunn* increases the probability that charge carriers can tunnel through the gate dielectric and increases the voltage of the floating gate.

To program the floating gate we first UV erase the chip and program the floating node at a specified high voltage. We use an off-chip power-FET to generate 1ms digital pulses with an amplitude of 14.15V. Each pulse increases the floating gate voltage by about 500μV. The programming resolution can be adjusted by varying the tunneling voltage or the tunneling pulse width. We note that in this implementation the floating gate voltage leaks at a rate of 1.5μV/sec, corresponding to $<1e^{-}$/sec. We believe that the cause for this is hot carrier injection current from the readout buffer of the floating node. This is significant for long integration times >1sec, but is negligible for more typical exposure times of 10-100ms. This effect can be mitigated by periodically "refreshing" the floating node during operation to keep it properly tuned.

Programming of the FG itself is limited only by GPIO speed of the connected micro-controller and can conservatively take a few milliseconds. The limiting factor in calibrating this circuit is the the long integration time of dark current required to determine the direction of FG adjustment. While this occurs on the time scale of minutes, it should only have to occur once prior to operation. We note that the photocircuit is calibrated at a specific temperature and, through Cadence simulations, found that a change in temperature of 1^{o}C shifts offset by approximately 1mV. Thus re-calibration would have to occur for significant changes in temperature.

Table 3.1: CTIA Dark Current Suppression Performance Summary

	Si	InGaAs
Non-Adaptive (best)	10.4fA	700pA
Adaptive (poor)	1.6pA	5.8nA
Adaptive (fine)	33.8fA	2.25pA

3.3.4 Summary of Results and Discussion

We compared dark currents measured for InGaAs and Si photodiodes using both adaptive and non-adaptive CTIA circuits over a range of photodiode bias conditions. Table 3.1 summarizes the results. We measured total power consumption of the photocircuit as 174μW for both the adaptive and non-adaptive circuits. The addition of a FG device and the tuning circuitry increase total amplifier area by about 7.5%.

As noted above, in this implementation we observed leakage on the floating node at a rate of 1.5μV/sec. The lowest observed dark current of 30fA in finely tuned Si photodiodes is limited by the floating gate leakage. These results were obtained for a discharge time of >1000 sec, which corresponds to a 1mV change in floating gate voltage. The lowest observed dark current of 2.25pA in finely tuned InGaAs photodiodes corresponds to a current density of 4.97nA/cm^2. This is an impressive result for a custom-fabricated sample with high defect density and demonstrates a reduction in dark current by 5 orders of magnitude relative to the measured dark current of 170nA for the same diode at a reverse bias of 100mV.

We intentionally designed a high gain amplifier to minimize offset accumulation, but that design choice introduces other types of noise and leads to trade-offs between dark current, amplifier gain, and dynamic range. It is important to consider all sources of noise including read noise, 1/f noise, kTC noise, shot noise, but this full analysis is beyond the scope of this

work. In future work we intend to examine the trade-offs between bias voltage, dark current, and quantum efficiency using responsivity and SNR measurements for SWIR detection.

These results provide the first demonstration of dark current reduction in SWIR sensors operating at room temperature by readout circuits that provide fine bias tuning. This provides an important proof of concept for a new SWIR readout concept.

3.4 Experimental Characterization: Signal-to-Noise Ratio

Using the same CTIA structures as discussed in Sec. 3.3, we extend the characterization of the non-adaptive CTIA structure to include SNR measurements. Again, circuits were placed in a box that provides electrical and optical shielding. We tested the CTIA structures with a different (less noisy) InGaAs photodiode and the same Si photodiode (OSRAM #SFH229FA). To enable optical measurements, a green LED is included in the shielded experimental setup to illuminate the photodiode. The light intensity of the LED is controlled with PWM output from a micro-controller (Teensy3.2) running at a frequency of 488Hz.

Using the non-adaptive structures configured as shown in Fig. 3.1a, we sweep the photodiode bias voltage from approximately 1mV to >1V. For each bias voltage, three consecutive measurements are taken. First, we measure pixel discharge in the dark to estimate dark current. Then, we measure pixel discharge at two different illumination intensities (LED1, LED2). Again we estimate diode current by measuring the discharge time (Δt) from 90% to 10% of the voltage range and calculate the average slope ($I = C \cdot \Delta V/\Delta t$). From these three measurement we are left with three diode currents (I_D, I_{L1}, and I_{L2}). I_D corresponds to the diode dark current, I_{L1} corresponds to the total diode current at illumination intensity LED1, and I_{L2} corresponds to the total diode current at illumination intensity LED2. We estimate the photocurrent for LED1

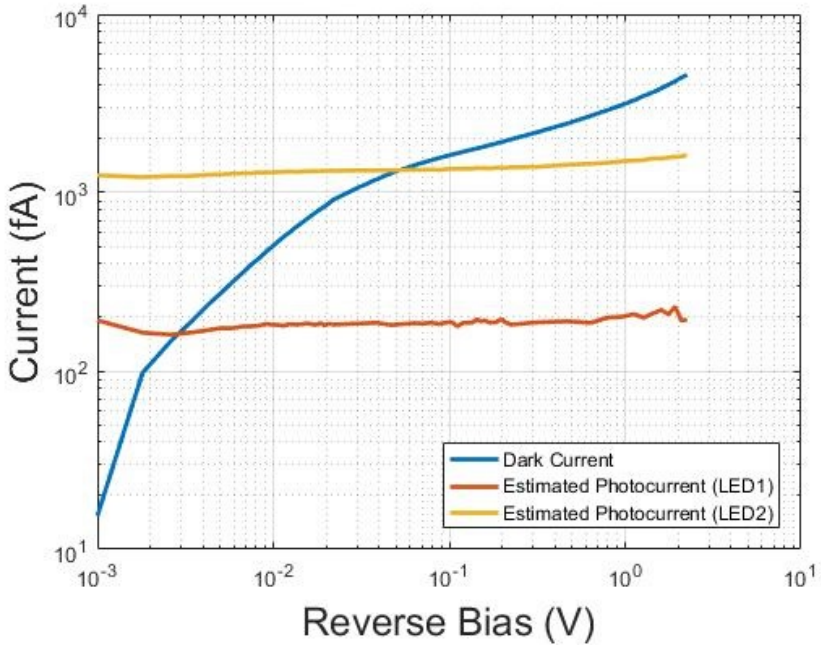

Figure 3.7: Estimated dark current and photocurrent for CTIA photocircuit with InGaAs photodiode. Reverse bias is controlled with a 12-bit DAC integrated into a Teensy3.2 micro-controller

as $I_{ph,1} = I_{L1} - I_D$ and the photocurrent for LED2 as $I_{ph,2} = I_{L2} - I_D$.

Fig. 3.7 and Fig. 3.8 show the plots for the dark current and photocurrents as a function of bias voltage for the InGaAs and Si diodes, respectively. We note that for both plots photocurrent is fairly constant with respect to bias voltage. This indicates that diode responsivity is not strongly dependent on bias voltage across these bias conditions. For the Si diode we note a reduction in dark current of approximately 1.5 orders of magnitude by tuning reverse bias from 20mV to 1mV. Similarly, we note a reduction in dark current of approximately two orders of magnitude by tuning reverse bias from 20mV to 1mV for the InGaAs diode. For both Si and InGaAs, dark current displays peak sensitivity to bias voltage near 1mV. This response is expected, following the exponential current-voltage diode relationship, and supports the need for accurate input-offset compensation with <1mV error.

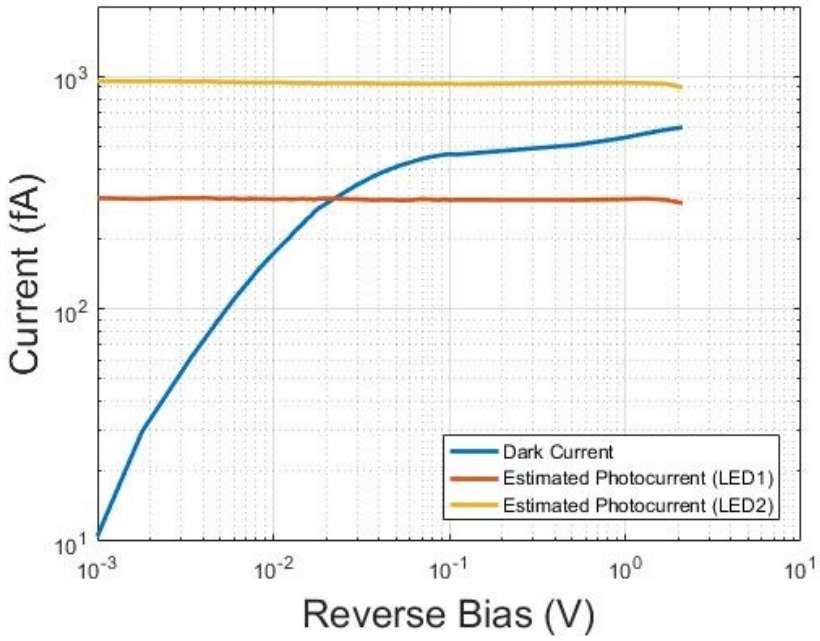

Figure 3.8: Estimated dark current and photocurrent for CTIA photocircuit with Si photodiode. Reverse bias is controlled with a 12-bit DAC integrated into a Teensy3.2 micro-controller

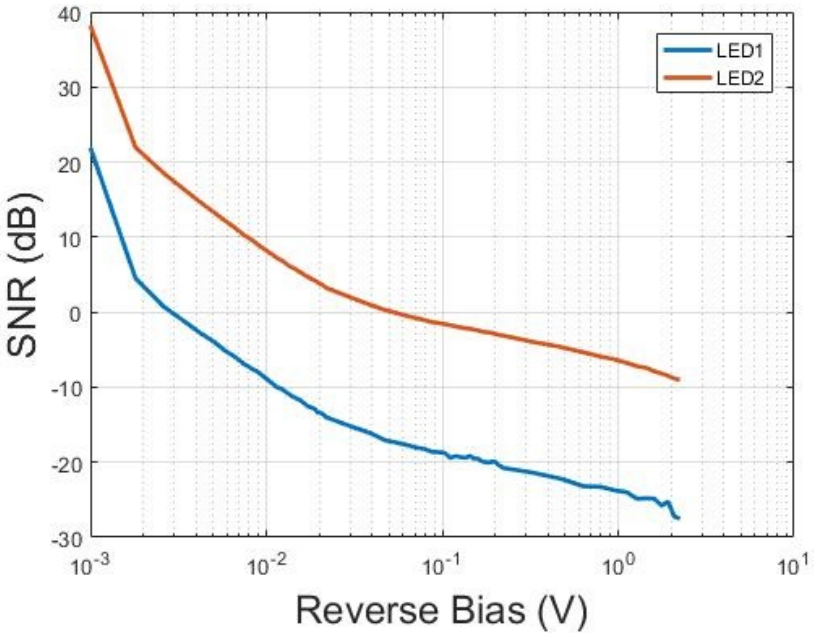

Figure 3.9: Estimated SNR for CTIA photocircuit with InGaAs photodiode illuminated at two different light intensities (LED1, LED2). Reverse bias is controlled with a 12-bit DAC integrated into a Teensy3.2 micro-controller

Fig. 3.9 and Fig. 3.10 show the plots for SNR as a function of bias voltage for the InGaAs and Si diodes, respectively. The two curves on each plot correspond to SNR measurements at different illumination intensities. SNR is estimated as the ratio of photocircuit to dark current. For the Si diode we note an improvement in SNR of approximately 30dB when reverse bias is adjusted from 20mV to 1mV. This 30dB gain agrees well with the 1.5 decade reduction in dark current and indicates that the dark current reduction does correspond to an improvement in SNR. For the InGaAs diode we note an improvement in SNR of approximately 40dB when reverse bias is adjusted from 20mV to 1mV. This 50dB gain agrees well with the 2 decade reduction in dark current and, again, indicates that the dark current reduction does correspond to an improvement in SNR.

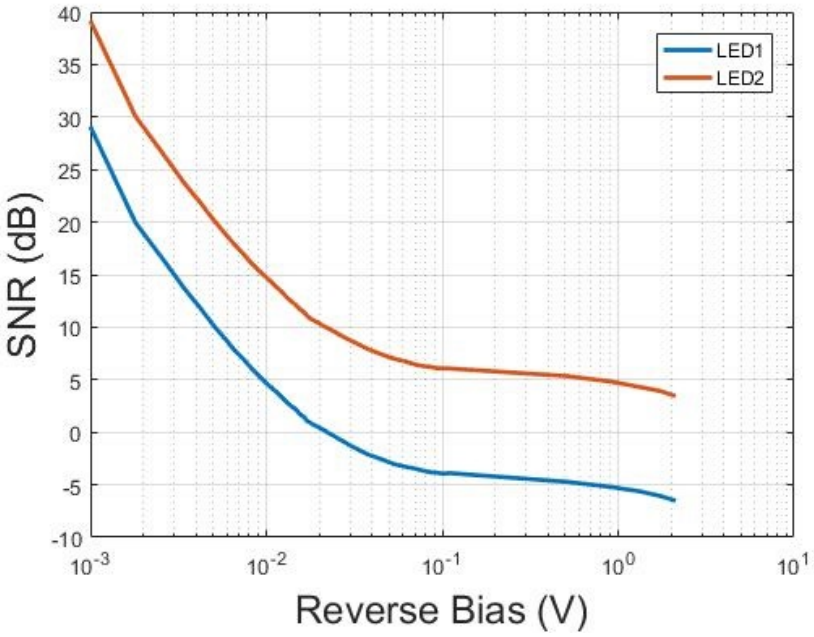

Figure 3.10: Estimated SNR for CTIA photocircuit with Si photodiode illuminated at two different light intensities (LED1, LED2). Reverse bias is controlled with a 12-bit DAC integrated into a Teensy3.2 micro-controller

3.5 Conclusions

In this chapter we demonstrate that we can precisely bias a CTIA photocircuit, using a 12-bit DAC, in order to compensate for input-offset and minimize photodiode bias (within 1mV). This reduction in photodiode bias voltage corresponds to a significant reduction in dark current of the photodiode while simultaneously boosting SNR. We demonstrate a 30dB and 40dB improvement for a Si and InGaAs diode, respectively, when comparing a 20mV bias voltage (typical amplifier input-offset) to a 1mV bias voltage. This is the first experimental demonstration of SNR improvement at near-zero bias conditions that we are aware of.

We also demonstrate that by incorporating adaptive circuit elements (floating gate transistor) to a differential amplifier in a CTIA photocircuit, we are able to compensate for input-offset with precision that exceeds 12-bit resolution (<1mV). This result is demonstrated by comparing minimum dark current achieved in a finely-tuned adaptive CTIA structure to minimum dark current achieved in a finely-tuned non-adaptive CTIA structure biased with a 12-bit DAC.

The combination of these results is significant because it suggests that floating gate compensation can be successfully implemented in room-temperature SWIR sensors. More importantly, this technique can be utilized to improve SNR for SWIR optical sensing. We show that adaptation can offer improved sensing capabilities (lower dark current) than programmable 12-bit DACs in each pixel. Moreover, this adaptive approach is scalable for high-resolution arrays, increases overall pixel size by <10%, and has negligible contribution to total power consumption of the photocircuit. A significant limitation of this approach is that it requires significant time for calibration of an entire array (minutes), and the photocircuit must be recalibrate, or adaptively tuned, if operating temperature deviate by more than 1-2^{o}C.

