A 10μJ/Frame, 1mm^2 128$\times$128 CMOS Active Pixel Image Sensor

by Jonathan Choy

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

Kristofer S. J. Pister
Research Advisor

12/18/03
Date

* * * * * *

Roger T. Howe
Second Reader

December 18, 2003
Date

Abstract

Just five years ago, digital cameras were considered a technological luxury appreciated by only a few, and it was said that digital image quality would always lag behind that of conventional film cameras. However, they have since become mainstream as technology has improved, and can now be found in many devices that we can buy, from personal digital assistants, to cellular phones, and even built-in on laptop computers. Clearly, digital imaging is not just a fad that will go away, but is here to stay. Looking toward the future, one possible application of a digital camera (or imager) is to serve as part of an autonomous sensor network that discretely monitors the environment. To function as a miniature sensor node, such an imager should be able to run for a long time without needing constant replacement (as when the battery dies); otherwise, this would defeat the purpose of the "autonomous" node. In this project, a black-and-white 10μJ-per-frame, 1mm^2 camera with a resolution of 128×128 pixels is designed, fabricated, and verified. It is found that even at low voltages (1V) and power (60μW), this chip is able to identify objects that are focused onto it. And though the imager's performance is not as good as was initially hoped, we can nevertheless conclude that the ultra-low-power CMOS image sensor needed as we head toward a 1mm^3 sand-grain-sized camera can indeed be realized.

Acknowledgements

A project of this magnitude is not accomplished alone, and I could not have come this far without the support of the following people whom I acknowledge now.

First I would like to thank my advisor, Dr. Kris Pister, for funding the research involved in this project and for sticking with me even during times when I would lose hope. It is his vision and dream of changing the world that gives me the assurance that with enough confidence and belief, anything is possible.

Next I would like to acknowledge all the members of Prof. Pister's research group: Brian, who advised me on parts of my design, wire-bonded some of my early test-chips, and taught me everything I knew about image sensors before I began to wade my way through mountains of literature on the subject; Al, for having an answer to any question that I could come up with, be it design-related or a Cadence problem; Mike, for the design advice, the LabVIEW assistance, and for being the captain of a championship softball team; Ben, for keeping us all sane during those all-nighers around tape-out; and everyone else in 471 Cory for the support and enthusiasm that helped me believe that everything would work out in the end.

In addition, I must acknowledge a number of my colleagues – Hanching, Eddie, Rich, Nate, and Janie – who have been there with me in graduate school from the start, and have often provided a much needed break from my research work.

Finally, I would not be anywhere without my friends and family – especially my parents who do not understand the details of any of the projects I have worked on for the past six years, yet encourage and support me with a blind faith I have not seen elsewhere.

Table of Contents

List of Figures

Chapter 1

Introduction & Motivation

With technology trends as they are, it is not difficult to imagine a world in which everything is monitored continuously by sand-grain-sized sensors that can be organized in an autonomous network. Yet the realization of such a network requires the development of "revolutionary advances in miniaturization, integration, and energy management" [48]. This concept is at the heart of the Smart Dust idea.

Sensor and circuit technologies have advanced at such an alarming rate in recent years that the minuscule components for this complex system of sensors can now be realized. For their part, these sensors would monitor a variety of elements of their environment – including temperature and the location of its neighboring sensor nodes – and report it to interested parties. Extending this idea further, one can imagine an image sensor (or in other words, a camera) whose purpose is to discretely capture images of the surrounding environment.

Such an imager can be used as (1) a standalone sensor (e.g., a button-sized device that can be concealed), (2) part of a more complex micro-robot system (in which the camera serves as the robot's eyes), or even (3) as an additional feature of another device such as a watch or a piece of jewelry. For any of these situations, the camera must meet a number of stringent requirements with regard to size and power consumption, and to this end, some performance may have to be traded off. Finally, it would be most desirable to implement the camera in a widely-available standard (e.g., CMOS) technology as this offers the lowest-cost solution with the highest compatibility of integration with other systems.

In this report, the design and test of such a camera system is explored. It begins with a survey of the different camera architectures in the literature and examines their suitability for the desired ultra-low-power standard-CMOS implementation. Then, a number of camera-related figures of

merit are defined and described. After that, the design of a CMOS Active Pixel Sensor-based camera is detailed, followed by the test results of the designed system. Finally, this report concludes with a discussion of future improvements to enhance the camera's performance.

Chapter 2

Architectures

At an abstract level, digital camera architectures consist of three stages: (1) optics: lenses, filters, etc. to focus the desired image onto the camera chip; (2) an image sensor, which takes the light information projected through the optics and generates a set of representative voltages or currents; and (3) circuitry to interpret and process these signals, eventually creating the visual image. Architectural decisions in each stage have direct consequences on camera performance metrics such as speed and power consumption.

The most important feature of any micro-scale device is low power consumption. Such a circuit could function for an extended period of time on either solar power or a small watch battery, but at the cost of performance: namely, less vivid images and a slower frame-rate in the case of a camera.

In modern digital camera implementations, many tricks are used to implement CMOS digital cameras. This chapter gives a broad overview of these techniques, and weighs them with respect to the low-power design target.

2.1 Optics

At minimum, a camera's optics system is a lens that focuses incoming light from an area of interest onto the image detector. Though the details of a particular lens configuration can involve complicated physics and optics theory, the end result is quite simple: the system takes an object of interest and projects a clear miniaturized image of it onto the image sensor, as shown in Figure 2.1. In more complex systems, filters may be used in conjunction with the lens assembly to (1) reduce the incidence of non-visible (e.g., infrared) wavelengths that are not of interest [44], and (2) create

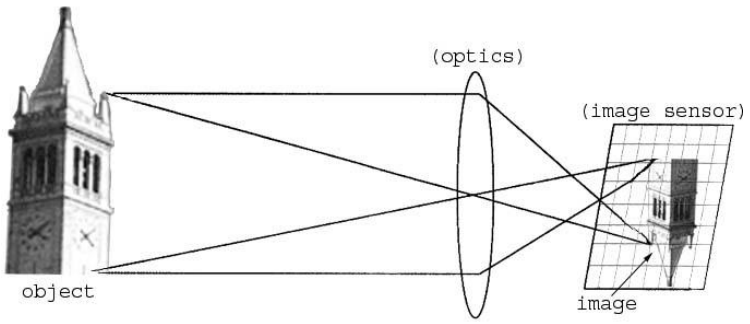

Figure 2.1: High-level depiction of a camera system's optics component. In this picture, a lens is used to form a miniaturized image of a clock tower onto a sensor. The image sensor itself is shown as a rectangular array of pixels. *(Image inspired by similar figure given in [18])*

color images (by using color filters).

2.2 Image Sensors

A camera's image sensor is usually a rectangular array of pixels, each of which contains a light-sensitive device (or photo-detector) that generates a quantity of electrical charge proportional to the intensity of the incident light. [43]. In a rectangular array, an entire image is represented by keeping track of the light intensity detected by each individual pixel. The two most common silicon-based photo-detectors today are the photogate and the photodiode, both of which are compatible with modern standard CMOS processes. In normal operation, both devices have a charge-free depletion region in which electrical charges are created and stored upon exposure to light. However, in all other aspects of operation, these two devices provide differing modes of operation.

2.2.1 Photogates

At its core, a photogate is just a metal-oxide-semiconductor (MOS) capacitor biased in depletion, a typical cross-section of which is shown in Figure 2.2. Under this bias condition, the substrate immediately below the MOS gate does not contain any conducting charges. Consequently, when light strikes the device (penetrating the gate and oxide layers into the depletion region) and generates an electron-hole pair, one of the generated charges will recombine in the bulk while the other remains trapped in the depletion region. The total number of trapped charges in the depletion

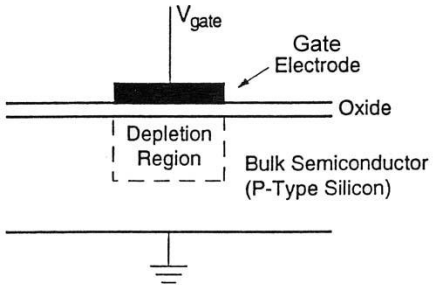

Figure 2.2: Cross-section of typical MOS capacitor. The gate electrode is commonly fabricated in polysilicon instead of metal. When biased in depletion, a region devoid of electrical charges forms underneath the gate at the surface of the bulk. More details about MOS capacitors can be found in any device physics reference and are beyond the scope of this project. *(Image inspired by similar figure in [44])*

region provides an indirect measure of incident light intensity.

2.2.2 Photodiodes

On the other hand, a photodiode is a simpler device: it is usually implemented as a reverse-biased p/n junction diode. Depending on the specific CMOS process, a number of p/n junction configurations are possible, as compared in Table 2.1. As described in the p/n junction literature, a depletion

Table 2.1: Comparison of the general performance of different types of PN junctions available in a standard n-well CMOS process [43].

	Time Response	**Current Level**	**Dynamic Range**	**Layout**
$n-well/p-sub$	Fast	Low	5-7 decades	Simple
$n^+/p-sub$	Fast	Low	5-7 decades	Simple
$p^+/n-well/p-sub$	Slow	High	3-5 decades	Complex

region forms at the p/n interface of a reverse-biased diode, and the charge storage capacity is determined by this region's width. When the device is exposed to light, incident photons generate charges in this depletion region that become reverse-bias diode current as they are swept by the induced electric field into the quasi-neutral P and N regions (depending on charge polarity). The amount of generated charge (and reverse-bias current) is proportional to the light intensity.

In a photodiode-based pixel, the photodiode must first be reverse-biased by a MOS switch, and then left to float. As current flows in response to incident light, the voltage across the photodiode drops, and either the reverse-bias current or the voltage drop on the photodiode can used as a measure of light intensity. After reading the photodiode voltage or current, control circuits reset the diode to the initial reverse bias [44].