Chapter 4

Optic Flow Benchmarking[‡]

4.1 Introduction

Optic flow estimation is fundamental to nearly all vision systems, biological and artificial. It provides motion fields that are used in primitive, visually-based control systems as well as high-level motion analysis tasks. In the animal kingdom it is used by flying birds and insects for tasks such as self-stabilization [135, 136], odometry [137, 138], obstacle avoidance [7, 139], and speed control [6, 139]. It is also used by humans and other animals for control of locomotion [140] and path planning [141]. In computer vision, optic flow is used for tasks such as video compression, motion-based image segmentation, image stabilization, and object tracking [142–144]. In nearly all visual systems, it is used as a low-level characterization of motion upon which higher-level visual processing tasks are built.

Optic flow detection has been an active research topic for more than 40 years [145, 146]. Decades of research has improved the performance of optic flow algorithms with respect to

[‡]In preparation for submission to *Biological Cybernetics* to appear as: **A. Berkovich**, G. Barrows, P. Abshire, "Benchmarking Photon-Limited Performance of Biologically-Inspired Optic Flow Processing Algorithms", 2017

large image displacement [147] and preservation of multi-scale image features [148, 149]. Several groups have benchmarked the performance of optic flow algorithms with respect to both accuracy [150] and efficiency [151], and several standardized data sets (MPI-Sintel [152], KITTI [153], Middlebury [154]) have been developed to support benchmarking efforts.

In spite of these efforts, one area of research not yet addressed is the performance of optic flow algorithms at low light levels. This is particularly relevant for robotic applications in photon limited environments such as deep inside a building or cave, or underneath a forest canopy at night time. These environments present ambient light levels in which just thousands of photons per second may be available for image sensors. Current technology is unable to meet such capabilities. It is important to understand the limitations of various optic flow algorithms in photon limited environments in order to inform the design of future low-light imaging technology.

In this work we present a framework for this analysis. We quantify the performance of optic flow processing algorithms in photon-limited conditions. In contrast to traditional computer vision systems, we draw inspiration from flying, nocturnal insects like the *Megalopta genalis* and adopt a minimalist approach. We simulate the acquisition of optical flow over a wide field of view with low-resolution sensors (<100k pixels). Our goal is not to obtain optimal performance metrics, but rather to quantify the minimum photon rates necessary to obtain "useful" optical flow measurements, understand the roles of spatial and temporal pooling in photon-limited optic flow processing, and explore the effects of shot noise on different types of optic flow algorithms. This will in turn support a better understanding of limitations in sensing optic flow for future MAV technologies.

4.2 Image Generation

We benchmark the performance of several algorithms with both natural and synthetic scenes. Benchmarking of each optic flow algorithm is performed through MATLAB-based simulations. Details about images are provided in Sec. 4.2.1 and Sec. 4.2.2. Each benchmarking trial begins by generating a "template" image. This image is a panorama that represents a full 360^o field-of-view (FOV). By using a panorama we can circularly shift the template horizontally along its x-axis, indefinitely, without generating discontinuities in the image. To preserve uniformity among templates, each template is re-sized to 251xm pixels. The importance of the value 251 is described later in this section. The number of pixels along the x-axis (m) varies for each natural scene but templates have aspect ratios ranging from 2:1 to 4:1. For synthetic images, each template is 251x251.

Once a template image is generated we then produce a series of frames that are processed by optic flow algorithms. Known optic flow is generated from frame to frame by circularly shifting the template image by a defined amount. Using this approach we generate 400 sequential frames for all benchmarking simulations. This provides a continuous visual stimulus for each algorithm that physically corresponds to an observer rotating clockwise while standing in place. This approach does include parallax, when using natural images, but does not generate the complex 2D flow fields that result from complex, dynamic environments with multiple moving objects (see Fig. 4.1(b). Instead, global 1D motion is generated along the x-axis (see Fig. 4.1(a)).

The role of spatial pooling is investigated by resizing each frame prior to computing optic flow to either 60x60, 120x120, 180x180, or 240x240. Re-sizing is implemented with the MATLAB function *imresize* which applies a nearest-neighbor interpolation to reduce image dimensions. Template images are shifted by strictly integer values. Thus, spatially-pooled frames are shifted

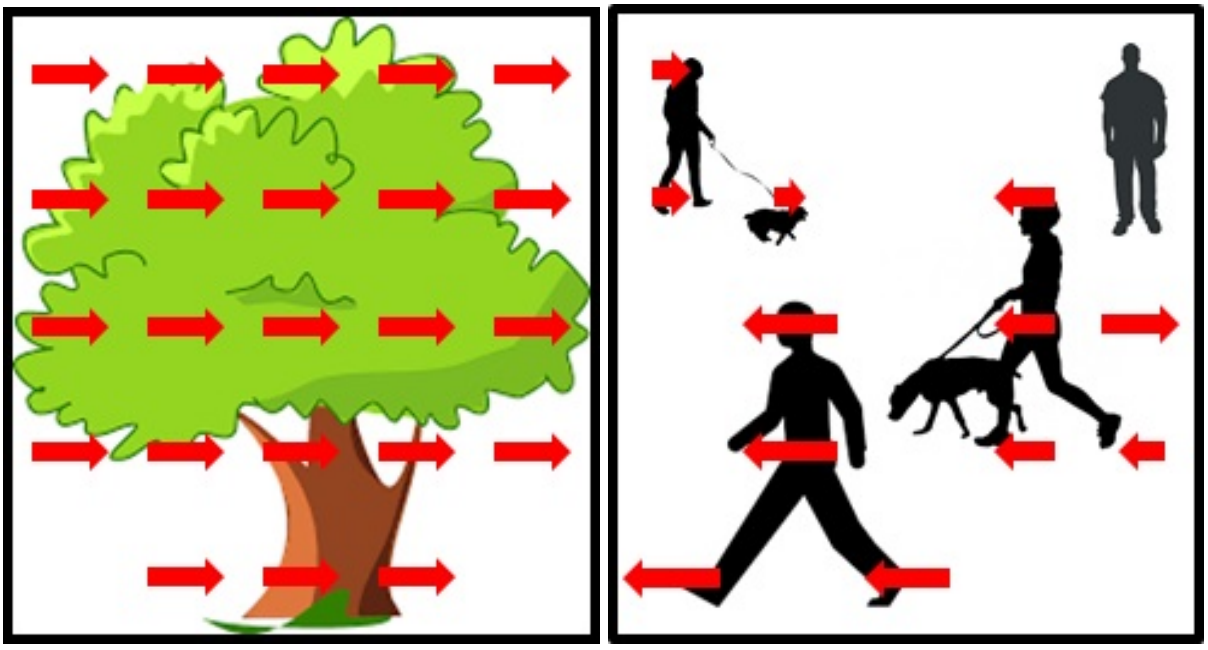

(a) (b)

Figure 4.1: Cartoon illustration of optic flow estimation in a simple scene (a) with global, 1D motion and a complex scene (b) with 2D flow generated by multiple objects moving at different velocities and depths.

by an effective image velocity of ν_{eff} which is given by $\nu_{eff} = \frac{\nu \cdot N}{N_{orig}}$. N_{orig} is the y-dimension of the template image and is set to a semi-arbitrary value of 251 to ensure that all effective image velocities are non-integer values which provides a more physically plausible simulation. We benchmark the accuracy of each algorithm over a range of velocities (ν = 1-10 pixels/frame). An image velocity of 5 pixels/frame translates to an effective velocity of 1.2 pixels/frame for a 60x60 image and 4.8 pixels/frame for a 240x240 image.

Using this approach we: 1) define a non-stochastic template image, 2) generate a sequence of non-stochastic frames based on known optic flow by circularly shifting the template image, 3) generate a stochastic image sequence from the non-stochastic image sequence produced in the previous step, 4) process both stochastic and non-stochastic image sequences using several optic flow algorithms, 5) statistically evaluate performance of each optic flow algorithm based on varying benchmarking parameters.

Photon arrival is modeled as a Poisson process. We incorporate dark count rates into our framework and ignore all other noise sources–thus we only explore the effects of shot noise and dark counts. Stochastic images are generated by converting pixel-level intensity values to Poisson random variables. The Poisson variables in each spatially-pooled frame are then normalized and scaled so that the sum of all Poisson variables in a frame is equal to the appropriate photon flux per frame. Next, each pixel-level Poisson variable is used to generate a random pixel value. This is done using MATLAB's Poisson random generator function (*poissrnd*). For a photon flux of 576k photons/frame, a spatially-pooled 60x60 frame will have an average Poisson value of 160 photons/pixel while a 240x240 frame will have an average Poisson value of 10 photons/pixel.

We assume all pixels are perfectly linear photon counting modules at low illumination intensities ($<10^7$ photons/sec). This assumption is reasonable since a Geiger-mode avalanche

photodiode operating at a modest detection rate of 20MHz has high linearity (R^2>0.999) up to 10^7 photons/sec. Thus even at "high" illumination intensities of 10^9 photons/frame, a single large detector in a 60x60 array will receive $<10^6$ photons/frame or $<10^7$ photons/sec at standard frame rates. For the simulations described in this paper we sweep target photon rate from 10^3 photons/frame to 10^{12} photons/frame with two data points per decade. While the upper end of this range would saturate the photodetectors, we include this data to ensure that we capture all relevant behaviors and effects in the simulations.

4.2.1 Synthetic Images

Synthetic images are used to study the narrow-band response of each algorithm and gain insight into their performance. While physically unrealistic, they allow us to test each algorithm at a well-defined image contrast and spatial frequency. All synthetic images are 2D sine wave gratings of the form $S(x, y) = C \cdot sin(2\pi \cdot f_x \cdot x + \phi_x) \cdot sin(2\pi \cdot f_y \cdot y + \phi_y) + 1$ where C is the Michelson contrast of the sinusoidal grating and f_x and f_y are the vertical and horizontal spatial frequencies of the gratings. We generate gratings of varying contrast and spatial frequency and always set $f_x = f_y$ for the sake of simplicity. We rotate sine gratings by 45^o to ensure all rows have non-zero spatial gradients.

4.2.2 Natural Images

Natural images provide physically realistic broadband visual stimuli. We use a variety of natural images shown in Figs. 4.3(a), 4.3(b), 4.3(c), 4.3(d). These images are 360^o panoramas gathered from Google Images capturing forest scenes. While our simulation framework enables benchmarking of any natural environment, we focus on forest scenes as they are particularly relevant in low-light applications for aerial platforms.

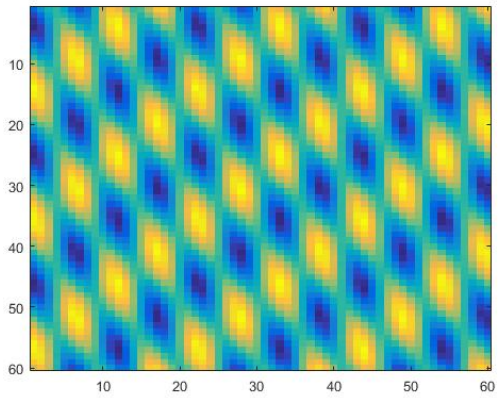

Figure 4.2: 2D sine grating used for synthetic images for $N_{pool} = 60$ and a spatial frequency of 4 cycles/frame.

4.2.3 DCR

For each benchmarking simulation we define a dark count rate (DCR). This DCR provides the average number of dark counts per pixel per frame in the original 251x251 pixel image. The DCR is incorporated in the spatial pooling process. Thus if the DCR for a simulation is 1 count/pixel/frame (for the 251x251 template image), a spatially-pooled image with $N_{pool} = 60$ would have 3600 total pixels each with a dark count rate of 17.5 counts/frame.

4.3 Optic Flow Algorithms

We explore three different classes of optic flow algorithms: block matching, gradient-based, and correlation-based algorithms. Our goal is not to optimize the performance of each algorithm but rather to benchmark the performance of each class of algorithm at low-light conditions. This work is in contrast to typical benchmarking experiments which look to optimize overall performance.

(a)

(b)

(c)

(d)

Figure 4.3: Gray scale panoramas of natural forest scenes. The original dimensions of each image are 1110x4000 for (a), 667x2000 for (b), 3327x7520 for (c), and 2180x6564 for (d). Each image is scaled down to create template images of size 251xm and a single 251x251 frame covers a FOV of 100.4^o for (a), 120.5^o for (b), 160.6^o for (c), and 120.5^o for (d).

4.3.1 Block Matching

Block matching algorithms are commonly used in video compression algorithms. They are traditionally considered to have few similarities with biologically-plausible circuits and algorithms. However, in their simplest form block matching algorithms may be realized using a structure similar to a neural network that implements a convolution layer and a winner-take-all layer.

The algorithm first divides each frame into sub-blocks. For each sub-block, a correlation is computed between that sub-block in consecutive frames. The location of the maximum correlation indicates the motion vector for that sub-block. The algorithm assumes uniform motion for all pixels within each block. We use an exhaustive search (ES) block matching algorithm adapted from Barjatya [155]. An implementation of this algorithm is provided on the MathWorks file exchange [156].

In the benchmarking experiments we use 100 sub-blocks for each frame and a search radius that varies depending on the level of pooling. We use a search radius of 3 pixels for $N_{pool} = 60$, 4 pixels for $N_{pool} = 120$, 6 pixels for $N_{pool} = 180$, and 7 pixels for $N_{pool} = 240$. We vary the search radius to account for the effect of spatial pooling on effective image velocity, as previously discussed. The choice of 100 sub-blocks per frame produces sub-blocks that are 6x6 pixels for a 60x60 spatially-pooled frame and sub-blocks that are 24x24 pixels for a 240x240 spatially-pooled frame.

4.3.2 Image Interpolation Algorithm

The Image Interpolation Algorithm (IIA) is a non-iterative gradient-based optic flow algorithm first proposed by Srinivasan [157] to measure egomotion. It computes six flow fields, three for translational motion and three for rotational motion, from consecutive frames (f_0, f). In this

work we implement a simplified adaptation of this algorithm that ignores rotational motion and focuses only on translational motion. An important assumption of this algorithm, like many gradient-based optic flow algorithms, is that image displacements are small between consecutive frames and motion is linear and continuous. Similar to block matching, we compute a motion vector for each of 100 sub-blocks for each frame.

4.3.3 Biological Correlation-Based Detectors

A number of biologically-inspired correlation-based detectors have been proposed based on neurophysiological and anatomical studies of insects. The best known example is the Hassenstein-Reichardt detector also known as the elementary motion detector (EMD). Unlike other types of OF algorithms, these algorithms produce a scalar number which does not correspond to a physical definitive velocity but a relative output of the algorithm—further, outputs are highly sensitive to contrast, spatial frequency, and temporal tuning. This, in addition to their non-linear bell-shaped transfer functions, means their outputs are ambiguous under conditions of unknown image contrast and spatial frequency.

In spite of these characteristics, biologically-inspired correlation detectors are widely used in mobile robotic systems [158–162]. These algorithms are computationally simple and can be efficiently implemented directly in the focal plane of an imager [163]. Furthermore, simple optic flow-based heuristics with these algorithms enable fairly sophisticated behaviors. Hardware implementations of these algorithms have been studied extensively and use of these algorithms in embedded systems goes back several decades [163].

Our analysis encompasses the classic elementary motion detector (EMD) as well as two non-directional variations characterized by Higgins and Dyhr [164, 165]. We hypothesize that the non-directional variations will outperform the directional EMD at low-light levels and thus look

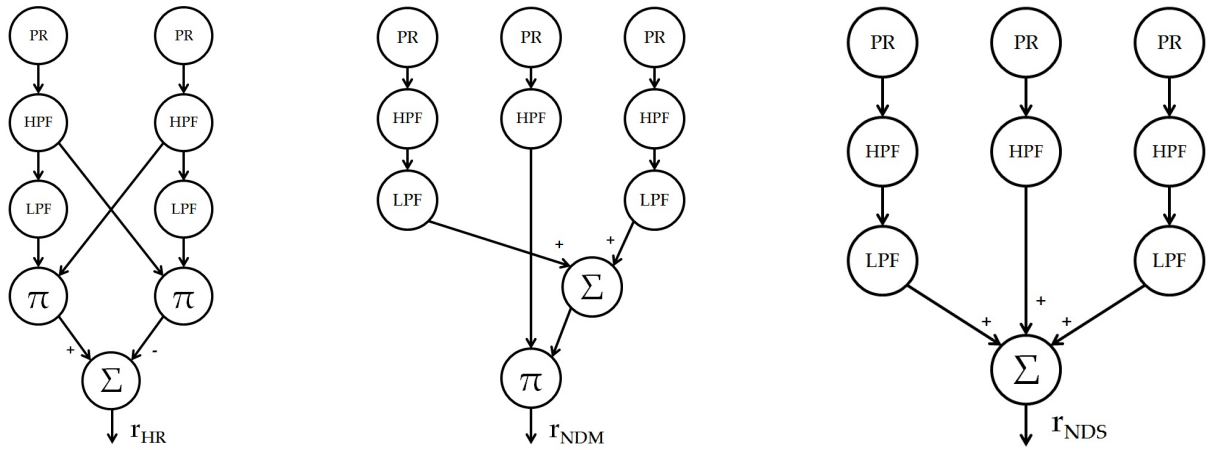

(a) (b) (c)

Figure 4.4: Computational model of the Hassenstein-Reichardt based elementary motion detector (EMD), shown in (a), non-directional multiplication (NDM) detector, shown in (b), and non-directional summation (NDS) detector, shown in (c). In each algorithm, the photoreceptor (PR) outputs first pass through a high-pass filter (HPF) to remove signal offset and produce zero-mean signals. Signals are then delayed with a low-pass filter (LPF) and correlated with either a multiplication block (Π) or a summation block (Σ).

to benchmark these algorithms in comparison to BM and IIA. The computational models for each correlation-based algorithm are shown in Fig. 4.4, and their transfer functions are provided in Sec. 4.3.3, Sec. 4.3.3, and Sec. 4.3.3. It is important to note that these transfer functions are all derived for moving sinusoidal gratings of the form $S(t,x) = .5 \cdot (1 + C \cdot sin(\omega_t \cdot t + \omega_x \cdot x))$, where C is the grating contrast, ω_t is temporal frequency, and ω_x is spatial frequency.

Elementary Motion Detector

The elementary motion detector (EMD), shown in Fig. 4.4(a), correlates the response of one photodetector to the delayed response of a neighboring photodetector. It is a "biologically-inspired" optic flow algorithm proposed by Hassenstein and Reichardt [166] to describe visual perception of insects. While the EMD and mathematically similar models do not estimate true image velocity, there is strong experimental evidence [167–169] that supports this model as a mechanism for visual motion detection in insects, humans, and other animals. The output of this algorithm R_{EMD} is given in Eq. 4.1 where C is contrast, h_1 is the magnitude response for a first-order high-pass filter, h_2 is the magnitude response for a first-order low-pass filter, ϕ_2 is the phase response of the first-order low-pass filter, and ϕ_x is the relative phase between neighboring photoreceptors.

$$R_{EMD} = -\frac{C^2}{4} \cdot h_1^2 \cdot h_2 \cdot sin\phi_2 \cdot sin\phi_x \tag{4.1}$$

The C^2 term in Eq. 4.1 indicates a strong, non-linear, dependence on image contrast while the $sin\phi_2$ makes the response an odd function and confers directionality to the algorithm. The output of the EMD is a sinusoidal wave with mean value R_{EMD}. This mean is extracted by applying a low-pass filter to the output of the detector.

In our simulations the HPF and LPF are implemented as infinite-impulse-response (IIR) single-pole low-pass filters in MATLAB through a recurrence relationship. These filters have no

explicit definition of time but rather perform temporal filtering with respect to frames. Thus our simulations are generalizable to imaging systems which may have varying frame rates.

Non-Directional Multiplication

The non-directional multiplication (NDM) detector, shown in Fig. 4.4(b), is mathematically similar to the EMD. It correlates the response of one photodetector to the delayed responses of two neighboring photodetectors. The output of this algorithm R_{NDM} is given in Eq. 4.2 where C, h_1, h_2, ϕ_2, and ϕ_x are the same parameters as in Eq. 4.1.

$$R_{NDM} = \frac{C^2}{4} \cdot h_1^2 \cdot h_2 \cdot cos\phi_2 \cdot cos\phi_x \tag{4.2}$$

Similar to the EMD, the output of the NDM algorithm is a sinusoidal signal with mean R_{NDM}. The transfer function is almost identical to the EMD with the exception of a $cos\phi_2$ term instead of $sin\phi_2$. This produces an even function instead of an odd function and removes directionality in its motion sensitivity. Benefits of NDM over the EMD algorithm include increased bandwidth with regard to spatial frequencies, smaller dependence on spatial frequency, and improved linearity with respect to speed.

Non-Directional Summation

Proposed by Rivera-Alvidrez [170], the non-directional summation (NDS) detector differs from the NDM detector only in its last stage, employing summation instead of multiplication (see Fig. 4.4(c)). The NDS model is linear and produces a zero-mean sinusoidal output. Thus unlike the NDM and EMD, it encodes velocity as an amplitude instead of a mean. The amplitude for a sinusoidal input is given by Eq. 4.3.

$$A_{NDS} = \frac{C}{2} \cdot h_1 \cdot \sqrt{4 \cdot h_2 \cdot (h_2 \cdot cos^2\omega_x + cos\omega_x + cos\phi_2) + 1} \tag{4.3}$$

In order to process the output signal in the same way as the EMD and NDM algorithms, we add an additional rectification layer after the output. This generates a DC offset that encodes velocity, similar to the other correlation-based detectors. This rectified output is given by Eq. 4.4.

$$R_{NDS} = \frac{2 \cdot A_{NDS}}{\pi} \tag{4.4}$$

Among the EMD, NDM, and NDS algorithms, the NDS offers the best linearity in its output with respect to velocity, the smallest sensitivity to spatial frequency, and the broadest bandwidth with respect to temporal frequency. In addition, it has a linear dependence on contrast in comparison with the quadratic dependence observed in the EMD and NDM algorithms. We hypothesize that the NDS algorithm will outperform the EMD and NDM algorithms when processing natural stimuli at low-light levels because it does not require multiplication. Multiplication serves the functional role to provide correlation of both signal and noise. For multiplication of two Poisson variables $\lambda_1, \lambda_2 > 1$ the SNR is >0 and multiplication improves SNR of the signal product. However, if $\lambda_1, \lambda_2 < 1$ then multiplication of correlated signals kills signal strength faster than noise. Addition, on the other hand, always results in an output SNR that is greater than that of the input signals. We expect that for photon-limited imaging, this difference will lead to improved low-light performance of NDS over the other algorithms.

4.3.4 Error Rate

In order to quantify the performance of each algorithm it is important to take into consideration the large differences that exist between the three classes of algorithms. One issue is the difference in baseline accuracy (how well algorithms perform in optimal conditions) of BM and IIA. Block-matching cannot detect sub-pixel motion and can only evaluate integral velocities while IIA handles large motion poorly and can evaluate non-integral values. Furthermore, the EMD,

NDM, and NDS algorithms all produce scalar values, not velocities, and are very sensitive to contrast and spatial frequency of the scene. Thus, defining baseline accuracy when processing a natural scene is impossible for these algorithms. For this reason we cannot directly compare baseline accuracy of each algorithm. Furthermore, rather than baseline performance we are primarily interested in quantifying the effects of shot noise on each optic flow algorithm. For this reason we propose the metric "error rate" to quantify low-light performance.

We define error rate as the difference between the stochastic output and non-stochastic output of a particular algorithm. This lets us eliminate differences in baseline accuracy and look solely at the effects of shot noise. This metric essentially suppresses any bias or error in baseline performance for a specific optic flow computation and isolates the effects of the noise (shot noise and DCR). We define R_{ideal} as the ideal response and R_{stoch} as the stochastic response of a particular algorithm. For a frame with n rows and m cols ($n = m = N_{pool}$) we compute the average error rate, E_{avg}, across a frame as shown in Eq. 4.5.

$$E_{avg} = \frac{1}{(n-2)(m-2)} \sum_{i=2}^{n-1} \sum_{j=2}^{m-1} \left| \frac{R_{ideal}(i,j) - R_{stoch}(i,j)}{R_{ideal}(i,j)} \right| \tag{4.5}$$

We express the average error rate E_{avg} corresponding to an optic flow measurement with n-bit precision. The accuracy in an n-bit measurement is 2^{-n} and the SNR of the output is $20 \cdot log_{10}(2^n)$. Thus 2-bit resolution corresponds to an error rate of 25% (12dB) and 3-bit resolution corresponds to an error rate of 12.5% (18dB). In subsequent analysis of each algorithm we define N_{2b}, N_{3b}, and N_{4b} as the photon flux at which the average optic flow measurement across a frame achieves 2-bit, 3-bit, and 4-bit resolution, respectively.

4.4 Results: Synthetic Images

Benchmarking results are provided for each algorithm when processing 2D sine gratings. Simulations have been performed, as described above, to quantify the effects of spatial frequency, image contrast, dark count rate, image velocity, and spatial pooling, and we present the results of each below. We show the results of each benchmarking experiment by plotting N_{2b} as a function of the parameter under investigation. We plot N_{2b} for each algorithm: EMD, NDM, NDS, BM, and IIA (see legends). The legend entry 'IM' corresponds to an intensity measurement. It is computed as the average error rate, E_{avg}, when the ideal response, R_{ideal}, and the stochastic response, R_{stoch}, correspond to pixel-level measurements of photon flux (number of photons detected per frame) rather than an optic flow estimate. We include this in the comparisons as it gives a sense of pixel-level SNR and enables a comparison between traditional imaging (intensity measurement) versus optic flow imaging (motion estimate) at low-light levels. As expected, we note that N_{2b} does not change as we sweep parameters such as spatial frequency or image contrast because those parameters have no effect on pixel-level intensity measurements.

4.4.1 Spatial Frequency

The effects of spatial frequency on optic flow estimation were investigated. The simulations used sine gratings with a fixed DCR of 0.1 counts/pixel/frame, Michelson contrast of 90% and image velocity of 5 pixels/frame. This corresponds to a full 360^o rotation every 50.2 frames. We simulate error rates for spatial frequencies ranging from 1 cycle/frame to 9 cycles/frame and measure N_{2b} for N_{pool} = 60,120,180,240. In Fig. 4.5 we plot N_{2b} versus spatial frequency for each algorithm.

Simulation results show a general trend that as spatial frequency increases N_{2b} decreases.

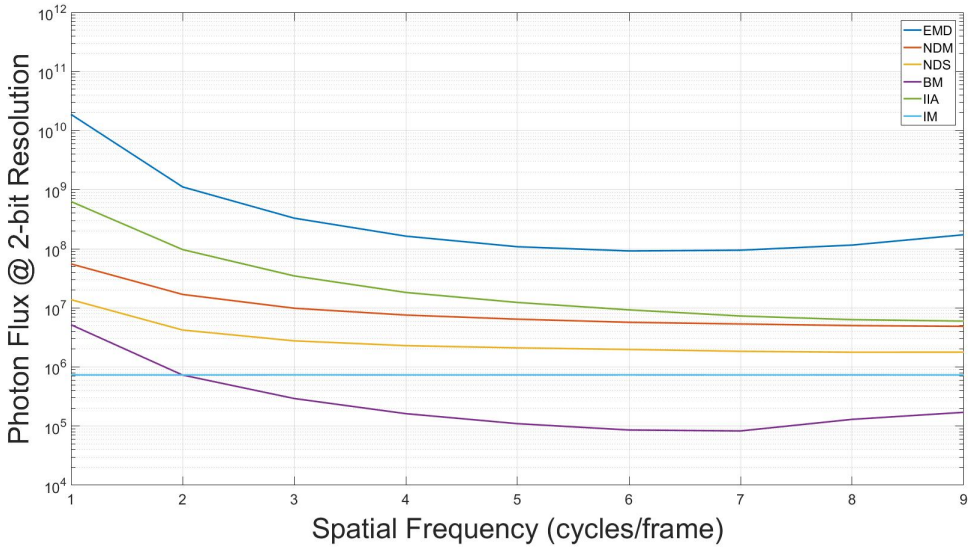

Figure 4.5: Plot of N_{2b} versus spatial frequency for synthetic images with C=90%, DCR=0.1, and image velocity of 5 pixels/frame. Spatial frequency is swept from 1 cycle/frame to 9 cycles/frame. We show results for $N_{pool} = 240$.

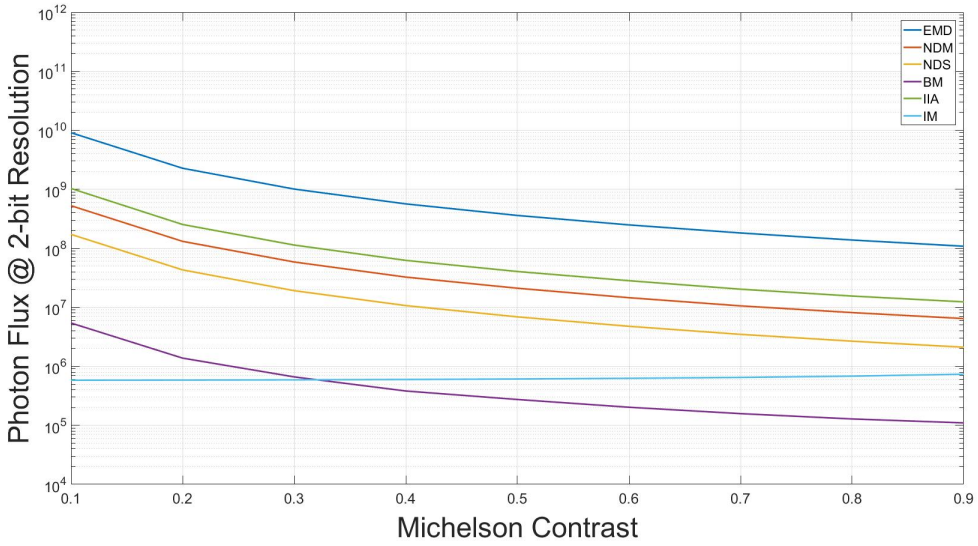

Figure 4.6: Plot of N_{2b} versus image contrast for synthetic images with f=5 cycles/frame, velocity = 5 pixels/frame, and DCR=0.1 counts/pixel/frame. Image contrast is swept from 10% to 90% in increments of 10%. We show results for $N_{pool} = 240$.