Photogates vs. Photodiodes

While either photogates or photodiodes can be used in all pixel architectures, the specifications of a particular application will determine which is better-suited. For the desired low-power standard-CMOS camera, a comparison of photo-detector complexity and expected performance is necessary before choosing one over the other.

First, photogates are more complex than photodiodes, requiring carefully-timed pulses to first bias a device into depletion and then transfer the collected charge to a neighboring charge-storage device [43]. On the other hand, photodiodes simply need to be reset, released, and reset once more. Second, photodiodes have been shown (in [33] and [37]) to exhibit higher responsivity [13] and quantum efficiency than photogates over the visible-light spectrum since they do not have an overlying polysilicon gate [40]. Third, in the fabrication of photogate devices, the process' overlying polysilicon gate and oxide layer must be thin enough to allow incident photons to pass through to the depletion region to generate measurable charge. At the same time, the presence of the overlying gate affects the photogate's response toward the blue end of the visible spectrum, as mentioned in [33], [43], and [44].

On the other hand, photogate structures lend themselves easily to a full correlated double sampling scheme [40], resulting in superior noise performance to that of a photodiode structure (in which only 1/f and fixed-pattern noise can be reliably eliminated – reset noise cannot be canceled completely [13]). On top of this, photogates demonstrate a higher conversion gain than do photodiodes [33].

In the end, a photodiode-type pixel architecture is chosen for this system, since the thin polysilicon gate (around 500Å thick [33]) necessary for a photogate is not available in the standard-CMOS process in which this chip will be fabricated. Nevertheless, since both photogates and photodiodes exhibit very similar characteristics of anti-blooming, image-lag, sensitivity, and dynamic range [33], [40], this choice should not limit the imager's performance.

2.2.3 Charge Conversion

Regardless of which photo-detector is used, the other circuitry in the pixel must convert the photo-generated charge into a useful signal (a current or voltage) at the image sensor output [18]. There are two levels at which to perform this conversion: (1) at the array or column level (one charge-to-voltage amplifier is shared by either all pixels in the array, or those in one column), or (2) at the pixel level (one amplifier per pixel). This choice involves a trade-off between pixel fill-factor (the percentage of pixel area exposed to light) and performance [14], [43]. In the area of solid-state imagers, two technologies have emerged: charge-coupled devices (CCD), and CMOS imagers (both passive and active). The CCDs use an array-level amplification scheme, and the passive CMOS pixels can use either an array-level or a column-level scheme, whereas the active CMOS pixels use pixel-level amplification.

Charge-Coupled Devices (CCD)

In a CCD implementation, photo-generated charges are transferred via charge coupling to an adjacent charge-storage node (often a neighboring pixel) and eventually, to an output amplifier, where a voltage can be read out. A typical CCD pixel contains three photogates, of which only one is exposed to light. The other two are used to store and/or transfer charge during readout, as illustrated in Figure 2.3. However, for the charge transfer to be efficient, high operating voltages are needed, as well as a modified (non-conventional CMOS) fabrication process [18]. In the end, the CCD imager provides very low readout noise, high dynamic range, and good detector sensitivity. But despite these performance advantages, it is difficult to integrate a CCD sensor on a chip with CMOS readout circuitry, as both parts must be fabricated under different processes and would require different operating supply voltages.

CMOS Image Sensors

CMOS imagers, on the other hand, do not rely on charge-transfer as heavily as a CCD. Instead, the amount of photo-generated charge is read out directly by addressing individual pixels via a pixel-selection pass-transistor. As with a DRAM, reading out values from a CMOS imager array is done by first activating a row of pixels and then selecting a column, as shown in Figure 2.4. However, when compared to CCD imagers, CMOS imager configurations introduce additional readout noise, resulting in reduced sensitivity and dynamic range [11]. Yet despite these drawbacks, CMOS

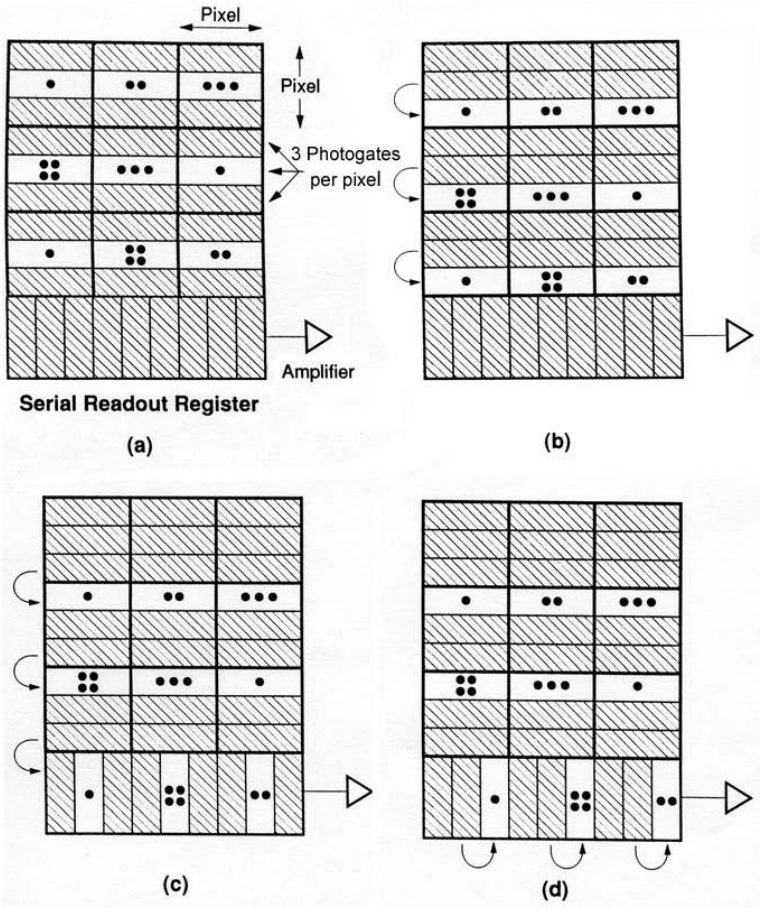

Figure 2.3: Diagram illustrating CCD array operation. **(a)** Each pixel really contains three photogate devices – one to be exposed to light, and the two surrounding ones to be used for charge shifting. **(b)** After exposure, charge is shifted (vertically) row by row to a readout buffer, which is a linear array of charge-storage states. **(c)-(d)** From there, they are shifted (horizontally) to a charge-to-voltage amplifier, eventually producing a serial voltage readout of the charge collected at each pixel. *(Image similar to one shown in [44])*

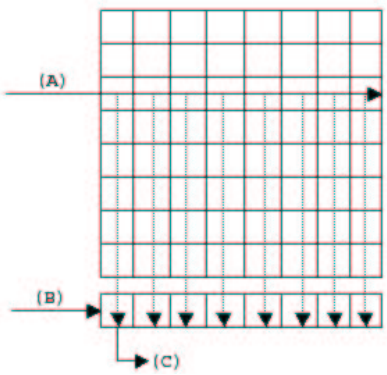

Figure 2.4: High-level diagram illustrating CMOS image sensor pixel readout. Much like reading out of a DRAM, first one row is activated (its values copied to the columns); and second, one particular value is chosen from all the columns in the row, using switches or an analog multiplexor.

sensors have the advantages of easy (single-chip) integration with ancillary CMOS circuitry, lower power consumption, and lower fabrication cost [18], [17].

CMOS imagers can further be subdivided into two classes: passive pixel sensors (PPS), and active pixel sensors (APS). To begin with, passive pixels operate much like single-transistor DRAM cells; in fact, a passive pixel is identical to a 1T DRAM cell in architecture except that the storage capacitor is replaced by a photo-detector [14], [18]. In a typical PPS array configuration (shown in Figure 2.5), each pixel has a selection transistor that governs charge flow from the photo-detector to a column-readout wire, where the charge is converted to a voltage by a column-shared output amplifier [14]. And as each pixel contains only this one transistor, high fill-factor can be achieved, resulting in high quantum efficiency. However, the passive-pixel design suffers from high readout noise and slow performance, both of which become progressively worse with increasing resolution [14]. To eliminate the shortcomings of a PPS, an active pixel sensor can be used. In an APS cell, the charge-to-voltage amplification stage is brought into each pixel, thus reducing readout noise while trading off fill factor [14]. In addition, since the highly capacitive column wires are now being driven by amplifiers, an APS system is much faster than its PPS counterpart, and scales to high resolutions much better. Though an APS pixel has a much smaller fill factor than a PPS pixel, the resulting reduction in optical signal is more than atoned for by the increase in dynamic range, resulting in a net performance gain with APS [14]. In fact, it has been shown that with advances in CMOS technology, APS system performance has become indistinguishable from CCD performance [15]. Given the inherent advantages of a CMOS implementation (e.g., lower cost and lower

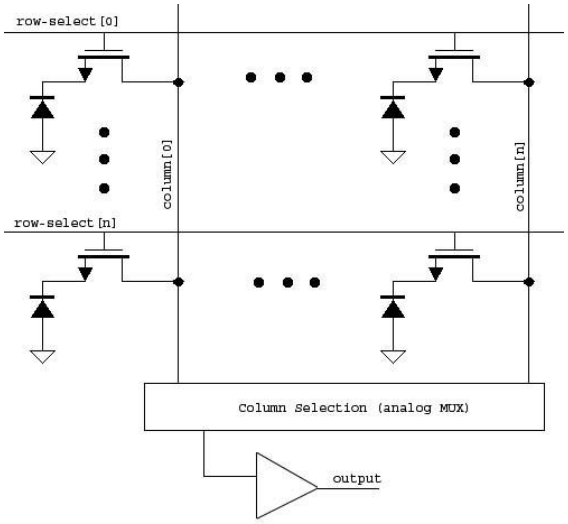

Figure 2.5: Schematic of a passive pixel array. In actual array operation, exactly one of the n row-select signals is asserted. Then, photo-generated charge from each of the pixels in this selected row will flow to the corresponding column. Finally, the column-select logic (often an analog multiplexor) will choose one of these columns to pass to an output amplifier.