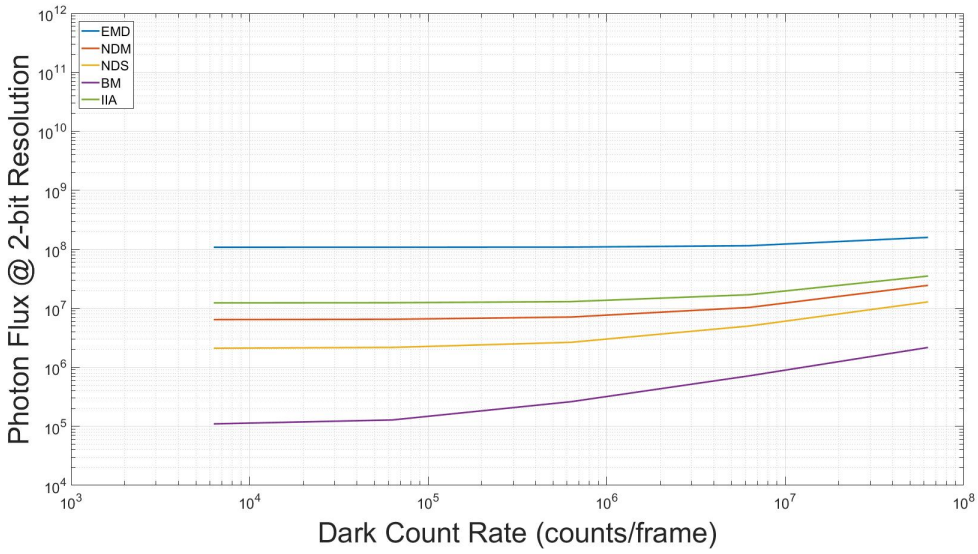

Figure 4.7: Plot of N_{2b} versus dark count rate for synthetic images with f=5 cycles/frame, velocity = 5 pixels/frame, and C=90%. DCR is swept from 0.1 to 1k counts/pixel/frame with one point per decade. We show results for $N_{pool} = 240$.

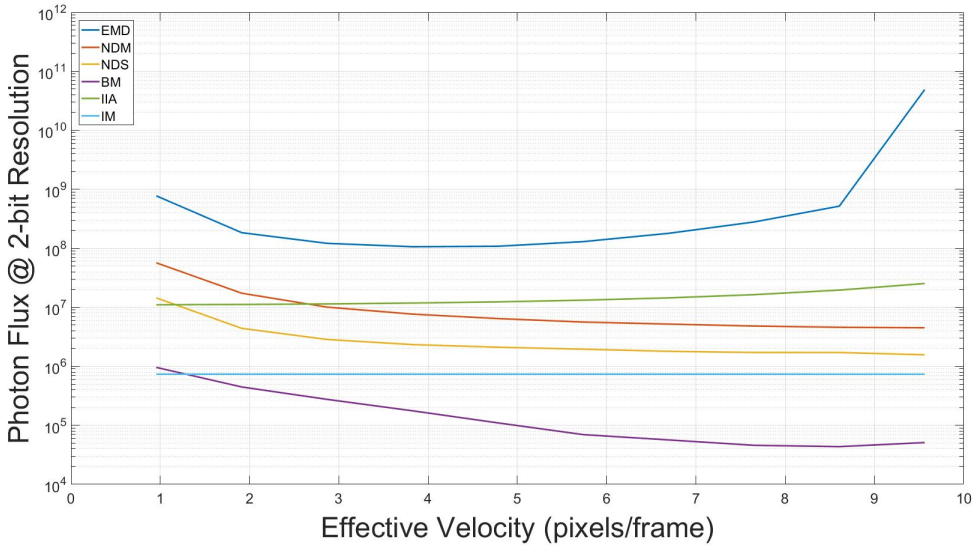

Figure 4.8: Plot of N_{2b} versus image velocity for a synthetic image with f=5, C=90%, DCR=0.1. Image velocity is swept from 1 pixel/frame to 10 pixels/frame. We show results for $N_{pool} = 240$.

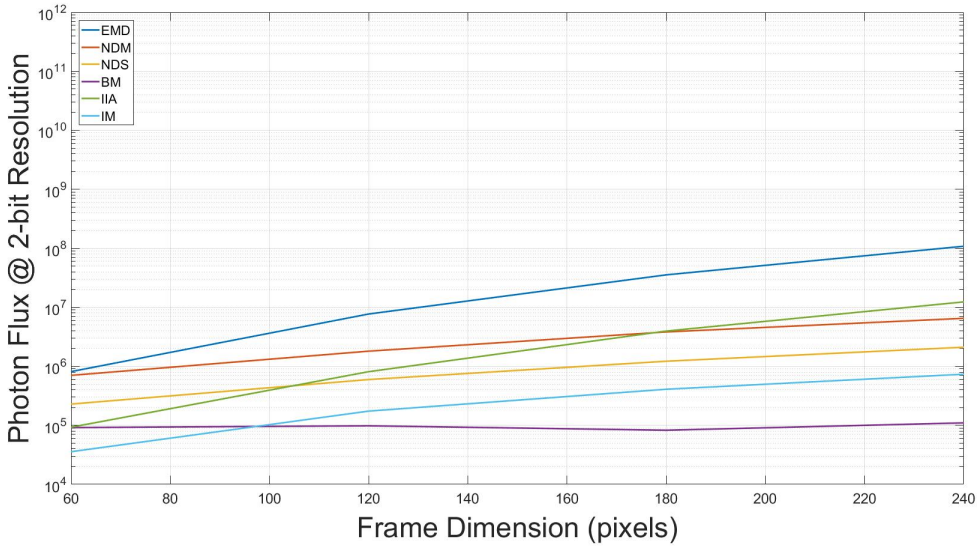

Figure 4.9: Plot of N_{2b} versus N_{pool}. N_{pool} is swept from 60 to 240 in steps of 60. DCR is 0.1 counts/pixel/frame, image velocity is 5 pixels/frame, and spatial frequency is 5 cycles/frame.

The EMD has the highest sensitivity to spatial frequency and experiences a shift in N_{2b} of about 3 orders of magnitude while the NDM and NDS experience a shift of <1 order of magnitude. At very high spatial frequencies (cycles/pixel) we note that N_{2b} begins to increase. This increase corresponds to saturation in the transfer functions of the correlation-based detectors. This is most notable for $N_{pool} = 60$ (as it has the highest spatial frequencies in units of cycles/pixel). The EMD transfer function becomes fully saturated at 7 cycles/frame which is when we see a sharp rise in N_{2b}. As previously noted, the NDM and NDS algorithms offer wider bandwidth than the EMD. This broadband response improves algorithm performance at high spatial frequencies and prevents N_{2b} from decreasing as quickly as the EMD at high spatial frequencies. The IIA algorithm experiences a decrease in N_{2b} of about 2 orders of magnitude as spatial frequency increases. Intuitively this makes sense as an increase in spatial frequency increases both the spatial and temporal gradients of the image. Since gradient-based algorithms extract motion

from gradients and not absolute pixel values, it makes sense that as these gradients becomes larger the effective SNR of these algorithms increase as well. The block matching algorithm experiences a similar shift downward in N_{2b} as we move from low frequencies to mid-range frequencies. Intuitively this makes sense as an increase in spatial frequency corresponds to an increase in spatial contrast within sub-blocks and more salient features to match. At high spatial frequencies, motion direction becomes ambiguous and the block matching algorithm becomes more susceptible to shot noise. This ambiguity is a feature of periodic sine gratings, as we see in following sections.

4.4.2 Contrast Effects

The effects of contrast on optic flow estimation were isolated by fixing the spatial frequency of the sine grating to 5 cycles/frame, the DCR to 0.1 counts/pixel/frame, and the image velocity to 5 pixels/frame as the contrast varied. Fig. 4.6 shows simulation results. We plot N_{2b} versus contrast for each algorithm as Michelson contrast is swept from 10% to 90%. Again, we repeat this measurement for N_{pool} = 60,120,180,240.

We see that as contrast increases, N_{2b} decreases monotonically for all of the algorithms. This result is fairly intuitive. Higher contrast images effectively produce more textured images. They produce a stronger response from the correlation-based detectors, and produce more salient features (bright pixels with higher SNR) to track or match. A surprising result is that the rate of change in N_{2b} with respect to contrast ($S_C = -\frac{\partial log_{10}(N_{2b})}{\partial C}$) is uniform across all algorithms regardless of spatial pooling. This result suggests that improvements are not a result of sensitivity of each algorithm to contrast but rather to improved SNR in image "texture" (being able to differentiate peaks and troughs). This is particularly surprising result because the responses of the EMD and NDM algorithms have a square dependence on contrast while the NDS response

has a linear dependence. We see that at low contrast (C=10-20%) S_C is approximately 0.5, while at high contrast S_C is around 0.1.

4.4.3 Dark Count Effects

Fig. 4.7 shows N_{2b} versus DCR as we sweep DCR from 0.1 to 1k counts/pixel/frame with one point per decade. The 2D sine grating has a fixed image frequency of 5 cycles/frame, image velocity of 5 pixels/frame, and contrast of 90%. Dark count rate has the effect of reducing contrast in a scene while also contributing to the shot noise in each pixel. Thus, it sets the noise floor for each pixel. A DCR of 1 counts/pixel/frame corresponds to a total of 63k counts/frame.

As expected, N_{2b} increases as DCR increases. At low DCR levels (0.1 and 1 counts/pixel/frame) N_{2b} is more than an order of magnitude larger than the total number of dark events per frame. Thus we see a negligible effect of DCR on N_{2b}. At higher DCR levels (>10 counts/pixel/frame) we find that N_{2b} increases at a rate of approximately $\frac{\partial log_{10}(N_{2b})}{\partial log_{10}(DCR)} \approx 0.5$. This is only observed if the DCR is close to or greater than N_{2b}. This result suggests that optic flow can be detected when noise events outnumber detection events, at least in high contrast images. We note that even for a sine grating with 6M photons/frame and 63M dark counts/frame we can visually inspect sequential frames (video) and detect motion in spite of the high noise levels.

4.4.4 Velocity

We isolate the effects of image velocity by fixing spatial frequency at 5 cycles/frame and sweeping image velocity from 1 to 10 pixels/frame. Again, contrast is set to 90% and DCR is set to 0.1 counts/pixel/frame. In Fig. 4.8 we plot N_{2b} versus velocity for each algorithm.

For sine gratings, an increase in image velocity has a similar effect to an increase in spatial frequency. Both have the effect of increasing optic flow and apparent image motion. Thus similar

to the frequency sweep, as we increase velocity we see a general downward trend of N_{2b}. This trend holds until the correlation-based detectors become saturated. For the narrow-band EMD we see this effect at velocities of around this occurs at a lower velocities of around 5 pixels/frame while the broad-band NDM and NDS do not saturate until velocities in excess of 12 pixels/frame for a spatial frequency of 5 cycles/frame. Similarly, the block matching algorithm displays a downward trend in N_{2b} as velocity increases. It is important to note that the block matching algorithm cannot detect sub-pixel motion or motion in excess of its search radius. Thus for $N_{pool} = 60$ and effective image velocities below 0.7 pixels/frame the BM algorithm outputs zero motion and we cannot compute error rates ($E_{avg} = \frac{0}{0}$) or extract N_{2b}. At large image velocities (>7 pixels/frame) we note a slight increase in N_{2b}. Uniquely for the BM algorithm, we note that certain pairs of velocity and spatial frequency generate ambiguous motion direction–this explains the spikes in N_{2b}. This occurs because of the periodic nature of sine gratings and is only observed in BM because each sub-block spans a large horizontal distance of $N_{pool}/10 + 2p$ where p is the search radius. For $N_{pool} = 60$ this corresponds to a full period of the sine grating, thus motion of approximately 0.5 cycles/frame, or 5-7 pixels/frame, is ambiguous. In contrast to BM, the IIA algorithm is designed for sub-pixel motion and displays poor baseline accuracy (high error rates in R_{ideal} compared to known optic flow). Thus N_{2b} generally increases for the IIA as velocity increases.

It is worth noting that we ignore image blurring for sine gratings because blurring effectively adds another layer of spatial pooling. This has minimal effect on sine gratings, as there is no high frequency content being filtered out, and has a more significant effect in natural scenes as we discuss later.

4.4.5 Spatial Pooling

In Fig. 4.9 we plot N_{2b} as we sweep N_{pool} from 60 to 240. We set the DCR to 0.1 counts/pixel/frame, image velocity to 5 pixels/frame, spatial frequency to 5 cycles/frame, and contrast to 90%.

We see that as spatial pooling increases, N_{2b} decreases for every algorithm except for BM. Spatial pooling has the effect of increasing signal strength at each pixel (expected to decrease N_{2b}). However, spatial pooling also has the effect of decreasing effective image velocity. We previously showed that an increase in velocity generally increases N_{2b} for the correlation-based detectors and decreases N_{2b} for the IIA. Thus, a decrease in N_{2b} as spatial pooling increases (N_{pool} decreases) makes sense. To explain the results for the correlation-based detectors we must note that an increase in spatial pooling also has the effect of increasing the spatial frequency in units of cycles/pixel (not cycles/frame). Thus for the correlation-based detectors, a decrease in effective image velocity is countered by an increase in spatial frequency (cycles/pixel) and cumulatively have no effect. Therefore, the primary effect of spatial pooling in the correlation-based detectors is an increase in signal strength which supports our finding that N_{2b} decreases with spatial pooling. BM is relatively unaffected by spatial pooling. One reason for this is the fact that block size scales with spatial pooling. Thus each sub-block has the same photon flux (photons/sub-block/frame) regardless of N_{pool}. This means that the SNR of each sub-block is unaffected by spatial pooling. The fact that N_{2b} slightly increases as spatial pooling decreases likely corresponds to the decrease in effective image velocity. We note that even for $N_{pool} = 60$, the effective image velocity is still >1 pixel/frame and BM can adequately estimate optic flow.

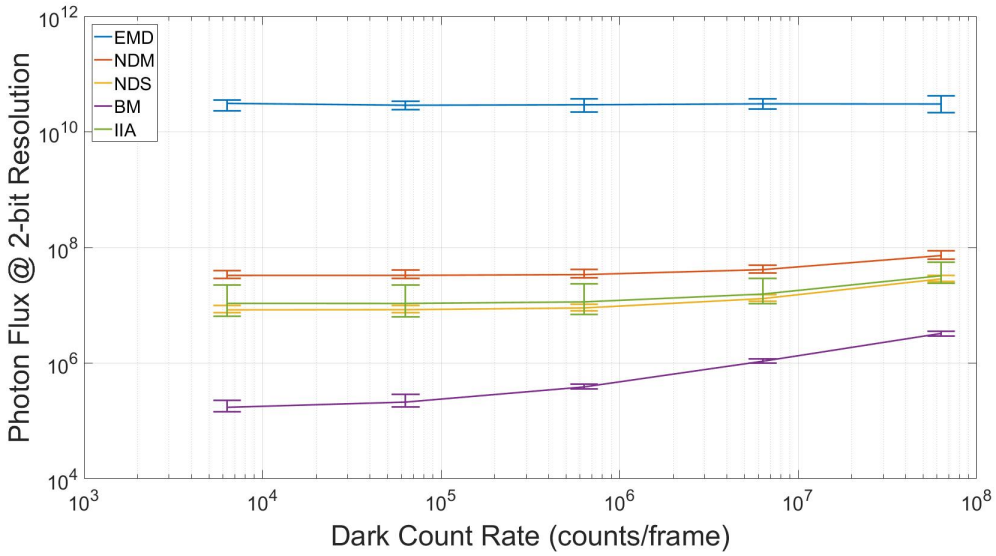

Figure 4.10: Simulation results showing N_{2b} averaged over each of the four natural panoramas with error bars indicating maximum and minimum values of N_{2b} among the panoramas. DCR is swept from .1 to 1k counts/pixel/frame with an image velocity of 5 pixels/frame and $N_{pool} = 240$.