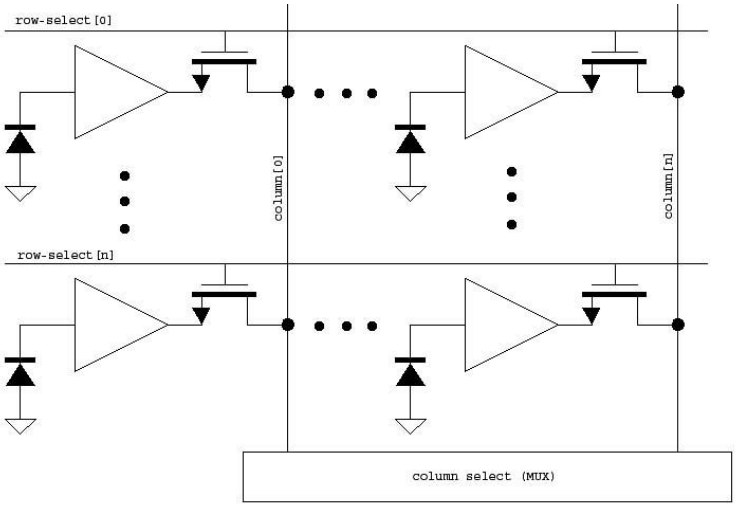

Figure 2.6: Schematic of an active pixel array. As with the passive pixel array of Figure 2.5, one row is selected at a time by asserting a row-select signal, and then one particular pixel's value chosen by the column-selection circuitry. The difference is that the amplification stage (often just a common-drain amplifier) is brought into the pixel level.

power), the best choice for a low-power camera would be an APS implementation.

2.2.4 Sampling & Charge Integration

The responsivity of silicon photo-detectors created in standard CMOS processes falls well short of that of the human eye. An eye can detect light over 14 decades of power (or 280dB); by contrast, the best standard CMOS photo-detectors only can detect 5-7 decades (100-140dB) [43]. To address this problem, a number of adaptation mechanisms are used to calibrate a photo-detector to its current environment. One such calibration method is charge integration.

In an active pixel, photo-generated charge is integrated on the total parasitic capacitance of the reverse-biased photodiode, and converted to a voltage. In turn, the integration time is controlled by a sample-and-hold switch that alternately resets (and pre-charges) the node to a known value, and then releases it during integration [8], [11] (and others). The voltage on this node then changes proportional to the amount of photo-generated charge integrated onto the capacitance. Intuitively, the magnitude of this voltage change can be controlled by varying the integration time of the pixel, but ultimately, the integration time is bound by the maximum desired clock frequency and the amount of dark current.

2.3 Readout Electronics & System-Level Architectures

After it is sensed by the imager, the image data must be processed into a visual image. In the chosen CMOS APS architecture, a number of decisions must be made regarding the readout circuitry that affect the camera's speed, power consumption, and image quality. Specifically, the imager-control circuitry (simultaneous vs. rolling shutter), the sampling policy (CDS, and how the data from the image sensor is read by the readout circuit), and post-processing (gain stages, ADC, etc.) must be determined.

2.3.1 Imager Control - Snapshot vs. Rolling Shutter

An important principle in photography is exposure – the period during which the image is imprinted onto film (conventional film camera), or image intensity sensed (solid-state imager). Ideally, all of the pixels in a digital imager are exposed simultaneously and for the same duration, in order to capture the same image in time and space. Following this exposure period, the corre-

sponding charge or voltage generated in each pixel is maintained to within a certain tolerance long enough for every pixel value to be read out and processed. As described, this is referred to as a simultaneous shutter, or a "snapshot" mode of operation [41], [23].

On the other hand, depending on the speed and pixel-size requirements, the amount of capacitance needed to hold the sampled voltage may not be achievable and thus, a snapshot mode of operation impossible. In such cases, one alternative is to use an electronic "rolling shutter" [23] (a technique used in most CMOS imagers today). In this scheme, the successive operations on a row (namely, reset, exposure, and readout) are staggered throughout the array. For example, in a typical rolling shutter architecture [11], row n is sampled onto the columns and read out at the same time as row $n-1$ is reset; meanwhile, all the other rows in the entire array are being exposed. A sample timing diagram for this is shown in Figure 2.7. As shown, the operations on a row are staggered from one to the next, and effectively "roll" down the array.

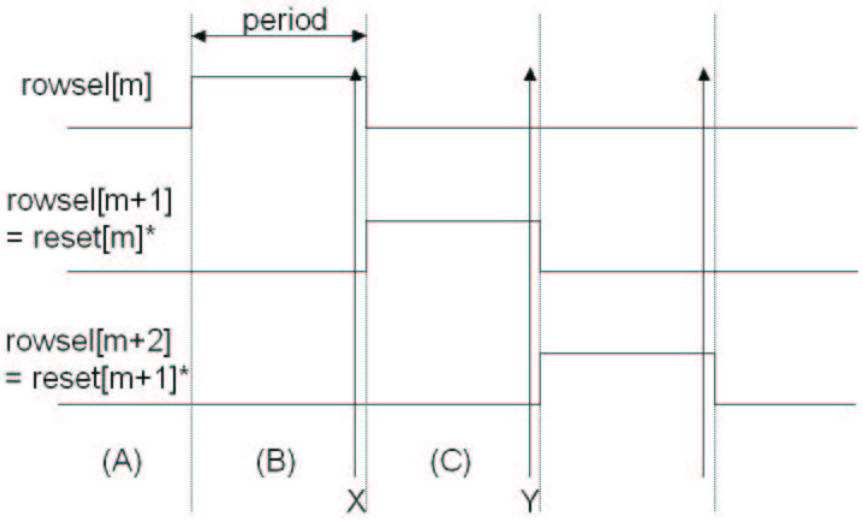

Figure 2.7: Relative row timing in rolling-shutter architecture. The regions are labeled as follows: **(A)-(B)** row m is exposed; during **(B)**, row $m-1$'s values are on column capacitances and are processed by read-out circuitry. At point **X**, the values of row m are sampled onto column capacitances, replacing the values from row $m-1$ (which are no longer needed). In **(C)**, row m is reset (but its sampled values are held on the column capacitances and are being processed by read-out) while at the same time, row $m+1$ is in its last stage of exposure (just as row m was during **(B)**). Finally, at point **Y**, the values of row $m+1$ are sampled onto the column capacitors, overwriting the values of row m that are no longer needed.

2.3.2 Data Sampling

In a system where pixels are successively reset, exposed, and sampled onto a column, the photo-generated signals must be sampled only when the system output is meaningful (since there will be times when the output is at some constant reset value and does not give a light intensity measurement, for instance). Naturally, the process of sampling data introduces extra complexity into the system with regard to signal timing and noise, but both of these can be dealt with quite easily: signal timing only requires more complex digital logic, and the increased sampling noise can be canceled out using such tricks as correlated double sampling [14].

2.3.3 Interfacing with other Electronics (analog-to-digital conversion)

So far, the system described produces a single analog output voltage stream that is proportional to the intensity of the light at each pixel in an array. For this to be useful in data transmission (as is the eventual goal of this low-power camera system), the analog voltage should be converted to a digital signal, and these digital bits transmitted. Naturally, an analog-to-digital converter (ADC) can be used [7], [15], [32], [33]; however, two issues must subsequently be addressed: (1) system speed, and (2) power/area. First, the speed of the imager may now be limited by the maximum speed of the converter, as one single converter must process every pixel's output serially. To get around this, a column-parallel ADC architecture can be employed, reducing the total amount of time needed to process a single image (or frame) [12], [14], [18], [22]. Taking this to a further extreme, ADCs can be shared among small groups of pixels [3], or there can even be one ADC per pixel [1]. However, adding an ADC increases the power consumed by the system, and also increases the necessary chip area. Since power consumption is the primary concern of the low-power system desired in this project, a single low-power ADC (as described in [38]) is used for the entire array.

2.4 Architectural Selection Summary

In summary, the following architecture is chosen for this system:

- **Photo-detector**: A p/n junction diode is selected over a photogate device since the process to be used does not provide sufficiently thin polysilicon gates.
- **Pixel topology**: CMOS Active Pixels are chosen to provide simple on-chip integration with other electronics, and to accommodate a low-power design, while providing adequate per-

formance.

- **Control scheme**: A rolling-shutter is chosen over a global shutter since the desired small pixel size does not provide enough capacitance to hold charge for an entire frame's read-out.
- **ADC**: For low-power design, a single ADC is to be used to serially read out all the pixel data, at the expense of frame rate.

Chapter 3

Figures of Merit

Before exploring the design of a CMOS camera system, it is important to describe some figures of merit by which any camera's performance can be measured and different cameras compared. The figures presented here are mostly relevant for a single pixel; however, since only one pixel is being read at any time during the camera's operation, this is gives a good estimate of the imager's overall performance metrics.

3.1 Noise and Dynamic Range

As a generalized concept in circuit design, noise is a random small signal occurring in devices due to the discrete nature of electronic charges. It often occurs as a function of the ambient temperature (thermal noise), and of the device manufacturing process (flicker noise), and limits the smallest signal that a circuit can process reliably [45], [47]. In practice, any time a signal is read through a circuit, noise is introduced, and a CMOS digital camera is no different. But in an imager, it is performance, particularly in low-lighting conditions, that is most affected by noise.

In a CMOS active-pixel cell, the dominant noise sources are reset noise, in-pixel amplifier noise, and readout noise. For the entire camera system, though, the eventual image data may come to be affected more by the quantization noise (among other sources) in the ADC; however, that will be a secondary consideration for now.

3.1.1 Reset Noise

During pixel reset, the photodiode sense node (hereafter known as the photonode) is pre-charged to (1) ensure that the photodiode is reverse-biased, and (2) effectively erase the previous-frame pixel value. The reset transistor switch is closed and operates in sub-threshold (though it may start out in saturation, depending on the previous integration) [11], during which the dominant noise sources are the reset transistor's shot noise and the photodiode shot noise. The reset transistor shot noise is given by [44] as

$$\overline{i_n^2} = kTC_{ph} \tag{3.1}$$

and diode shot noise by [47] as

$$\overline{i_d^2} = 2q\left(i_{dc} + i_{ph}\right)\Delta f \tag{3.2}$$

where k is the Boltzmann constant, i_{dc} is the dark current (current due to thermal charge generation, even in the absence of incident light), i_{ph} is the photocurrent, T is the temperature, and C_{ph} is the photonode capacitance. Including photodiode shot noise, the following reset noise figures are given in [11]:

If steady-state is achieved during reset, the reset noise $\overline{V_n^2}$ (referred to the photonode) is (kT/C_{ph}). However, in many cases, steady-state is not reached, and further detailed analysis gives a total mean squared noise at the end of reset of

$$\overline{V_n^2\left(t_r\right)} = \frac{1}{2}\frac{kT}{C_{ph}}\left(1 - \frac{\delta^2}{\left(t_r - t_1 + \delta\right)^2}\right), \tag{3.3}$$

where t_r is the reset time, t_1 is the rise/fall time of the reset pulse (depending on polarity), and δ is the time required to charge the photonode capacitance to the thermal voltage (v_T).

3.1.2 Integration/Exposure Noise

During exposure, noise is dominated by photodiode shot noise (again, due to dark current and photocurrent). Assuming a constant photodiode capacitance, the noise is often quoted as

$$\overline{V_n^2\left(t_{int}\right)} = \frac{q\left(i_{ph} + i_{dc}\right)}{C_{ph}^2}t_{int}. \tag{3.4}$$

However, the photonode voltage does not remain constant during integration; instead, it drops proportional to the photocurrent. Further analysis (as given in [11]) gives a mean squared noise value (referred to the photonode) at the end of integration of

$$\overline{V_n^2\left(t_{int}\right)} = \frac{q\left(i_{ph} + i_{dc}\right)}{C_{ph}^2\left(v_{ph}(0)\right)}t_{int} \cdot \left(1 - \frac{1}{2\left(v_{ph}(0) + \phi\right)}\frac{i_{ph} + i_{dc}}{C_{ph}\left(v_{ph}(0)\right)}\right)^2, \tag{3.5}$$

where $v_{ph}(0)$ is the photonode voltage at the beginning of the integration period (in this case, it's the reset voltage), ϕ is the built-in junction potential, t_{int} is the integration time, and all other variables are as described previously.

3.1.3 Source-Follower Readout Noise

During readout, a pixel operates like a single-stage source-follower amplifier, as illustrated in Figure 3.1. This circuit can be modeled by the small-signal equivalent given in Figure 3.2. For this

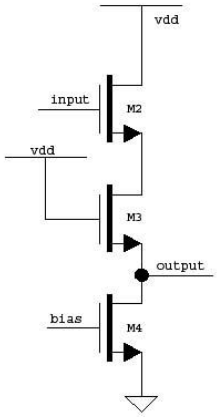

Figure 3.1: Equivalent schematic diagram of a pixel in readout mode. During readout, the signal path is basically just a single-NMOS source follower device, and so the noise contributors are the amplifying device (M2), the row-select device (M3 - which acts like a resistor), and the bias/load device (M4).

figure, (applying standard noise analysis as in [45] and [47]):

$$\overline{i_2^2} = 4kT \cdot \gamma \cdot g_{m2} \tag{3.6}$$

$$g_{d3} = \frac{1}{k'_n \frac{W}{L} V_{dsat}} \tag{3.7}$$

$$\overline{v_3^2} = \frac{4kT}{g_{d3}}, \tag{3.8}$$

with $\gamma = 2/3$.

Based on this, further analysis [11] gives the following noise contributions from each transistor, referred to the pixel's output node:

$$\overline{V_{n,sf}^2} = \frac{2}{3}\frac{kT}{C_o}\frac{1}{1+\frac{g_{m,sf}}{g_{d,rs}}} \tag{3.9}$$

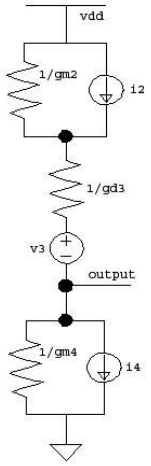

Figure 3.2: Small-signal model if the schematic in Figure 3.1 used to calculate readout noise. $1/f$ (flicker) noise is ignored in the analysis; otherwise, the noise sources are as described in [45].

$$\overline{V_{n,rs}^2} = \frac{kT}{C_o} \frac{1}{g_{d,rs}\left(\frac{1}{g_{d,rs}} + \frac{1}{g_{m,sf}}\right)} \tag{3.10}$$

$$\overline{V_{n,b}^2} = \frac{2}{3}\frac{kT}{C_o} g_{m,b}\left(\frac{1}{g_{d,rs}} + \frac{1}{g_{m,sf}}\right) \tag{3.11}$$

3.1.4 Optical Dynamic Range

Dynamic range is the ratio of the maximum signal power to the maximum noise power. In this imager, the "signal" to be concerned with is photodiode current since it is directly affected by light intensity (which is what is really of interest). Mathematically [39],

$$DR = 20\log\frac{i_{sat}}{i_{dc}} \tag{3.12}$$

where i_{sat} is the saturation current – the smallest current that causes the pixel reading to saturate, and i_{dc} is the dark current in the photodiode. In a standard CMOS process, the photo-detector dynamic range is expected to be around "5 to 7 decades of incident light intensity" [43], or 100-140dB.

3.2 Quantum Efficiency & Responsivity

The quantum efficiency (QE) of a photodiode is the fraction of incident photons on a photosensitive device that create charges (usually given in electrons/photon), or mathematically:

$$QE = \frac{\text{number of detected charges}}{\text{number of incident photons}}. \tag{3.13}$$

A sensor's QE usually varies with the wavelength of incident light [43].

A photodiode's responsivity is a similar metric, but is often given in A/W instead of electrons/photon. It too varies with wavelength, and can usually be derived from the QE (and vice versa).

3.3 Scaling Issues/Trends

As CMOS technologies advance and feature sizes shrink, the resulting changes in the preceding figures of merit often put into question the feasibility of image sensors in the first place. First, however, there are a number of advantages to be gained with scaling, including increased chip density, and possible speed increases. With regard to CMOS imagers, the scaling of a standard CMOS process could provide for an increased pixel resolution at a given imager size [43]. In general, however, the biggest gains from scaling come in the form of increased pixel fill factor and increased signal-processing electronics per pixel [34].

On the other hand, spectral responsivity and sensitivity worsen as geometries scale at the 0.25μm technology and beyond. First, the response to longer-wavelength (red) light decreases as junction depths decrease and implants increase [43], [42]. If junction depths get small enough, long-wavelength light will penetrate into the bulk beyond the depletion region and generate charges in a quasi-neutral region of the p/n junction. While this will still contribute to photocurrent, it is not nearly as efficient as depletion-region charge-generation. In addition, with increases in substrate doping, carrier mobility is reduced, further affecting the quantum efficiency of a photosensor [34]. Second, the increased off-current of (in-pixel) transistors due to thinner gate oxides, and the increased transistor junction leakage introduced with the use of shallow trench isolation and salicides will eventually exceed the dark current of the photodiode and dominate the pixel's dark current, thereby reducing the optical dynamic range [34], [43], [42]. Even though standard CMOS technologies "provide adequate imaging performance at the 2-1 μm generation without *any* process change, some modifications to the fabrication process and innovations of the pixel architecture are

needed to enable...good quality imaging at the 0.5μm technology generation and beyond," and in short, it is reasonable to expect that CMOS imagers can be scaled down to 0.25μm-0.18μm geometries if they "are willing to depart from 'standard' CMOS technologies by tailoring the junction and/or channel implants and...removing the silicide module" [34].

Chapter 4

System Design

This chapter details the design of the CMOS camera system. First, an overview of the entire system is given, followed by an explanation of the design of its individual parts. Finally, some hand calculations are included to predict the electrical and optical performance of the designed imager.

4.1 The Big Picture

System Constraints:

- Supply Voltage: 1V
- Power Consumption: $\approx$50 μW static power
- Maximum Pixel Response Time: 1.28ms. In the rolling-shutter architecture chosen, a pixel's sample will be changed at fastest, every 1.28ms when operating in short-exposure at 100kSample/sec.
- Imager Resolution: 128$\times$128
- Pixel Size: 8μm $\times$ 8μm (total imager array size is 1.024mm $\times$ 1.024mm)
- Frame Rate: limited by 100kS/s ADC to a maximum of 6 fps (assuming only one ADC processes all the pixels serially)

A high-level block diagram of the system to be designed is given in Figure 4.1, and is similar to the systems described in [7] and [32]. In addition, a detailed signal path from photon to ADC in

this system is given in Figure 4.2. As shown, four components must be designed: (1) a pixel array, (2) column circuitry, (3) an analog multiplexor, and (4) control logic.