4.5 Results: Natural Images

Unlike synthetic images, we are unable to control parameters such as spatial frequency and contrast in natural images. Instead we benchmark the performance of each optic flow algorithm for four different panoramas. Each panorama has slightly different image properties (image contrast and frequency content). Again, we look at the effects of dark count rate, image velocity, and spatial pooling. Additionally, we explore the effects of image blurring. Since we are estimating N_{2b} for four different panoramas, we plot the mean value of N_{2b} (averaged across the four panoramas) in each figure and indicate the maximum and minimum value of N_{2b} with error bars.

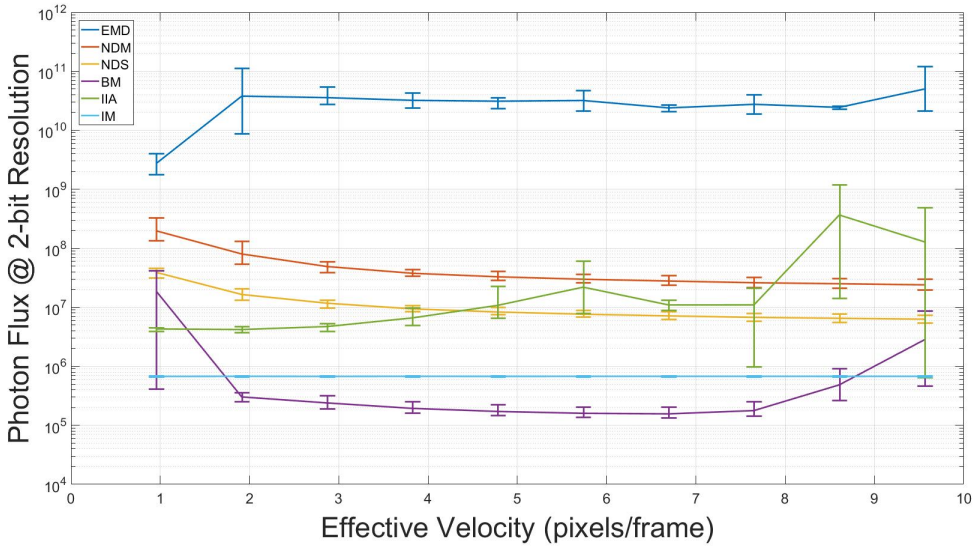

Figure 4.11: Simulation results showing N_{2b} averaged over each of the four natural panoramas with error bars indicating maximum and minimum values of N_{2b} among the panoramas. Velocity is swept from 1 to 10 pixels/frame with DCR=0.1 counts/frame and $N_{pool} = 240$.

4.5.1 Dark Count Effects

We swept DCR from 0.1 to 1k counts/pixel/frame with one point per decade. This corresponds to a total of 6.3k to 63M dark counts/frame. We set image velocity to 5 pixels/frame and compute N_{2b} for each of the four panoramas. In Fig. 4.10 we plot N_{2b} on the y-axis and DCR (dark counts/frame) on the x-axis.

The results are similar to what we found for the synthetic images. When DCR is smaller than N_{2b} we see little effect on N_{2b}. As DCR increase to levels comparable and larger than N_{2b}, N_{2b} increases. A notable difference is that the mean value of N_{2b} is higher in natural images than in synthetic images. Thus some algorithms, such as the EMD, experience no noticeable difference in N_{2b}. This is because the DCR remains orders of magnitude lower than N_{2b} even at high levels of DCR.

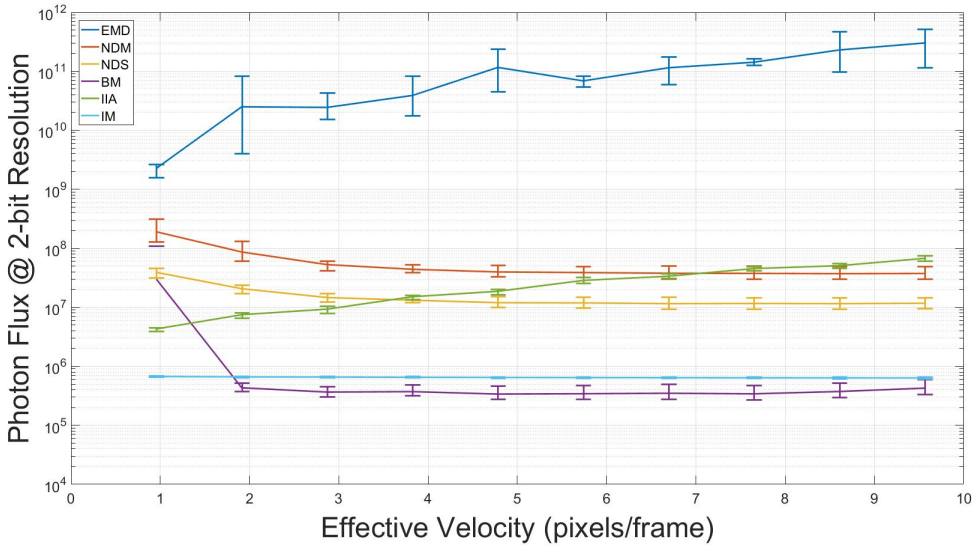

Figure 4.12: Simulation results showing N_{2b} averaged over each of the four natural panoramas with error bars indicating maximum and minimum values of N_{2b} among the panoramas. Velocity is swept from 1 to 10 pixels/frame with DCR=0.1 counts/frame and $N_{pool} = 240$. Image blurring is simulated with a spatial filter in MATLAB.

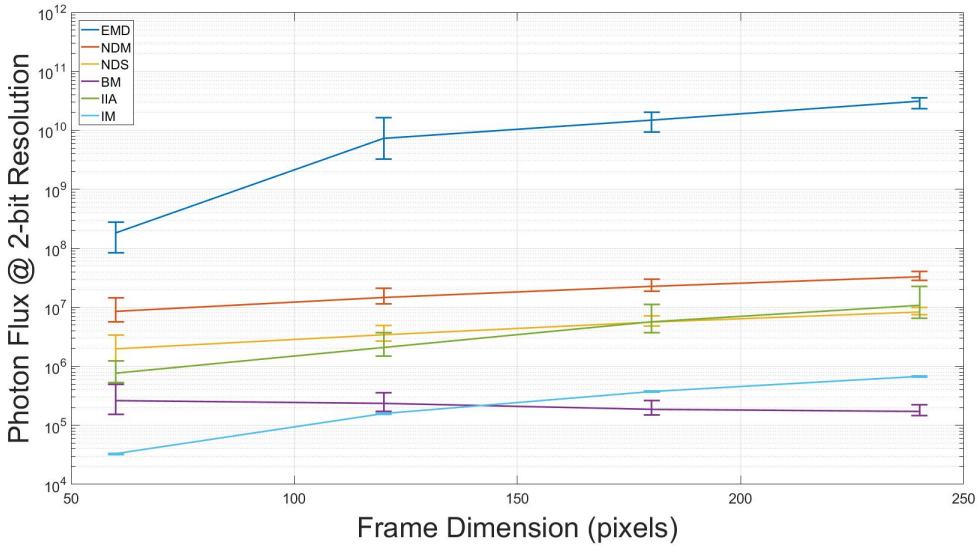

Figure 4.13: Simulation results showing N_{2b} averaged over each of the four natural panoramas with error bars indicating maximum and minimum values of N_{2b} among the panoramas. N_{pool} is swept from 60 to 240 in steps of 60 with an image velocity of 5 pixels/frame and DCR=0.1 counts/pixel/frame.

4.5.2 Velocity

We repeated the same velocity sweep for natural scenes. We did this with a DCR of 0.1 counts/pixel/frame and N_{pool} = 60 and 240. We swept over image velocities from 1 to 10 pixels/frame. We did this for all four panoramas and for each algorithm. In Fig. 4.11 we plot the average value of N_{2b} across each of the panoramas on the y-axis and velocity on the x-axis. Again the error bars indicate maximum and minimum values of N_{2b} among the panoramas.

Results show similar trends to what we see in synthetic images. Generally, as effective velocity increases, N_{2b} decreases for the NDM and NDS detectors. In contrast, N_{2b} quickly increases for the EMD and remains fairly high. The reason for this is that the narrow bandwidth of the EMD causes it to saturate at a very low effective velocity of 1-2 pixels/frame while the NDS and NDM saturate at velocities in excess of 10 pixels/frame. Thus once the EMD saturates, N_{2b} remains high. These finding agree with our findings with synthetic images and accentuates both the high sensitivity of the EMD to spatial frequency and the benefits of using a detectors with a wider bandwidth.

Results for the IIA also mimic those found using synthetic images. As velocity increases, N_{2b} also increases. At high image velocities, baseline accuracy of the IIA degrades substantially and produces high variance in N_{2b}. This is primarily because it is operating beyonds its optimal operating point. This is nearly identical behavior to what we find in the EMD. The IIA essentially saturates and cannot compute optic flow with high accuracy making it more susceptible to noise. N_{2b} decreases for BM until the velocity exceeds its search radius parameter. At that point we see N_{2b} increase fairly quickly. Again, the search radius essentially sets the velocity at which the BM algorithm saturates and becomes more susceptible to noise.

4.5.3 Motion Blurring

Motion blurring is a particularly important effect in low-light imaging. In such environments temporal pooling is a common technique used to increase the optical signal strength and boost SNR. Blurring removes high frequency image features (both spatial and temporal), decreases image contrast, and can be viewed as an additional layer of spatial pooling. To explore the effects of image blurring we replicate the experiment conducted in Sec. 4.5.2 and incorporate a spatial filter during image generation to simulate the effects of blurring. This spatial filter assumes linear motion of the camera and is implemented using the MATLAB function H = fspecial('motion',ν,0). We plot the result of this simulation in Fig. 4.12.

Blurring has the effect of causing the EMD, NDM, and NDS to saturate at slightly lower velocities. It also has the effect of decreasing contrast between neighboring pixels. As expected, at low velocities blurring (<3 pixels/frame) has minimal effect on N_{2b}. At higher velocities we see that blurring has a negative effect on all of the correlation-based detectors. This effect is fairly small for the NDM and NDS, as blurring causes them to just begin to saturate around a velocity of 7 pixels/frame, and is much larger for the EMD which is further pushed into saturation. The performance of BM and IIA are actually improved by blurring at high velocities. This is because blurring improves the baseline accuracy of both algorithms when they begin to saturate. This result is not intuitive but leads to an improvement in N_{2b} over non-blurring simulation results.

4.5.4 Spatial Pooling

In Fig. 4.13 we plot N_{2b} as we sweep N_{pool} from 60 to 240. We set the DCR to 0.1 counts/pixel/frame and image velocity to 5 pixels/frame. The results are nearly identical to those from synthetic images. One notable difference is that N_{2b} is generally larger for natural scenes than for syn-

thetic scenes. As with synthetic images we see that all of the algorithms benefit from spatial pooling except for BM. In fact the degradation in performance of BM is even more substantial for natural scenes.

4.6 Discussion

4.6.1 Effects of Baseline Accuracy

A common trend we see in our benchmarking results is that as baseline performance decreases, N_{2b} generally increases and optic flow algorithms become more susceptible to noise. As previously stated, we can only analyze BM and IIA in this manner because the correlation-based detectors have no baseline accuracy. This trend makes sense with block matching because its baseline accuracy falls at low velocities (sub-pixel motion) and high velocities (greater than search parameter). At these velocities, the algorithm fails to find a good match between consecutive frames. The output signal of BM is generated by the cross-correlation function between sub-blocks. Thus when the algorithm fails to find good matches between consecutive frames, the cross-correlation decreases and the output signal strength and SNR drop. Intuitively this can be thought of as the algorithm having low confidence in its motion estimate and thus becoming more susceptible to noise. Baseline accuracy drops for the IIA only at high velocities. In this operating regime the algorithm fails to compute spatial derivatives accurately (or undersamples gradients). This is important because the IIA operates on spatial and temporal gradients, not raw pixel values. As the accuracy of spatial derivative estimates decreases so does the effective input signal strength of the algorithm. This reduces overall SNR for the algorithm and increases sensitivity to noise.

It is important to note that modifications can be made to these algorithms to improve

baseline accuracy where they struggle. However, generally these algorithms are used in different operating regimes (small displacement vs. large displacement).

The overall conclusion from this analysis is that low-light performance of these optic flow algorithms is closely linked to their operating conditions. Thus to maximize low-light performance, it is important to maximize baseline accuracy. While this result is fairly obvious, it is an important finding.

4.6.2 Correlation-Based Detectors

Upon inspection of our benchmarking results it quickly becomes obvious that performance of the NDM and NDS algorithms are quite similar while the EMD is substantially worse across all benchmarking experiments. This result shows three things. First, scaling of output magnitude by contrast has negligible effect. We would expect that since EMD and NDM have C^2 dependence in their optic flow response they would be more sensitive to contrast than NDS which has linear dependence. Second, mathematically the EMD and NDM are nearly identical. The major difference between the two is kernel size. The EMD takes inputs from only two pixels while the NDM takes inputs from three. The larger kernel has a higher SNR than the smaller kernel and would be expected to improve low-light performance. Third, in spite of the differences in algorithm design the non-directional algorithms consistently outperform the directional algorithm. From an SNR point-of-view this makes sense. Summation universally improves SNR, since signals add linearly and Poisson noise adds sub-linearly, at its output while subtraction universally decreases SNR at its output. This conclusion is supported by theoretical and experimental studies of nocturnal insects [171–173] that discuss various forms of neural and optical summation throughout their visual systems—subtraction is rarely observed in the visual pathways of nocturnal insects. Achieving directionality in a correlation-based detector

inevitably requires subtraction at some stage to differentiate between motion to the left and motion to the right. As noise is introduced into the visual system, the ability to differentiate between those directions decreases and subtraction of the "left" and "right" signals can all but eliminate signal strength at the output of the EMD and diminishes SNR.

Another conclusion is that the transfer functions of the algorithms are closely related to low-light performance. The EMD has the narrowest tuning curves (with respect to velocity) of the three correlation-based detectors. The EMD's output quickly falls to zero once aliasing and saturation occur. In contrast, the NDS and NDM algorithms display much wider tuning curves and higher bandwidth. Even after their outputs saturate they remain high, near their peak value, for much higher velocities. Practically, this means that the NDS and NDM fail to differentiate large image velocities but are very sensitive to a much wider range of high image velocities than the EMD. This behavior is evident in our benchmarking results as they offer "peak" low-light performance over a wider range of velocities and operating conditions. We can draw parallels with baseline accuracy of BM and IIA. The EMD has a fairly narrow set of operating conditions over which it can reliably detect optic flow while the NDM and NDS have a broader set of operating conditions. This makes the NDM and NDS algorithms more resilient to fluctuations in spatial frequency and image velocity than the EMD. This translates to a wider range of image velocities over which the NDM and NDS can reliably detect optic flow at low-light levels.

An important distinction between the correlation-based detectors, BM, and the IIA is that the correlation-based detectors are the only detector that provide dense flow fields with one vector per pixel. BM and IIA generate sparser fields and, in this instance, only provide 100 vectors per frame. While they can be programmed to generate higher density fields they generally

cannot generate flow for each pixel, or can only do so with a substantial increase in computing time and power.