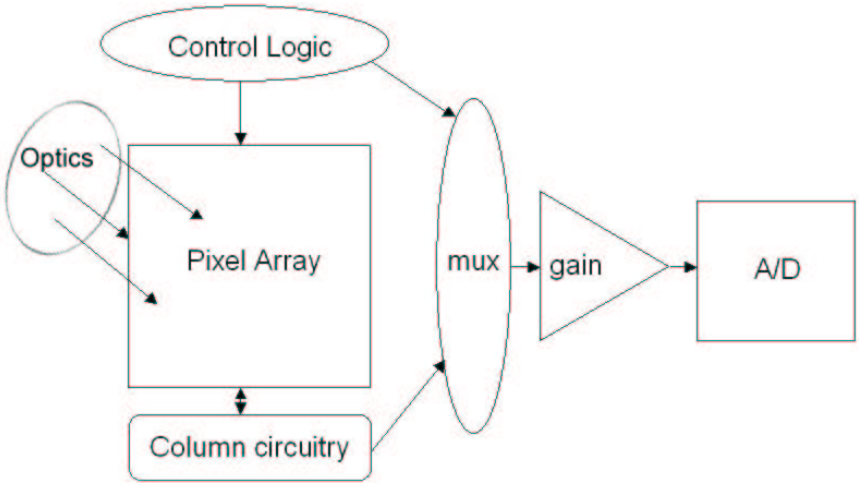

Figure 4.1: High-level Block Diagram of Camera System, showing the main functional blocks: (1) The Pixel Array, containing a 128×128 array of APS cells; (2) Column Circuitry, containing biasing for each column of the pixel array, sample-and-hold switches, and capacitors for each column; (3) Analog MUX, used to select one column's output from the single activated row; and (4) Control logic to cycle through the addressing of the rows and columns in the pixel array and analog mux. The other two blocks – a gain stage, followed by an ADC – are designed elsewhere and are not emphasized in this report.

4.2 Pixel Array

The imager is built as a rectangular array of pixel cells, each of which consists of a photodiode, in-pixel source follower, and reset and row-select switches.

4.2.1 The Pixel

The CMOS APS pixel schematic chosen for this design (and depicted in Figure 4.3) is the photodiode-based pixel cell described in [7], [9], [11], [13], [22], [28], [43], and [40]. It consists of a p-substrate/n^+-diffusion region as the photodiode and three transistors to be used for reset, amplifier, and row-select. In normal operation, the pixel acts in two different modes: reset and exposure.

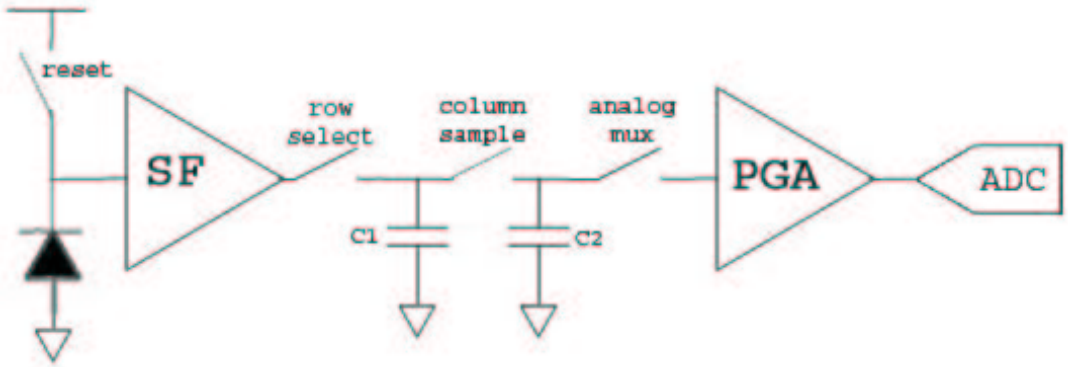

Figure 4.2: Signal path from photon to ADC. Incident light (photons) strike the photodiode, causing a voltage change that is amplified through the in-pixel source-follower (SF). The row-select switch, when closed, allows this source-follower to charge the column bus C1, and the column-select switch will sample this voltage onto the column capacitance C2. From there, the signal passes through an analog multiplexor to a programmable-gain amplifier which drives the analog-to-digital converter.

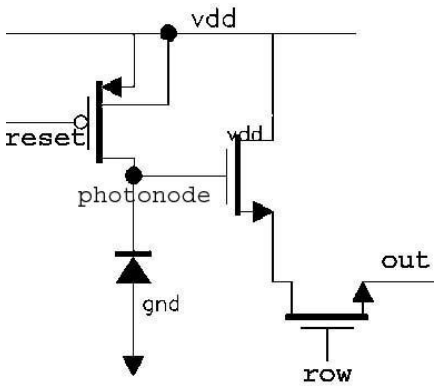

Figure 4.3: Schematic of a single 3-transistor active pixel sensor (APS) cell. The photosensitive element is a reverse-biased p-substrate/n^+-diffusion diode connected between the photonode and ground. Operation of the pixel is depicted later in Figures 4.4 and 4.5.

Reset Mode

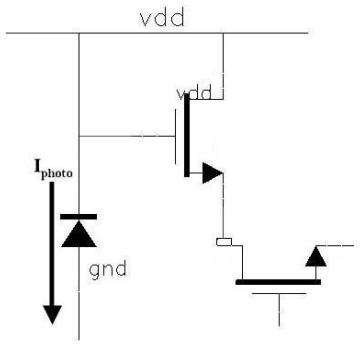

Figure 4.4: Schematic of a single pixel in reset mode. The reset transistor is effectively a short, thus charging the photonode to V_{dd}. In this state, the shutter of the camera is essentially off, since regardless of the intensity of incident light, the pixel produces the same output voltage. Photocurrent (I_{photo}) flows, but has no effect on the pixel output as long as the photonode is held at V_{dd}.

In reset mode, as shown in Figure 4.4, the internal capacitance of the photonode is charged to V_{dd}, effectively "closing the shutter" on the pixel. Though incident light on the pixel induces photocurrent, the photonode voltage remains constant at V_{dd}.

Exposure Mode

After the reset period, the pixel is exposed by turning off the reset transistor, effectively releasing the photonode from V_{dd}. Photocurrent then discharges the photonode.

Pixel Design Methodology

Since the primary concern of this camera is power consumption, the pixel design begins with the 50μW system specification. At a supply of 1 volt (approximately two solar cells), this corresponds to 400nA of bias current for each of 128 columns, and hence, for each pixel. In addition, the pixel should have a -3dB bandwidth of at least 1kHz to match the longest amount of time between pixel samples (namely, one row period, or 128 samples at 100kSample/sec). For a safety margin, the pixel will be designed for a 100kHz -3dB bandwidth.

Given these constraints, design of the single pixel follows in two steps: (1) design of the in-pixel source-follower circuit, and (2) sizing of reset and row-select switches. The size of the

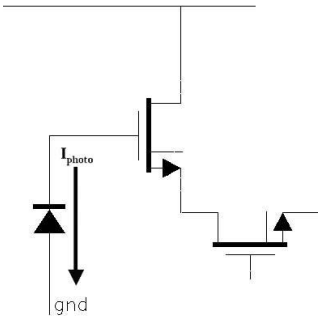

Figure 4.5: Schematic of a single pixel during exposure (during which the reset transistor is an open switch). The generated photocurrent I_{photo} will discharge the photonode (previously charged to V_{dd} during reset). Correspondingly, the photonode voltage drops, and the pixel output voltage follows accordingly.

photodiode will then be the remaining area available in the 64μm^2 pixel footprint after all other elements have been laid out.

4.2.2 In-Pixel Source Follower Design

The design of the in-pixel source follower applies when a pixel is exposed and its row-select signal is asserted. In this case, it boils down to the circuit shown in Figure 4.7, in which, C_{23} is composed of the diffusion-to-bulk capacitance of the active device (M2) and the row-select device (M3), and the gate-to-diffusion capacitance of M3. C_{34} is the drain-to-bulk capacitance of bias transistor (M4), and the total capacitance hanging off of the column wire, including 127 row-select devices that are "off" (modeled as diffusion-to-bulk of a row-select transistor). R_3 is given by the small-signal resistance of M3 in triode, or (including body effect):

$$R_3 = \frac{1}{k_n' \frac{W}{L} \left[V_{GS,3} - \left(V_{Tn0,3} + \gamma \left(\sqrt{V_{SB} + 2\phi} - \sqrt{2\phi} \right) \right) \right]}. \tag{4.1}$$

For preliminary calculations, it is assumed that all row-select transistors are sized at 0.44μm $\times$ 0.44μm for minimum NMOS leakage (as found in simulations). To first order, then, it can be shown that $C_{23} = 2.3fF$, $C_{34} = 108fF$, and $R_3 = 19.5k\Omega$ using numbers from the CMOS process' electrical design rules.

Small-signal, zero-value time-constant analysis of the circuit in Figure 4.7 yields an estimated

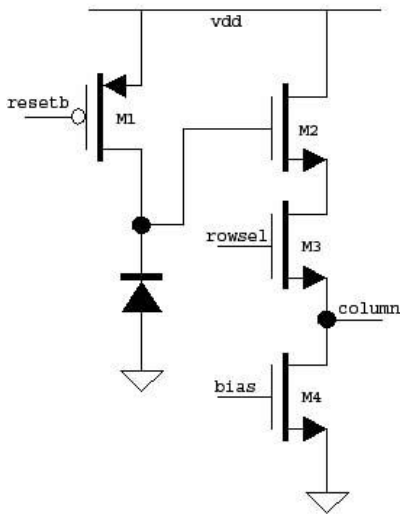

Figure 4.6: Schematic of basic pixel circuit including current source load (device M4). The main pieces to be designed are the in-pixel source-follower amplifier (device M2), and the two switches: reset (M1) – designed for minimum leakage during exposure; and row-select (M3) – designed so that its on-resistance does not degrade the source-follower performance.

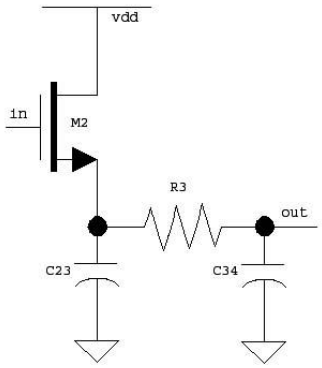

Figure 4.7: Schematic of equivalent pixel used for source-follower design. It is assumed that row-select is asserted (and thus device M3 becomes a resistance R3). Further, parasitic capacitances C23 (from the node between M2 and M3, and ground), and C34 (from the output/column) node are included for frequency response calculations.

-3dB frequency (neglecting body effect) of

$$f_{3dB} = \frac{g_{m2}}{2\pi * [C_{23} + C_{34}\,(1 + g_{m2}R_3)]}. \tag{4.2}$$

For a bias current of 400nA and a chosen V_{dsat} of 100mV (for output swing considerations), g_{m2} is found to be 7μS, and f_{3dB}, from equation 4.2, is > 50MHz which is more than enough for this application.

From this analysis, device M2 is sized with a W/L ratio of 0.471. Choosing $L = 2\mu$m for improved matching and linearity across multiple pixels, the channel width W is computed as 0.94μm.

4.2.3 Reset and Row-Select Switches

The PMOS transistor (M1 in Figure 4.6) used as a reset switch, and the NMOS transistor (M3 in Figure 4.6) used as a row-selection switch are designed to optimize pixel performance during exposure and readout, respectively.

Reset Device

First, a PMOS is used for the reset device since it can charge photonode all the way to V_{dd}, which is an important swing consideration given the low supply voltage. Next, two considerations for the reset switch are the speed of reset, and the off-current (leakage) during exposure. As mentioned in [11], a PMOS reset device provides sufficient reset speed, especially if the reset period is at least 1ms. Therefore, the off-current during exposure and its contribution to the total dark current is more of a concern than the reset speed. Preliminary simulations indicate that a minimum off-current (for $V_{SG} = V_{dd}$) occurs when $W/L = 0.54\mu$m$/0.54\mu$m, so this size is chosen.

Row-Select Device

Low readout on-resistance is the primary concern for the row-select device, as it is during readout that the pixel is active and drives the column wire through this switch; and because this resistance influences the pixel's output range and bandwidth. Leakage through an "off" row-select device (that could discharge a shared column wire) is only a secondary concern, since one pixel per column will always drive each column wire. In addition, charge injection from the row-select device will have little effect on the system output, since column-sampling occurs and de-couples

the value to be processed at the output before row-select is de-asserted (full timing details will be given later in this chapter). Therefore, for area considerations, the minimum length of 0.24μm is used; and for low on-resistance, a W/L ratio of 10 is chosen.

4.2.4 Layout Considerations

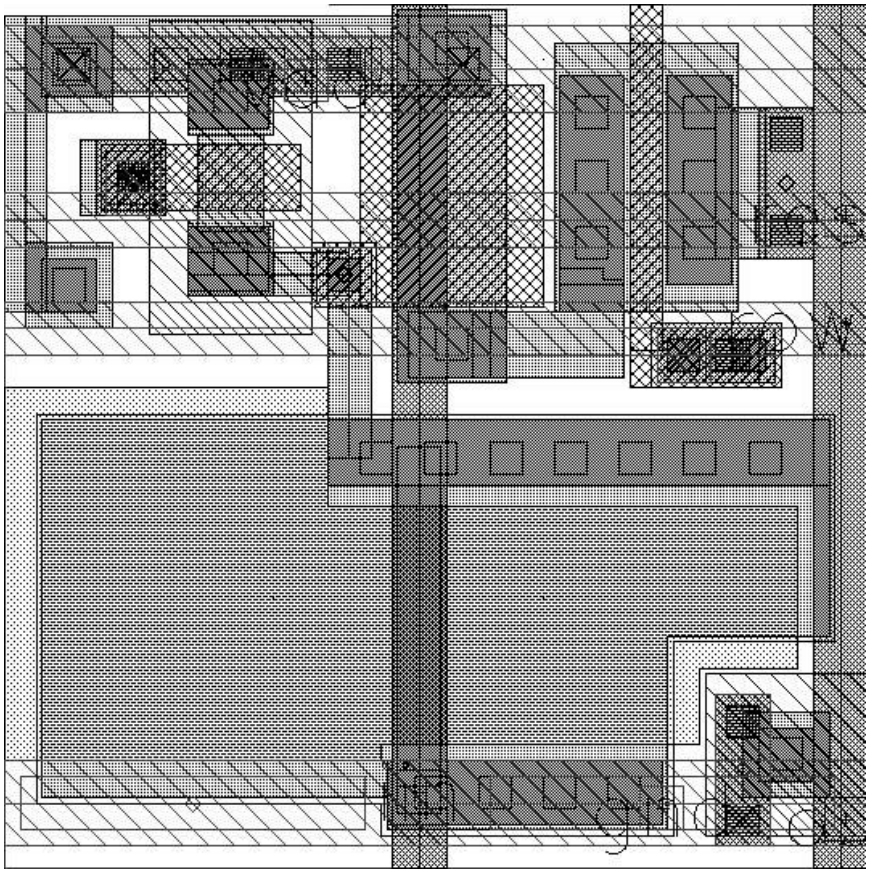

Figure 4.8: Layout of a single pixel as designed. Total pixel size is 8μm$\times$8μm, with a fill factor of 26.4%. Some fill-factor is given up for a substrate contact in the lower-right corner of each pixel to reduce the effects of blooming and cross-talk. All non-photodiode devices are covered by a grounded top-metal light shield (not shown in this picture).

As shown in the single-pixel layout of Figure 4.8, each pixel contains a substrate contact to reduce crosstalk between neighboring pixels [35], and shared reset, row-select, supply, and column lines are laid out for simple array construction. The photodiode is a n^+/p-substrate diode to provide an optimal combination of layout simplicity and performance [43], as compared previously in Table 2.1.

4.2.5 Array Configuration

The pixel array is a 128×128 rectangular array of the pixel cells described above. Pixels in a row share row-select and reset signals, and all pixels in a column share an output wire, current-sink load, and column-sample circuitry, as depicted in Figure 4.10. Once the array has been thus connected, it is similar to a conventional DRAM cell – namely, there are n row-select inputs (equivalent to word lines), and m column outputs (equivalent to bitlines), as illustrated conceptually in Figure 4.9. At this point, the array can be operated in either a snapshot or a rolling-shutter mode, depending only on the sequence of reset[0..n] and row-select[0..n] inputs.

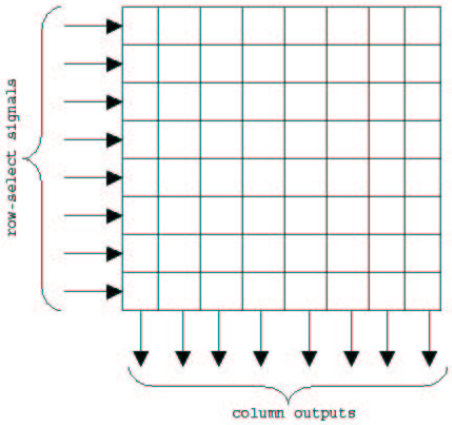

Figure 4.9: Block diagram of pixel array, shown for a sample 8-by-8 array. Inputs are the row-select signals, of which only one can be active at a time; outputs are the parallel columns, containing values determined by the selected row.

4.3 Column Circuits and Biasing

The column-common circuitry necessary to implement the rolling-shutter algorithm previously described consists of (1) a sample-and-hold switch/capacitor (to store a row's values for processing), and (2) a current-sink used to bias one row at a time. These elements are depicted in Figure 4.10.

4.3.1 Sample-and-Hold

A 1pF column capacitor is chosen for the sample and hold since it should be able to hold charge to within one-half of an LSB of the ADC (2mV) over the row period (128 sample-process times).

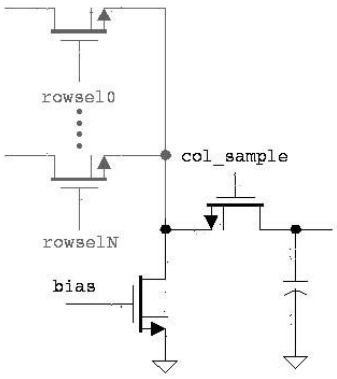

Figure 4.10: Schematic of column circuitry. Each column of the pixel array will share one current-sink load device, a sample-and-hold column-sample switch, and a column capacitor on which the selected row's output is sampled. Note that the architecture guarantees that only one row at a time is connected to the column circuitry.

Additionally, the sample switch is chosen to be a minimum-sized NMOS to minimize area and charge-injection.

4.3.2 Biasing

To properly bias the in-pixel source-followers of each active pixel, an appropriate bias voltage and current-sink load is needed for each column. And since only one row is ever connected to a column wire at a time, one current-sink per column is sufficient, as commonly shown in the literature.

For this design, a resistor and a diode-connected NMOS device (shown in Figure 4.11), are used to generate the bias voltage for the current-sinks. In such a current-mirror topology, power can be saved by scaling the width of the diode-connected FET relative to the current sink FETs that it will drive. However, since the desired bias current is a mere 400nA, a 1:1 ratio of current is chosen for the bias FET. Preliminary hand calculations show that in order to attain a 100mV V_{dsat} for the column loads, a W/L of 0.63 is needed. Correspondingly, the W/L of each column current-sink is chosen to be 1.23μm/2μm. Simulations show that a bias resistance (R_{bias}) of 1.3MΩ provides the right amount of bias current.