4.6.3 Block Matching

A major difference between the BM algorithm and the correlation-based detectors is kernel size. Unlike the correlation-based detectors that estimate optic flow for each pixel, BM estimates optic flow for sub-blocks. Thus it aggregates photons over an area of the image in order to estimate motion. The BM algorithm relies upon a Mean Absolute Difference (MAD) cost function to estimate maximum "similarity" or correlation between sub-blocks. Maximum correlation between sub-blocks corresponds to minimum MAD between sub-blocks. This means a single "bright" pixel or high contrast region in a sub-block can provide enough SNR to estimate motion in a noisy scene. The large kernel size of this algorithm lets it "search" a sub-block for the best pixel to match with. Intuitively it makes sense that this algorithm provides the best overall low-light performance. As long as one pixel in a sub-block overcomes some noise threshold, it is able to accurately estimate motion. At low-light levels this improves performance with respect to algorithms with smaller kernel sizes and contributes to it outperforming the other algorithms in photon-limited environments. Due to its large kernel, BM is able to achieve 2-bit accuracy in optic flow measurements at lower light levels than traditional intensity measurements can achieve 2-bit resolution. The downside of these large kernels is that BM benefits from high-density pixel arrays that are more costly to implement in terms of power consumption, complexity, and computing cost.

4.6.4 Computational Considerations

As a first-order approximation of computing costs, we measured the average computing time of each algorithm for a single frame. Computations and timing measurements were performed in MATLAB on an Intel Core i5-5300U CPU @ 2.3GHz. For a 240x240 pixel frame, the average computing times were 338ms for BM, 9ms for IIA, and <3ms for the correlation-based detectors. For a 60x60 pixel frame, the average times drop to 17ms for BM, 6ms for IIA, and <.3ms for the correlation-based detectors. It is worth noting that MATLAB is optimized for fast matrix operations over loop-based operations which introduces some bias into these measurements. Regardless, the discrepancy amongst algorithms is quite large. BM has, by far, the highest computational load of the algorithms. The EMD, NDM, and NDS algorithms are computationally simple and can be implemented in analog circuitry fairly efficiently [174]. Furthermore, the BM algorithm offers peak performance when low-levels of spatial pooling are employed which further increases disparities in computing times—the difference in processing a 60x60 frame with NDS and a 240x240 frame with BM is around 3 orders of magnitude.

4.6.5 Conclusions and Results

In Table 4.1 we summarize peak performance of each algorithm in both natural and synthetic scenes. This table provides the minimum value of N_{2b}, in units of photons/frame, of all benchmarking experiments. A notable result is that BM is able to measure optic flow with 2-bit resolution at a photon flux of 156k photons/frame. This corresponds to an average of just 2.7 photons/pixel/frame for a 240x240 pixel image. This is a remarkable result and roughly translates to an illuminance of a few hundred μlux when operating at a frame-rate of 30fps. This result is promising for low-light vision systems and suggests that there is room for improve-

ment in current night vision systems. While this work does not address the effects of mismatch and non-uniformities in image sensors, these issues are addressed in current literature and can generally be minimized through post-processing and image corrections [175–177]. Photon loss from quantum efficiency and optics can be accounted for based on the technology of the image sensor and can be used to scale our results as appropriate. Even a 10x increase in N_{2b} would still represent a notable improvement over current state-of-the-art optic flow sensors.

Table 4.1: Peak Photon-Limited Optic Flow Performance Summary

	Synthetic	Natural
BM	83k	156k
IIA	68k	739k
EMD	813k	154M
NDM	681k	4.38M
NDS	228k	1.02M

In order to extrapolate our findings to higher resolution optic flow imaging, we note that for the EMD, NDM, NDS, IIA increasing photon flux by a factor of 10 roughly improves SNR at the output of the optic flow detectors by 10dB. This result holds regardless of spatial pooling but requires that algorithms are operating in "optimal conditions" (sub-pixel motion for IIA and not in saturation). This relationship holds once these algorithms exceed 2-bit resolution. This slope of 10dB/decade matches the slope of SNR for a Poisson process $SNR = 10 \cdot log_{10} N$, where N is the number of Poisson events. This shows that once each optic flow algorithm overcomes some noise threshold, they all improve as a typical Poisson process. Before they exceed this threshold, improvements in SNR are non-linear for all of the algorithms except NDS. In this work we are quantifying what this noise threshold is for each algorithm. BM does not follow this trend.

SNR of this algorithm increases at a much faster rate, closer to 30dB/decade and, unlike the other algorithms, cannot be extended to very high photon rates. This is because BM is the only algorithm that strictly outputs integer values. At some point error rates drop to exactly 0%, or an infinite SNR, because optic flow estimates between stochastic and non-stochastic images are approximated to precisely the same values—this never occurs in the other algorithms.

In conclusion, there are many trade-offs that need to be considered when developing low-light optic flow sensors. These considerations include computational complexity, whether large displacement or small displacement optic flow needs to be computed, density requirements of optic flow estimates, high-level image processing requirements of low-level optic flow estimates, and more. Each of the five algorithms present unique characteristics in low-light environments. The goal of this work is not to optimize a particular algorithm or to suggest an optimal approach for low-light optic flow imaging. Instead we benchmark the performance of these five algorithms at low-light levels, explore the effects of visual stimuli on each algorithm, and quantify differences between each algorithm.

Bibliography

[1] L. Petricca, P. Ohlckers, and C. Grinde, "Micro- and nano-air vehicles: State of the art," *International Journal of Aerospace Engineering*, vol. 2011, pp. 1–17, 2011.

[2] R. R. Harrison, "An analog VLSI motion sensor based on the fly visual system," Ph.D. dissertation, California Institute of Technology, 2000.

[3] E. Warrant and M. Dacke, "Visual orientation and navigation in nocturnal arthropods," *Brain Behav. Evol.*, vol. 75, pp. 156–173, 2010.

[4] S. Johnsen, A. Kelber, E. Warrant, A. M. Sweeney, E. A. Widder, J. Raymond L. Lee, and J. Hernández-Andrés, "Visual control of flight speed in honeybees," *Journal of Experimental Biology*, vol. 209, no. 5, pp. 789–800, 2006.

[5] R. Frederiksen, W. T. Wcislo, and E. J. Warrant, "Visual reliability and information rate in the retina of a nocturnal bee," *Current Biology*, vol. 18, no. 5, pp. 349–353, 2008.

[6] E. Baird, M. V. Srinivasan, S. Zhang, and A. Cowling, "Visual control of flight speed in honeybees," *Journal of Experimental Biology*, vol. 208, no. 20, pp. 3895–3905, 2005.

[7] E. Baird, E. Kreiss, W. Wcislo, E. Warrant, and M. Dacke, "Nocturnal insects use optic flow for flight control," *Biology letters*, pp. 499–501, 2011.

[8] M. Dacke, M. Byrne, E. Baird, C. Scholtz, and E. Warrant, "How dim is dim? precision of the celestial compass in moonlight and sunlight," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 366, no. 1565, pp. 697–702, 2011.

[9] E. J. Warrant, "Seeing better at night: life style, eye design and the optimum strategy of spatial and temporal summation," *Vision research*, vol. 39, no. 9, pp. 1611–1630, 1999.

[10] E. Warrant, "Vision in the dimmest habitats on earth," *Journal of Comparative Physiology A*, vol. 190, no. 10, pp. 765–789, 2004.

[11] E. Baird, M. V. Srinivasan, S. Zhang, R. Lamont, and A. Cowling, "Visual control of flight speed and height in the honeybee," in *From Animals to Animats 9*. Springer, 2006, pp. 40–51.

[12] E. Warrant and M. Dacke, "Visual navigation in nocturnal insects," *Physiology*, vol. 31, no. 3, pp. 182–192, 2016.

[13] A. L. Stöckl, D. C. OCarroll, and E. J. Warrant, "Neural summation in the hawkmoth visual system extends the limits of vision in dim light," *Current Biology*, vol. 26, no. 6, pp. 821–826, 2016.

[14] S. M. Tierney, M. Friedrich, W. F. Humphreys, T. M. Jones, E. J. Warrant, and W. T. Wcislo, "Consequences of evolutionary transitions in changing photic environments," *Austral Entomology*, 2017.

[15] E. J. Warrant, "Seeing in the dark: vision and visual behaviour in nocturnal bees and wasps," *J. Exp. Biol.*, vol. 211, pp. 1737–1746, 2008.

[16] B. Greiner, "Adaptations for nocturnal vision in insect apposition eyes," ser. International Review of Cytology, K. W. Jeon, Ed. Academic Press, 2006, vol. 250, pp. 1 – 46. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0074769606500014

[17] J. van Hateren, "Three modes of spatiotemporal preprocessing by eyes," *Journal of Comparative Physiology A*, vol. 172, no. 5, pp. 583–591, 1993. [Online]. Available: http://dx.doi.org/10.1007/BF00213681

[18] J. C. Theobald, B. Greiner, W. T. Wcislo, and E. J. Warrant, "Visual summation in night-flying sweat bees: a theoretical study," *Vision research*, vol. 46, no. 14, pp. 2298–2309, 2006.

[19] A. Narendra, A. Alkaladi, C. A. Raderschall, S. K. A. Robson, and W. A. Ribi, "Compound eye adaptations for diurnal and nocturnal lifestyle in the intertidal ant, *Polyrhachis sokolova*," *PLoS ONE*, vol. 8, no. 10, 10 2013. [Online]. Available: http://dx.doi.org/10.13712Fjournal.pone.0076015

[20] E. Warrant, "Visual tracking in the dead of night," *Science*, vol. 348, no. 6240, pp. 1212–1213, 2015. [Online]. Available: http://www.sciencemag.org/content/348/6240/1212.short

[21] S. Sponberg, J. P. Dyhr, R. W. Hall, and T. L. Daniel, "Luminance-dependent visual processing enables moth flight in low light," *Science*, vol. 348, no. 6240, pp. 1245–1248, 2015. [Online]. Available: http://www.sciencemag.org/content/348/6240/1245.abstract

[22] P. Yuan, K. Anselm, C. Hu, H. Nie, C. Lenox, A. Holmes, B. Streetman, J. Campbell, and R. McIntyre, "A new look at impact ionization-part ii: Gain and noise in short avalanche photodiodes," *IEEE Transactions on Electron Devices*, vol. 46, no. 8, pp. 1632–1639, Aug 1999.

[23] R. McIntyre, "Multiplication noise in uniform avalanche diodes," *IEEE Transactions on Electron Devices*, vol. ED-13, no. 1, pp. 164–168, Jan 1966.

[24] O. Shcherbakova, "3D camera based on gain-modulated CMOS avalanche photodiodes," Ph.D. dissertation, University of Trento, 2013.

[25] B. Nouri, M. Dandin, and P. Abshire, "Large-area low-noise single-photon avalanche diodes in standard cmos," in *Sensors, 2012 IEEE*. IEEE, 2012, pp. 1–5.

[26] P. Jerram, P. J. Pool, R. Bell, D. J. Burt, S. Bowring, S. Spencer, M. Hazelwood, I. Moody, N. Catlett, and P. S. Heyes, "The llccd: low-light imaging without the need for an intensifier," in *Photonics West 2001-Electronic Imaging*. International Society for Optics and Photonics, 2001, pp. 178–186.

[27] J. Hynecek, "Impactron-a new solid state image intensifier," *IEEE Transactions on Electron Devices*, vol. 48, no. 10, pp. 2238–2241, Oct 2001.

[28] Raptor Photonics. EMCCD - low light sensitivity and speed. [Online]. Available: http://www.raptorphotonics.com/technology/emccd-1

[29] Hamatsu. PMT (photomultiplier tubes). [Online]. Available: https://www.hamamatsu.com/us/en/3001.html

[30] Photonis. MCP-PMTs. [Online]. Available: http://www.photonis.com/en/ism/104-mcp-pmt.html

[31] B. Seitz, "Position sensitive photon detectors for nuclear physics, particle physics and healthcare applications," *Journal of Instrumentation*, vol. 7, no. 1, 2012. [Online]. Available: http://stacks.iop.org/1748-0221/7/i=01/a=C01031

[32] A. Eisele, R. Henderson, B. Schmidtke, T. Funk, L. Grant, J. Richardson, and W. Freude, "185 MHz count rate 139 dB dynamic range single-photon avalanche diode with active quenching circuit in 130 nm CMOS technology," in *Proc. Int. Image Sensor Workshop*, 2011, pp. 278–280.

[33] T. Delbruck and C. A. Mead, "Adaptive photoreceptor with wide dynamic range," in *IEEE International Symposium on Circuits and Systems*, vol. 4. IEEE, 1994, pp. 339–342.

[34] E. R. Fossum, "Photon counting without avalanche multiplication-progress on the quanta image sensor," *Proceedings of the Image Sensors*, 2016.

[35] E. R. Fossum, J. Ma, S. Masoodian, L. Anzagira, and R. Zizza, "The quanta image sensor: Every photon counts," *Sensors*, vol. 16, no. 8, p. 1260, 2016.

[36] E. R. Fossum, J. Ma, and S. Masoodian, "Quanta image sensor: Concepts and progress," in *SPIE Commercial+ Scientific Sensing and Imaging*. International Society for Optics and Photonics, 2016, pp. 985 804–985 804.

[37] J.-Y. Wu, S.-C. Li, F.-Z. Hsu, and S.-D. Lin, "Two-dimensional mapping of photon counts in low-noise," in *International Image Sensor Workshop*, Utah, 2013.

[38] D. Bronzi, F. Villa, S. Bellisai, S. Tisa, A. Tosi, G. Ripamonti, F. Zappa, S. Weyers, D. Durini, W. Brockherde *et al.*, "Large-area cmos spads with very low dark counting rate," in *SPIE OPTO*. International Society for Optics and Photonics, 2013, pp. 86 311B–86 311B.

[39] J. A. Richardson, E. A. Webster, L. A. Grant, and R. K. Henderson, "Scaleable single-photon avalanche diode structures in nanometer cmos technology," *IEEE Transactions on Electron Devices*, vol. 58, no. 7, pp. 2028–2035, 2011.

[40] B. Nouri, "Integrated single-photon sensing and processing platform in standard CMOS," Ph.D. dissertation, University of Maryland, College Park, Maryland, 2013.

[41] C. Niclass, A. Rochas, P.-A. Besse, and E. Charbon, "Design and characterization of a CMOS 3-D image sensor based on single photon avalanche diodes," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 9, pp. 1847–1854, 2005.

[42] C. Veerappan, J. Richardson, R. Walker, D.-U. Li, M. W. Fishburn, Y. Maruyama, D. Stoppa, F. Borghetti, M. Gersbach, R. K. Henderson *et al.*, "A 160× 128 single-photon image sensor with on-pixel 55ps 10b time-to-digital converter," in *IEEE International Solid-State Circuits Conference*. IEEE, 2011, pp. 312–314.

[43] D. Stoppa, D. Mosconi, L. Pancheri, and L. Gonzo, "Single-photon avalanche diode cmos sensor for time-resolved fluorescence measurements," *IEEE Sensors Journal*, vol. 9, no. 9, pp. 1084–1090, 2009.

[44] A. Torricelli, D. Contini, A. Pifferi, M. Caffini, R. Re, L. Zucchelli, and L. Spinelli, "Time domain functional NIRS imaging for human brain mapping," *NeuroImage*, vol. 85, Part 1, pp. 28 – 50, 2014.

[45] M. A. Karami, M. Gersbach, H.-J. Yoon, and E. Charbon, "A new single-photon avalanche diode in 90nm standard CMOS technology," *Opt. Express*, vol. 18, no. 21, pp. 22 158–22 166, Oct 2010.

[46] M. Gersbach, J. Richardson, E. Mazaleyrat, S. Hardillier, C. Niclass, R. Henderson, L. Grant, and E. Charbon, "A low-noise single-photon detector implemented in a 130nm CMOS imaging process," *Solid-State Electronics*, vol. 53, no. 7, pp. 803–808, 2009.

[47] M. Marwick and A. Andreou, "Fabrication and testing of single photon avalanche detectors in the TSMC 0.18um CMOS technology," in *Conf. on Information Sciences and Systems*, 2007, pp. 741–744.

[48] A. Rochas, A. R. Pauchard, P.-A. Besse, D. Pantic, Z. Prijic, and R. S. Popovic, "Low-noise silicon avalanche photodiodes fabricated in conventional cmos technologies," *IEEE Transactions on Electron Devices*, vol. 49, no. 3, pp. 387–394, 2002.

[49] A. Pauchard, P.-A Besse, and R. S. Popovic, "Simulations of a new CMOS compatible method to enhance the breakdown voltage of highly-doped shallow PN junctions," in *International Conference on Modeling and Simulation of Microsystems*, 1998.

[50] M. Dandin, A. Akturk, B. Nouri, N. Goldsman, and P. Abshire, "Characterization of single-photon avalanche diodes in a $0.5\mu m$ standard CMOS process–Part 1: Perimeter breakdown suppression," *IEEE Sens. J.*, vol. 10, pp. 1682–1690, 2010.

[51] M. Dandin, M. Habib, B. Nouri, P. Abshire, and N. McFarlane, "Characterization of single-photon avalanche diodes in a $0.5\mu m$ standard CMOS process–Part 2: Equivalent circuit model and geiger mode readout," *IEEE Sens. J.*, vol. (in press), 2016.

[52] M. Dandin and P. Abshire, "High signal-to-noise ratio avalanche photodiodes with perimeter field gate and active readout," *IEEE Electr. Device Lett.*, vol. 33, no. 4, pp. 570–572, 2012.

[53] H. Finkelstein, M. Hsu, and S. Esener, "STI-bounded single-photon avalanche diode in a deep submicrometer CMOS technology," *IEEE Electron Device Lett.*, vol. 27, pp. 887–889, 2006.

[54] J. A. Richardson, E. A. G. Webster, L. A. Grant, and R. K. Henderson, "Scaleable single-photon avalanche diode structures in nanometer cmos technology," *IEEE Trans. Electron Devices*, vol. 58, no. 7, pp. 2028–2035, July 2011.

[55] I. Rech, A. Ingargiola, R. Spinelli, I. Labanca, S. Marangoni, M. Ghioni, and S. Cova, "In-depth analysis of optical crosstalk in single-photon avalanche diode arrays," in *Optics East 2007*. International Society for Optics and Photonics, 2007, pp. 677 111–677 111.

[56] H. Xu, L. Pancheri, L. H. Braga, G.-F. Dalla Betta, and D. Stoppa, "Crosstalk characterization of single-photon avalanche diode (SPAD) arrays in CMOS 150nm technology," *Procedia Engineering*, vol. 87, pp. 1270–1273, 2014.

[57] A. Ficorella, L. Pancheri, G.-F. Dalla Betta, P. Brogi, G. Collazuol, P. Marrocchesi, F. Morsani, L. Ratti, and A. Savoy-Navarro, "Crosstalk mapping in CMOS SPAD arrays," in *European Solid-State Device Research Conference (ESSDERC)*. IEEE, 2016, pp. 101–104.

[58] D. Bronzi, S. Tisa, F. Villa, S. Bellisai, A. Tosi, and F. Zappa, "Fast sensing and quenching of CMOS SPADs for minimal afterpulsing effects," *IEEE Photonics Technology Letters*, vol. 25, no. 8, pp. 776–779, 2013.

[59] L. Pancheri and D. Stoppa, "Low-noise cmos single-photon avalanche diodes with 32 ns dead time," in *Solid State Device Research Conference, 2007. ESSDERC 2007. 37th European*. IEEE, 2007, pp. 362–365.

[60] E. A. Webster, L. A. Grant, and R. K. Henderson, "A high-performance single-photon avalanche diode in 130-nm CMOS imaging technology," *IEEE Electron Device Letters*, vol. 33, no. 11, pp. 1589–1591, 2012.

[61] F. Zappa, S. Tisa, A. Gulinatti, A. Gallivanoni, and S. Cova, "Monolithic CMOS detector module for photon counting and picosecond timing," in *Proceeding of the European Solid-State Device Research Conference*. IEEE, 2004, pp. 341–344.

[62] A. Dalla Mora, A. Tosi, F. Zappa, S. Cova, D. Contini, A. Pifferi, L. Spinelli, A. Torricelli, and R. Cubeddu, "Fast-gated single-photon avalanche diode for wide dynamic range near infrared spectroscopy," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 16, no. 4, pp. 1023–1030, July 2010.

[63] A. Tosi, A. D. Mora, F. Zappa, and S. Cova, "Single-photon avalanche diodes for the near-infrared range: detector and circuit issues," *Journal of Modern Optics*, vol. 56, no. 2-3, pp. 299–308, 2009.

[64] A. Ingargiola, M. Assanelli, A. Gallivanoni, I. Rech, M. Ghioni, and S. Cova, "Avalanche buildup and propagation effects on photon-timing jitter in si-spad with non-uniform electric field," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2009, pp. 73 200K–73 200K.

[65] M. Assanelli, A. Ingargiola, I. Rech, A. Gulinatti, and M. Ghioni, "Photon-timing jitter dependence on injection position in single-photon avalanche diodes," *IEEE Journal of Quantum Electronics*, vol. 47, no. 2, pp. 151–159, Feb 2011.

[66] J. A. Richardson, L. A. Grant, and R. K. Henderson, "Low dark count single-photon avalanche diode structure compatible with standard nanometer scale cmos technology," *IEEE Photonics Technology Letters*, vol. 21, no. 14, pp. 1020–1022, 2009.

[67] P. Magnan, "Detection of visible photons in ccd and cmos: A comparative view," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 504, no. 1, pp. 199–212, 2003.

[68] A. Fenigstein, A. Birman, and A. Lahav, "Night vision cmos image sensors pixel for sub-mililux light conditions," in *Proc. Of*, 2015.

[69] S. K. Mendis, S. E. Kemeny, R. C. Gee, B. Pain, C. O. Staller, Q. Kim, and E. R. Fossum, "Cmos active pixel image sensors for highly integrated imaging systems," *IEEE Journal of Solid-State Circuits*, vol. 32, no. 2, pp. 187–197, 1997.

[70] E. Webster, L. Grant, and R. Henderson, "A high-performance single-photon avalanche diode in 130-nm CMOS imaging technology," *IEEE Electron Device Letters*, vol. 33, no. 11, pp. 1589–1591, Nov 2012.

[71] C. Niclass, M. Sergio, and E. Charbon, "A single photon avalanche diode array fabricated in 0.35-μm cmos and based on an event-driven readout for tcspc experiments," in *Optics East 2006*. International Society for Optics and Photonics, 2006, pp. 63 720S–63 720S.

[72] C. Niclass, A. Rochas, P. Besse, and E. Charbon, "Design and characterization of a CMOS 3-D image sensor based on single photon avalanche diodes," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 9, pp. 1847–1854, Sept 2005.

[73] S. Mandai, M. W. Fishburn, Y. Maruyama, and E. Charbon, "A wide spectral range single-photon avalanche diode fabricated in an advanced 180 nm CMOS technology," *Opt. Express*, vol. 20, no. 6, pp. 5849–5857, Mar 2012. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-20-6-5849

[74] M. Gersbach, J. Richardson, E. Mazaleyrat, S. Hardillier, C. Niclass, R. Henderson, L. Grant, and E. Charbon, "A low-noise single-photon detector implemented in a 130nm CMOS imaging process," *Solid-State Electronics*, vol. 53, no. 7, pp. 803 – 808, 2009, papers Selected from the 38th European Solid-State Device Research Conference ESSDERC08. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0038110109000896

[75] J. A. Richardson, L. A. Grant, and R. K. Henderson, "Low dark count single-photon avalanche diode structure compatible with standard nanometer scale cmos technology," *IEEE Photonics Technology Letters*, vol. 21, no. 14, pp. 1020–1022, 2009.

[76] M. Liu, C. Hu, J. C. Campbell, Z. Pan, and M. M. Tashima, "Reduce afterpulsing of single photon avalanche diodes using passive quenching with active reset," *IEEE Journal of quantum electronics*, vol. 44, no. 5, pp. 430–434, 2008.

[77] S. Tisa, F. Guerrieri, A. Tosi, and F. Zappa, "100 kframe/s 8 bit monolithic single-photon imagers," in *ESSDERC*, Sept 2008, pp. 274–277.

[78] F. Zappa. 64x32 3D SPAD camera. [Online]. Available: http://www.everyphotoncounts.com/files/Datasheet_64x32_SPAD_camera.pdf

[79] L. Pancheri, N. Massari, and D. Stoppa, "SPAD image sensor with analog counting pixel for time-resolved fluorescence detection," *IEEE Transactions on Electron Devices*, vol. 60, no. 10, pp. 3442–3449, 2013.

[80] D. Chitnis and S. Collins, "Compact readout circuits for SPAD arrays," in *IEEE Transactions on Circuits and Systems*, 2010, pp. 357 – 360.

[81] B. Dierickx, B. Dupont, A. Defernez, and N. Ahmed, "Indirect X-ray photon-counting image sensor with 27T pixel and 15e-rms accurate threshold," in *IEEE Journal of Solid-State Circuits*, 2011, pp. 114 – 116.

[82] L. Gasparini, N. Massari, M. Perenzoni, L. Pancheri, and D. Stoppa, "Compact time-gated analog counting SPAD-based pixels for high resolution, single-photon, time-resolved imagers," in *Proc. Int. Image Sensor Workshop*, 2015.

[83] I. M. Antolovic, S. Burri, C. Bruschini, R. Hoebe, and E. Charbon, "Nonuniformity analysis of a 65-kpixel cmos spad imager," *IEEE Transactions on Electron Devices*, vol. 63, no. 1, pp. 57–64, 2016.

[84] J. P. Dyhr and C. M. Higgins, "Non-directional motion detectors can be used to mimic optic flow dependent behaviors," *Biological Cybernetics*, vol. 103, pp. 433– 446, 2010.

[85] M. Almasri, Z. Celik-Butler, D. P. Butler, A. Yaradanakul, and A. Yildiz, "Uncooled multimirror broad-band infrared microbolometers," *Journal of Microelectromechanical Systems*, vol. 11, no. 5, pp. 528–535, Oct 2002.

[86] J. P. Estrera, T. E. Ostromek, A. V. Bacarella, W. Isbell, M. J. Iosue, M. R. Saldana, and T. R. Beystrum, "Advanced image intensifier night vision system technologies: status and summary 2002," in *International Symposium on Optical Science and Technology*. International Society for Optics and Photonics, 2003, pp. 49–59.

[87] E. Heves, C. Ozturk, V. Ozguz, and Y. Gurbuz, "Solution-Based PbS Photodiodes, Integrable on ROIC, for SWIR Detector Applications," *IEEE Electron Device Letters*, vol. 34, no. 5, pp. 662–664, May 2013.

[88] X. Jiang, M. Itzler, K. ODonnell, M. Entwistle, M. Owens, K. Slomkowski, and S. Rangwala, "InP-Based Single-Photon Detectors and Geiger-Mode APD Arrays for Quantum Communications Applications," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 21, no. 3, pp. 5–16, May 2015.

[89] J. Beck, C. Wan, M. Kinch, J. Robinson, P. Mitra, R. Scritchfield, F. Ma, and J. Campbell, "The HgCdTe electron avalanche photodiode," in *2006 Digest of the LEOS Summer Topical Meetings*, 2006, pp. 36–37.

[90] B. Chen, W. Jiang, J. Yuan, A. L. Holmes, and B. M. Onat, "SWIR/MWIR InP-Based p-i-n Photodiodes With InGaAs/GaAsSb Type-II Quantum Wells," *IEEE Journal of Quantum Electronics*, vol. 47, no. 9, pp. 1244–1250, Sept 2011.

[91] H. W. Shin, S. J. Lee, D. G. Kim, M.-H. Bae, J. Heo, K. J. Choi, W. J. Choi, J.-w. Choe, and J. C. Shin, "Short-wavelength infrared photodetector on Si employing strain-induced growth of very tall InAs nanowire arrays," *Scientific reports*, vol. 5, p. 10764, 2015.

[92] S. Unlimited. SWIR Array Cameras. [Online]. Available: http://www.sensorsinc.com/products/area-cameras/

[93] M. LLC. Xenics Bobcat-640: Smallest SWIR Low Noise CL/GigE Camera Now in Production. [Online]. Available: www.xenics.com

[94] I. FLIR Systems. Tau SWIR. [Online]. Available: http://www.flir.com/tauswir/

[95] P. Richards, "Bolometers for infrared and millimeter waves," *Journal of Applied Physics*, vol. 76, no. 1, pp. 1–24, 1994.

[96] R. Bhan, R. Saxena, C. Jalwani, and S. Lomash, "Uncooled infrared microbolometer arrays and their characterisation techniques (review paper)," *Defence Science Journal*, vol. 59, no. 6, p. 580, 2009.

[97] F. Niklaus, C. Vieider, and H. Jakobsen, "Mems-based uncooled infrared bolometer arrays: a review," in *Photonics Asia 2007*. International Society for Optics and Photonics, 2007, pp. 68 360D–68 360D.

[98] F. Zappa, A. Tosi, and S. Cova, "Ingaas spad and electronics for low time jitter and low noise," in *International Congress on Optics and Optoelectronics*. International Society for Optics and Photonics, 2007, pp. 65 830E–65 830E.

[99] C. H. Tan, A. Marshall, M. J. Steer, and J. P. David, "The development of extremely low noise inas electron apds for photon counting applications in swir/mwir wavelengths," in *SPIE Europe Optics+ Optoelectronics*. International Society for Optics and Photonics, 2009, pp. 73 550Z–73 550Z.

[100] N. Calandri, M. Sanzaro, L. Motta, C. Savoia, and A. Tosi, "Optical crosstalk in ingaas/inp spad array: Analysis and reduction with fib-etched trenches," *IEEE Photonics Technology Letters*, vol. 28, no. 16, pp. 1767–1770, 2016.

[101] A. Tosi, A. Ruggeri, F. Villa, M. Sanzaro, M. Buttafava, N. Calandri, and F. Zappa, "Short-gate techniques for high-speed photon counting with ingaas/inp spads," in *CLEO: QELS_Fundamental Science*. Optical Society of America, 2016, pp. FW1C–4.

[102] C. Rafferty, C. King, B. Ackland, J. O'Neill, I. Aberg, T. Sriram, A. Mackay, and R. Johnson, "Monolithic germanium swir imaging array," in *Technologies for Homeland Security, 2008 IEEE Conference on*. IEEE, 2008, pp. 577–582.

[103] J. E. Carey and J. Sickler, "Black silicon sees further into the ir," *Laser Focus Word*, pp. 39–44, 2009.

[104] F. Koppens, T. Mueller, P. Avouris, A. Ferrari, M. Vitiello, and M. Polini, "Photodetectors based on graphene, other two-dimensional materials and hybrid systems," *Nature nanotechnology*, vol. 9, no. 10, pp. 780–793, 2014.