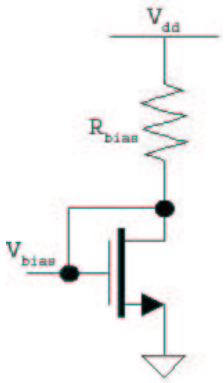

Figure 4.11: Schematic of the bias circuit used to generate the bias voltage for the column current sink loads. In order to provide the desired 400nA of current, R_{bias} is chosen at 1.3MΩ, and W/L of the FET is chosen to be 1.23/2. While this may not the most robust configuration for a bias network, it nevertheless will suffice for this application.

4.4 Digital Control Circuits

To realize a rolling-shutter pixel array, a number of external control circuits are needed, namely (1) row decoder, (2) analog multiplexor, (3) binary counter to drive the decoder and multiplexor, and optionally, (4) digital logic to generate control signals for the post-processing units.

In any rolling shutter architecture, there are two possible row exposure timings. In one case (long exposure), all row-select signals are tied to the previous row's reset signal. In a 128-row array, one row is reset for the row-period immediately after it is selected and sampled, and subsequently exposed for the remaining 127 row periods, as described earlier in Figure 2.7 and shown in detail in Figure 4.12. Alternatively, the exposure period can be limited to a single row selection period for a shorter exposure time (and longer reset time) by setting a row's reset signal equal to the inverse of its row-select signal (logically: the row is being reset when it is not selected). This is shown in Figure 4.13

4.4.1 Row Decoder

In either of the timing schemes above, a decoder generates the row-select signals, and a 7-bit binary counter addresses this decoder. Each row can then be reset/selected according to whichever exposure scheme is desired. For this system, a 7-to-128 bit decoder (that outputs both the 128 select bits and their complements) is implemented using digital logic gates from the standard cells

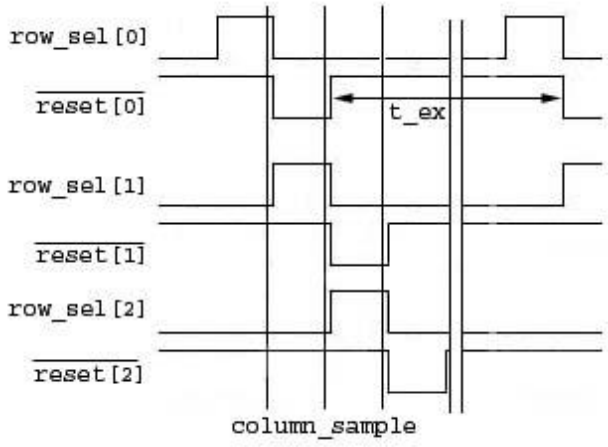

Figure 4.12: Reset and Row-Select signal timing for long-exposure scheme. Since a PMOS device is used for the reset switch in my pixels, $\overline{\text{Reset}}$ signals are depicted here. As indicated, the reset PFETs can be driven by the inverse of the next row's row-select signal (or in other words, one row's select signal is tied to the previous row's reset signal). The exposure time t_{ex} is shown to be 127 row periods.

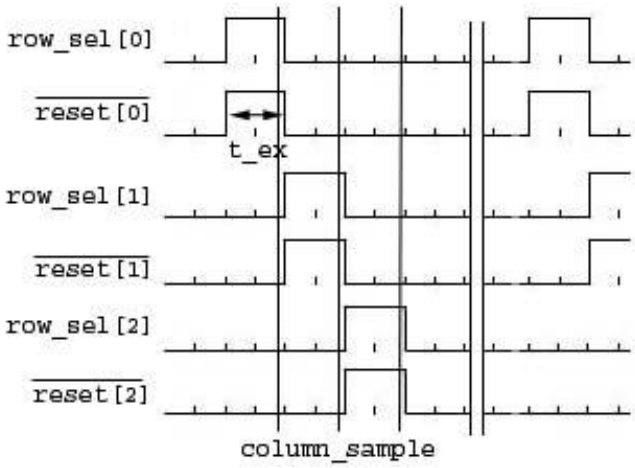

Figure 4.13: Reset and Row-Select signal timing for short-exposure scheme. Since a PMOS device is used for the reset switch in my pixels, $\overline{\text{Reset}}$ signals are depicted here. As indicated, the reset PFETs can be driven by the exact same signals that drive the row-select devices, giving an exposure time t_{ex} equal to one row period.

library provided with National Semiconductor's CMOS8 process.

4.4.2 Analog Multiplexor

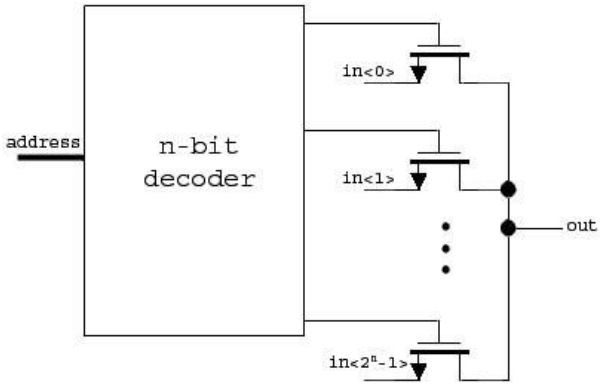

Figure 4.14: Schematic of analog multiplexor. In this system, a 7-bit address drives a 7-to-128-bit decoder, which then activates one pass-gate for output.

At the output of the pixel array, a 128-to-1 analog multiplexor selects one pixel output at a time from the row that is currently being processed. Shown schematically in Figure 4.14, the analog multiplexor consists of a 7-to-128 decoder and 128 pass-gates. The same decoder designed to drive the rows is used, and each of the pass-gates is a minimum-sized complementary pass-gate to provide full-rail voltage swing. The 128 samples per row can be accessed sequentially during readout by driving the address bits with a 7-bit binary counter. To achieve this, the same counter used for the row decoder is expanded to fourteen bits to also generate column address signals for the analog multiplexor – the seven least-significant bits will drive the analog multiplexor, and the seven top bits the row decoder.

Charge Sharing Through Analog Multiplexor

Unfortunately, the analog multiplexor designed above (using pass gates, with capacitive loads on both sides) introduces a potential charge-sharing problem that causes unwanted image lag among neighboring pixel samples. This is largely an artifact of the readout timing (see Figure 4.2) and is described and illustrated in Figure 4.15.

Using Figure 4.15 as an example, two neighboring pixels in the same row drive the output circuitry in the following sequence: **(A)** Column-Sample signal at the end of the row's exposure

period causes column capacitance $C_{2,n}$ to be charged up to the pixel's output voltage V_n. **(B)** The next row's row-select signal is asserted, disconnecting this row from the columns; however, the charges acquired in (A) remain on the column capacitors $C_{2,i}$. Meanwhile, the column wire capacitors $C_{1,i}$ are driven to new values based on the next row's samples. **(C)** The first column in this row is selected, by connecting $C_{2,n}$ to the implicit capacitance on the output node of the analog mux. Charge sharing between the output-node capacitance C_3 and the $C_{2,n}$ results in a V_n^* value that is dependent on the node's previous value of V_{n-1}^*! **(D)** The second column is selected for output, causing even more charge sharing; in fact, the value at the output of the analog mux now is a function of both the first and second column's voltage value (V_{n+1}^* is now a function of V_n^* too) – wherein lies the problem: if the third column's light intensity is identical to that of the second column, but different than the first, then the readings for columns 2 and 3 will be different, even though they have the same light intensity!

First-order hand calculations place the potential error due to this charge sharing anywhere up to 40mV (>10 LSB), which is unacceptable, especially since this error will be potentially doubled through the gain stage before reaching the ADC. Fortunately, two circuit techniques can eliminate charge sharing: (1) buffer each input to the analog mux or (2) periodically reset (or pre-charge) the output node of the analog mux after a sample's output has been read, and before the next sample is selected. The second of these two options is more attractive, as it will consume less area (only one reset switch for the entire mux is needed, plus digital circuitry to generate the control pulse), and power (no new active components drawing static current).

4.4.3 Binary Counter

The previously-mentioned 14-bit counter is implemented on-chip as a ripple-counter with resynchronization, shown conceptually in Figure 4.16.

4.4.4 Post-Processing Controls

Two digital control signals to drive the programmable-gain amplifier are needed: (1) a reset signal, and (2) a track signal, the details of which are beyond the scope of this report, as the amplifier is designed by Al Molnar and has been tested independently of this project. However, to generate these pulses for each sample, the on-chip counter for this system is further expanded to sixteen bits: the upper fourteen bits drive the row and columns of the array as before, and now the lowest