[105] E. Heves, C. Ozturk, V. Ozguz, and Y. Gurbuz, "Solution-based pbs photodiodes, integrable on roic, for swir detector applications," *IEEE Electron Device Letters*, vol. 34, no. 5, pp. 662–664, 2013.

[106] A. Rogalski, "Hgcdte infrared detector material: history, status and outlook," *Reports on Progress in Physics*, vol. 68, no. 10, p. 2267, 2005.

[107] ——, "Infrared detectors: an overview," *Infrared Physics & Technology*, vol. 43, no. 3, pp. 187–210, 2002.

[108] J. Boisvert, T. Isshiki, R. Sudharsanan, P. Yuan, and P. McDonald, "Performance of very low dark current SWIR PIN arrays," in *SPIE Defense and Security Symposium*. International Society for Optics and Photonics, 2008, pp. 69 400L–69 400L.

[109] B. M. Onat, W. Huang, N. Masaun, M. Lange, M. H. Ettenberg, and C. Dries, "Ultra-low dark current InGaAs technology for focal plane arrays for low-light level visible-shortwave infrared imaging," in *Defense and Security Symposium*. International Society for Optics and Photonics, 2007, pp. 65 420L–65 420L.

[110] M. MacDougal, J. Geske, C. Wang, S. Liao, J. Getty, and A. Holmes, "Low dark current ingaas detector arrays for night vision and astronomy," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2009, pp. 72 983F–72 983F.

[111] J. Boisvert, T. Isshiki, R. Sudharsanan, P. Yuan, and P. McDonald, "Design of very low dark current SWIR pin arrays," in *IEEE LEOS Summer Topical Meetings*, 2006, pp. 54–55.

[112] B. M. Onat, X. Jiang, and M. Itzler, "A systematic approach to dark current reduction in ingaas-based photodiode arrays for shortwave infrared imaging," in *IEEE LEOS Annual Meeting*, Oct 2009, pp. 231–232.

[113] R. F. Cannata, R. J. Hansen, A. N. Costello, and W. J. Parrish, "Very wide dynamic range SWIR sensors for very low background applications," in *AeroSense'99*. International Society for Optics and Photonics, 1999, pp. 756–765.

[114] Y. Ni, B. Arion, Y. Zhu, P. Potet, O. Huet, J. L. Reverchon, J. P. Truffer, J. A. Robo, and E. Costard, "Toward a single-chip TECless/NUCless InGaAs SWIR camera with 120-dB intrinsic operation dynamic range," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2011, pp. 80 121W–80 121W.

[115] M. H. Ettenberg, M. A. Blessinger, M. T. O'Grady, S.-C. Huang, R. M. Brubaker, and M. J. Cohen, "High-resolution SWIR arrays for imaging at night," in *Proc. SPIE*, vol. 5406, 2004, pp. 46–55.

[116] M. D. Enriquez, M. A. Blessinger, J. V. Groppe, T. M. Sudol, J. Battaglia, J. Passe, M. Stern, and B. M. Onat, "Performance of high resolution visible-ingaas imager for day/night vision," in *SPIE Defense and Security Symposium*. International Society for Optics and Photonics, 2008, pp. 69 400O–69 400O.

[117] Y. M. Jo, D. H. Woo, S. G. Kang, and H. C. Lee, "Very wide dynamic range ROIC with pixel-level ADC for SWIR FPAs," *IEEE Sensors J.*, vol. 16, no. 19, pp. 7227–7233, Oct 2016.

[118] R. F. Cannata, R. J. Hansen, A. N. Costello, and W. J. Parrish, "Very wide dynamic range SWIR sensors for very low background applications," in *AeroSense'99*. Intl. Society for Optics and Photonics, 1999, pp. 756–765.

[119] D. Dzahini and H. Ghazlane, "Auto-zero stabilized cmos amplifiers for very low voltage or current offset," in *Nuclear Science Symposium Conference Record, 2003 IEEE*, vol. 1. IEEE, 2003, pp. 6–10.

[120] C. A. Mead and T. P. Allen, "Cmos amplifier with offset adaptation," Nov. 26 1991, uS Patent 5,068,622.

[121] ——, "Subthreshold cmos amplifier with offset adaptation," Jun. 19 1990, uS Patent 4,935,702.

[122] J. S. Jandu and I. Miller, "Operational amplifier having improved digitally adjusted null offset," Jun. 12 1990, uS Patent 4,933,643.

[123] J. Hasler and H. B. Marr, "Finding a roadmap to achieve large neuromorphic hardware systems," *Frontiers in Neuroscience*, vol. 7, no. 118, 2013. [Online]. Available: http://www.frontiersin.org/neuromorphic_engineering/10.3389/fnins.2013.00118/abstract

[124] C. Diorio, S. Mahajan, P. Hasler, B. Minch, and C. Mead, "A high-resolution non-volatile analog memory cell," in *Circuits and Systems, 1995. ISCAS'95., 1995 IEEE International Symposium on*, vol. 3. IEEE, 1995, pp. 2233–2236.

[125] Y. L. Wong, M. Cohen, and P. Abshire, "A 128x128 floating gate imager with self-adapting fixed pattern noise reduction," in *IEEE International Symposium on Circuits and Systems*, vol. 5, May 2005, pp. 5314–5317.

[126] Y. L. Wong, M. H. Cohen, and P. A. Abshire, "A 750-MHz 6-b adaptive floating-gate quantizer in 0.35-μm CMOS," *IEEE TCAS I: Reg. Papers*, vol. 56, no. 7, pp. 1301–1312, 2009.

[127] A. Basu, S. Brink, C. Schlottmann, S. Ramakrishnan, C. Petre, S. Koziol, F. Baskaya, C. M. Twigg, and P. Hasler, "A floating-gate-based field-programmable analog array," *IEEE J. Solid-State Circuits*, vol. 45, no. 9, pp. 1781–1794, Sept 2010.

[128] P. Hasler, "Foundations of learning in analog VLSI," Ph.D. dissertation, California Institute of Technology, 1997.

[129] T. S. Hall, "Field-programmable analog arrays: A floating-gate approach," Ph.D. dissertation, Georgia Institute of Technology, 2004.

[130] P. Hasler, B. A. Minch, and C. Diorio, "An autozeroing floating-gate amplifier," *IEEE TCAS II, Analog Digit. Signal*, vol. 48, no. 1, pp. 74–82, Jan 2001.

[131] C. Mead, "Adaptive retina," in *Analog VLSI implementation of neural systems*. Springer, 1989, pp. 239–246.

[132] P. Hasler, B. A. Minch, and C. Diorio, "Adaptive circuits using pfet floating-gate devices," in *Advanced Research in VLSI, 1999. Proceedings. 20th Anniversary Conference on*. IEEE, 1999, pp. 215–229.

[133] ——, "Floating-gate devices: they are not just for digital memories any more," in *Circuits and Systems, 1999. ISCAS'99. Proceedings of the 1999 IEEE International Symposium on*, vol. 2. IEEE, 1999, pp. 388–391.

[134] M. Islam, J. Y. Feng, A. Berkovich, P. Abshire, G. Barrows, and F.-S. Choa, "InGaAs/InP PIN photodetector arrays made by MOCVD based zinc diffusion processes," in *Proc. SPIE*, vol. 9819, 2016, pp. 98 190G–98 190G–7.

[135] H. G. Krapp, R. Hengstenberg *et al.*, "Estimation of self-motion by optic flow processing in single visual interneurons," *Nature*, vol. 384, no. 6608, pp. 463–466, 1996.

[136] G. K. Taylor and H. G. Krapp, "Sensory systems and flight stability: what do insects measure and why?" *Advances in insect physiology*, vol. 34, pp. 231–316, 2007.

[137] M. V. Srinivasan, S. Zhang, M. Altwein, and J. Tautz, "Honeybee navigation: Nature and calibration of the" odometer"," *Science*, vol. 287, no. 5454, pp. 851–853, 2000.

[138] M. Srinivasan, S. Zhang, and N. Bidwell, "Visually mediated odometry in honeybees," *Journal of Experimental Biology*, vol. 200, no. 19, pp. 2513–2522, 1997.

[139] P. S. Bhagavatula, C. Claudianos, M. R. Ibbotson, and M. V. Srinivasan, "Optic flow cues guide flight in birds," *Current Biology*, vol. 21, no. 21, pp. 1794–1799, 2011.

[140] H. Bruggeman, W. Zosh, and W. H. Warren, "Optic flow drives human visuo-locomotor adaptation," *Current Biology*, vol. 17, no. 23, pp. 2035 – 2040, 2007.

[141] A. S. Etienne, R. Maurer, and V. Séguinot, "Path integration in mammals and its interaction with visual landmarks." *Journal of Experimental Biology*, vol. 199, no. 1, pp. 201–209, 1996.

[142] M. Narayana, A. Hanson, and E. Learned-Miller, "Coherent motion segmentation in moving camera videos using optical flow orientations," in *Proceedings of the IEEE International Conference on Computer Vision*, 2013, pp. 1577–1584.

[143] A. M. Tekalp, *Digital video processing*. Prentice Hall Press, 2015.

[144] J.-Y. Chang, W.-F. Hu, M.-H. Cheng, and B.-S. Chang, "Digital image translational and rotational motion stabilization using optical flow technique," *IEEE Transactions on Consumer Electronics*, vol. 48, no. 1, pp. 108–115, 2002.

[145] K. Prazdny, "Egomotion and relative depth map from optical flow," *Biological cybernetics*, vol. 36, no. 2, pp. 87–102, 1980.

[146] G. Johansson, "Visual motion perception." *Scientific American*, 1975.

[147] T. Brox and J. Malik, "Large Displacement Optical Flow: Descriptor Matching in Variational Motion Estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 3, pp. 500–513, March 2011.

[148] L. Xu, J. Jia, and Y. Matsushita, "Motion Detail Preserving Optical Flow Estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 9, pp. 1744–1757, Sept 2012.

[149] T. Amiaz, E. Lubetzky, and N. Kiryati, "Coarse to over-fine optical flow estimation," *Pattern Recognition*, vol. 40, no. 9, pp. 2496 – 2503, 2007.

[150] B. McCane, K. Novins, D. Crannitch, and B. Galvin, "On benchmarking optical flow," *Computer Vision and Image Understanding*, vol. 84, no. 1, pp. 126 – 143, 2001.

[151] H. Liu, T.-H. Hong, M. Herman, T. Camus, and R. Chellappa, "Accuracy vs efficiency trade-offs in optical flow algorithms," *Computer Vision and Image Understanding*, vol. 72, no. 3, pp. 271 – 286, 1998.

[152] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black, "A naturalistic open source movie for optical flow evaluation," in *Proceedings of the 12th European Conference on Computer Vision - Volume Part VI*. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 611–625.

[153] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.

[154] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. J. Black, and R. Szeliski, "A database and evaluation methodology for optical flow," *International Journal of Computer Vision*, vol. 92, no. 1, pp. 1–31, 2011.

[155] A. Barjatya, "Block matching algorithms for motion estimation," *IEEE Transactions Evolution Computation*, vol. 8, no. 3, pp. 225–239, 2004.

[156] ——. Block matching algorithms for motion estimation. [Online]. Available: http://www.mathworks.com/matlabcentral/fileexchange/8761-block-matching-algorithms-for-motion-estimation

[157] M. Srinivasan, "An image-interpolation technique for the computation of optic flow and egomotion," *Biological Cybernetics*, vol. 71, no. 5, pp. 401–415, 1994.

[158] W. E. Green, P. Y. Oh, and G. Barrows, "Flying insect inspired vision for autonomous aerial robot maneuvers in near-earth environments," in *International Conference on Robotics and Automation (ICRA)*, vol. 3, April 2004, pp. 2347–2352 Vol.3.

[159] P. Xu, P. Abshire, and J. S. Humbert, "Motion image sensor with on-chip adaptation and programmable filtering," in *IEEE International Symposium of Circuits and Systems (ISCAS)*, May 2011, pp. 2425–2428.

[160] F. Iida and D. Lambrinos, "Navigation in an autonomous flying robot by using a biologically inspired visual odometer," in *Intelligent Systems and Smart Manufacturing*. International Society for Optics and Photonics, 2000, pp. 86–97.

[161] S. B. Fuller and R. M. Murray, "A hovercraft robot that uses insect-inspired visual autocorrelation for motion control in a corridor," in *IEEE International Conference on Robotics and Biomimetics*, Dec 2011, pp. 1474–1481.

[162] J. S. Humbert, A. Hyslop, and M. Chinn, "Experimental validation of wide-field integration methods for autonomous navigation," in *International Conference on Intelligent Robots and Systems*, Oct 2007, pp. 2144–2149.

[163] G. Orchard and R. Etienne-Cummings, "Bioinspired visual motion estimation," *Proceedings of the IEEE*, vol. 102, no. 10, pp. 1520–1536, 2014.

[164] C. M. Higgins, "Nondirectional motion may underlie insect behavioral dependence on image speed," *Biological Cybernetics*, vol. 91, no. 5, pp. 326–332, 2004.

[165] J. P. Dyhr and C. M. Higgins, "Non-directional motion detectors can be used to mimic optic flow dependent behaviors," *Biological Cybernetics*, vol. 103, no. 6, pp. 433–446, 2010.

[166] B. Hassenstein and W. Reichardt, "Systemtheoretische analyse der zeit-, reihenfolgenund vorzeichenauswertung bei der bewegungsperzeption des rsselkfers chlorophanus," *Z. Naturforschg.*, vol. 11, pp. 513–524, 1956.

[167] J. Haag, W. Denk, and A. Borst, "Fly motion vision is based on reichardt detectors regardless of the signal-to-noise ratio," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 101, no. 46, pp. 16 333–16 338, 2004.

[168] W. Reichardt and T. Poggio, "Visual control of orientation behaviour in the fly," *Quarterly Reviews of Biophysics*, vol. 9, no. 3, pp. 311–375, 1976.

[169] E. Buchner, "Behavioural analysis of spatial vision in insects," in *Photoreception and Vision in Invertebrates*, 1984, vol. 74, pp. 561–621.

[170] Z. R. Alvidrez, "Computational modeling of neurons involved in fly motion detection," Master's thesis, The University of Arizona, August 2005.

[171] J. C. Theobald, B. Greiner, W. T. Wcislo, and E. J. Warrant, "Visual summation in night-flying sweat bees: A theoretical study," *Vision Research*, vol. 46, no. 14, pp. 2298 – 2309, 2006. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0042698906000071

[172] P. Ala-Laurila, "Visual neuroscience: How do moths see to fly at night?" *Current Biology*, vol. 26, no. 6, pp. R231–R233, 2016.

[173] B. Greiner, "Adaptations for nocturnal vision in insect apposition eye," Ph.D. dissertation, Lund University, Lund, Sweden, 12 2005.

[174] C. M. Higgins and S. A. Shams, "A biologically inspired modular vlsi system for visual measurement of self-motion," *IEEE Sensors Journal*, vol. 2, no. 6, pp. 508–528, 2002.

[175] S. Lim and A. El Gamal, "Gain fixed pattern noise correction via optical flow," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 51, no. 4, pp. 779–786, 2004.

[176] M. Perenzoni, N. Massari, D. Stoppa, L. Pancheri, M. Malfatti, and L. Gonzo, "A 160x120-pixels range camera with in-pixel correlated double sampling and fixed-pattern noise correction," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 7, pp. 1672–1681, 2011.

[177] R. Panicacci, B. Mansoorian, and E. Fossum, "Ram line storage for fixed pattern noise correction," Oct. 29 2002, uS Patent 6,473,124.