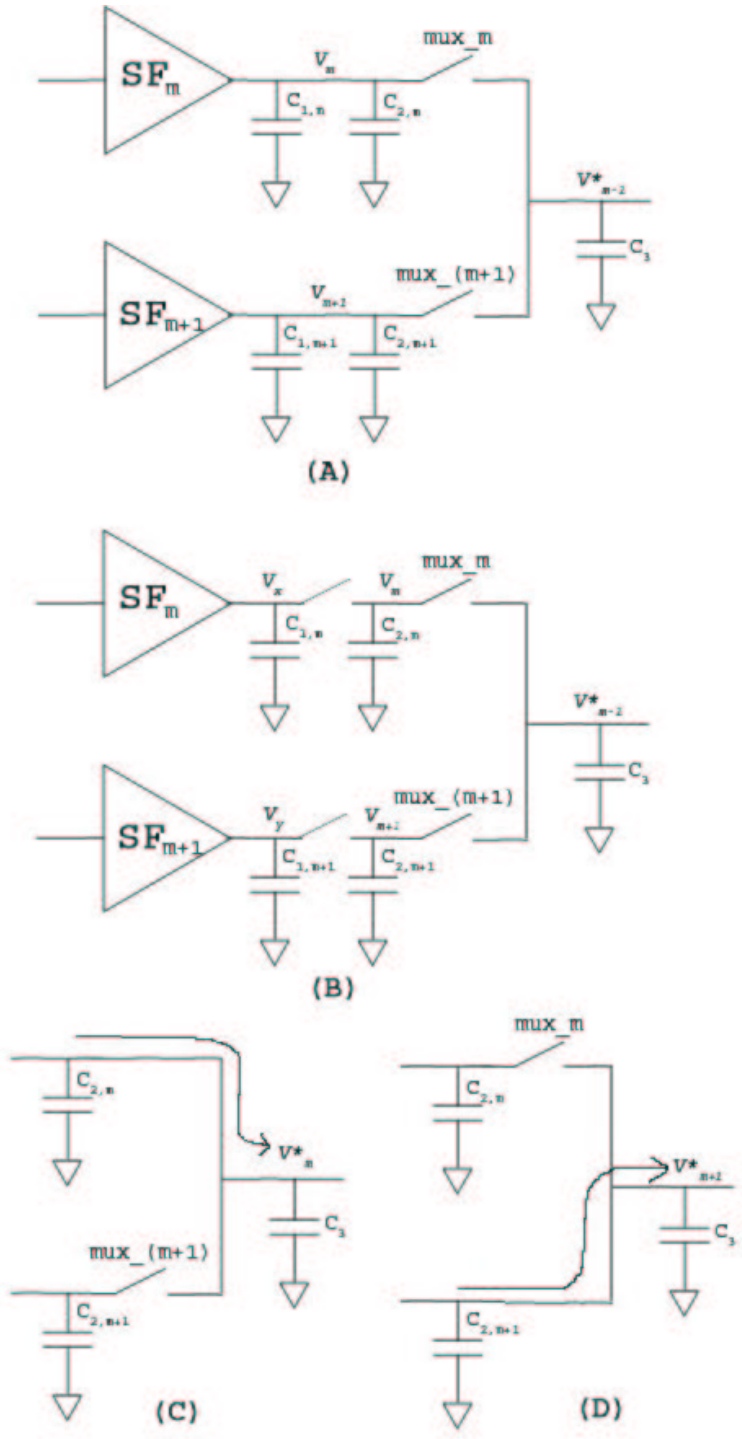

Figure 4.15: Illustration of charge-sharing problem through the analog mux.

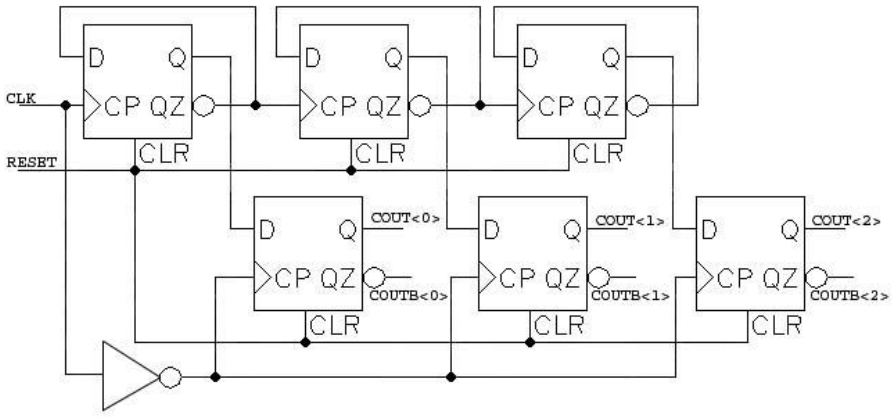

Figure 4.16: Schematic of binary ripple-counter with resynchronization. A 3-bit counter is shown here just for conceptualization, but this idea can be extended easily to a larger counter.

two are used to generate the PGA controls, the analog multiplexor's pre-charge signal (to avoid the charge-sharing described above), and the column-sample signal, as illustrated in Figure 4.17. In practice, since there are four clock cycles per pixel access, a clock frequency of $4N$ Hz is needed to achieve a sampling rate of N samples per second *(e.g., a system input clock of 400kHz is needed to achieve the maximum ADC sample rate of 100kS/s in this system)*.

4.5 Expected System Performance

Finally, it is useful to describe the expected electrical characteristics at each point along the signal path from photon to ADC. This section contains graphs of expected waveforms and summarizes the hand-calculated figures of merit for the designed camera system, providing a purely electrical basis for comparison when testing the fabricated chip.

4.5.1 Waveform Characteristics

Referring to Figure 4.18, it is expected that the waveform characteristics at each node will be:

- **At the source-follower output (A)**: The voltage level will hold constant during pixel reset (at around 400-500mV), and will slew with linear slope during exposure (or until the pixel saturates under intense light).

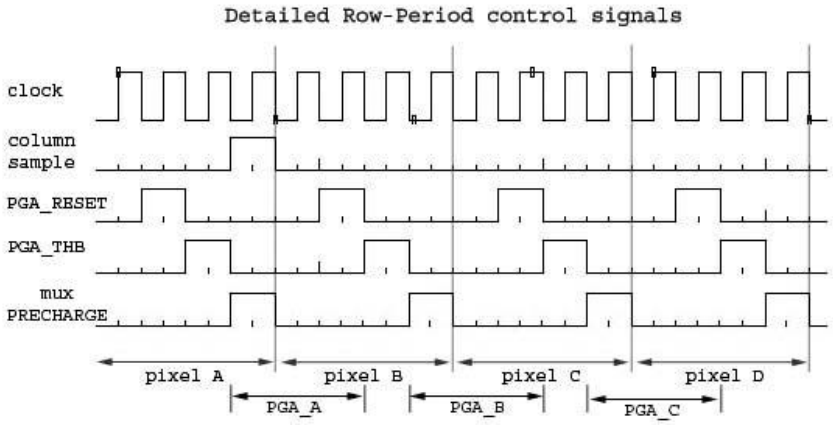

Figure 4.17: Pixel A, B, C, and D represent four pixels that straddle a row border (as determined by the location of the column-sample pulse). For example, pixel A could be the 128th (last) pixel in row number 5; after that, pixels B, C, and D would be the 1st, 2nd, and 3rd pixels in row number 6. The time periods denoted by PGA-X signify the times when the output of the programmable-gain amplifier would correspond to the value of pixel X (and thus it would then be safe for sampling at the input of an ADC).

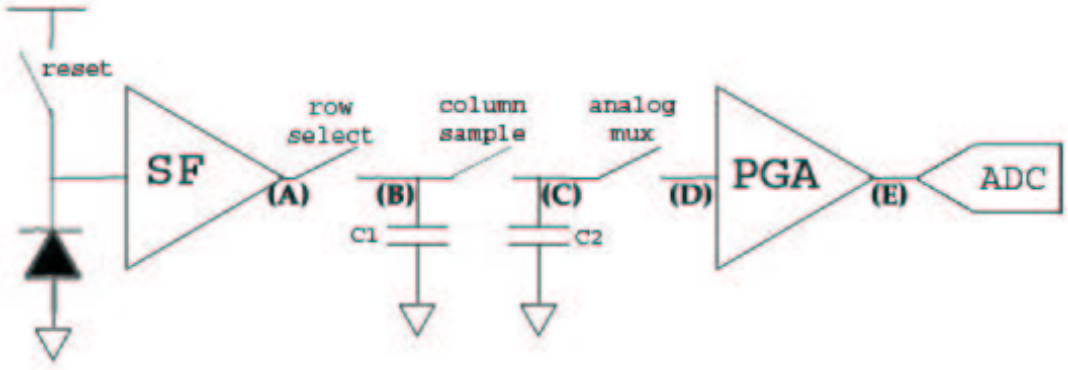

Figure 4.18: Signal path from photon to ADC. This is the same image as given in Figure 4.2, but with labeled nodes for reference in the expected waveform output graphs.

- **At the shared column output (B)**: Since the row-select signal is never asserted during reset, and the columns are connected to a row only while it is slewing, this should be a series of slewing voltages that looks like a sawtooth wave (assuming no saturation).

- **On the sampled column capacitors (C)**: The column voltage from **(B)** is sampled onto this node at the end of a row exposure period. This value will then be held constant until the analog multiplexor selects the column.

- **At the output of the analog multiplexor (D)**: This should be a series of "samples," each of which consists of pre-charge (from the charge-sharing avoidance scheme developed earlier in this chapter) for one quarter of the sample, and drops to an evaluate/hold for the rest of the period, at a voltage proportional to pixel light intensity.

- **At the output of the programmable-gain amplifier (E)**: The output of the gain stage should take two forms in sequence: (1) holding a steady value that corresponds to twice the value in **(D)** (when set to a gain of 2), and (2) unknown behavior when the output is left to float. The "floating" voltage is not a problem, however, as the ADC will be synchronized to sample the PGA output only when it is holding a steady value.

- **ADC Output (not shown graphically)**: Ideally, the ADC will sample the PGA output **(E)** while it is steady, and produce eight parallel output bits before the next steady PGA sample arrives to be sampled.

These are summarized visually in Figures 4.19 (A, B, and C) and 4.20 (D and E). To put these two figures in perspective, waveforms A-C are shown across a little over two row periods (256 samples); by contrast, waveforms D and E are shown across three individual samples.

4.5.2 Effects of Charge Redistribution

While the passive charge-sharing avoidance scheme developed previously will eliminate image lag, it also introduces a capacitive divider that reduces the analog output swing, as described in Figure 4.21. Given the previously-chosen 1pF column capacitance (C_1), and computing the analog multiplexor output's parasitic capacitance to be 213fF (C_2) based on the electrical design rules, it is found that an array-level analog output voltage range of 0-500mV will be compressed to 88-500mV via charge redistribution (assuming a pre-charge voltage V_{pre} of 500mV).

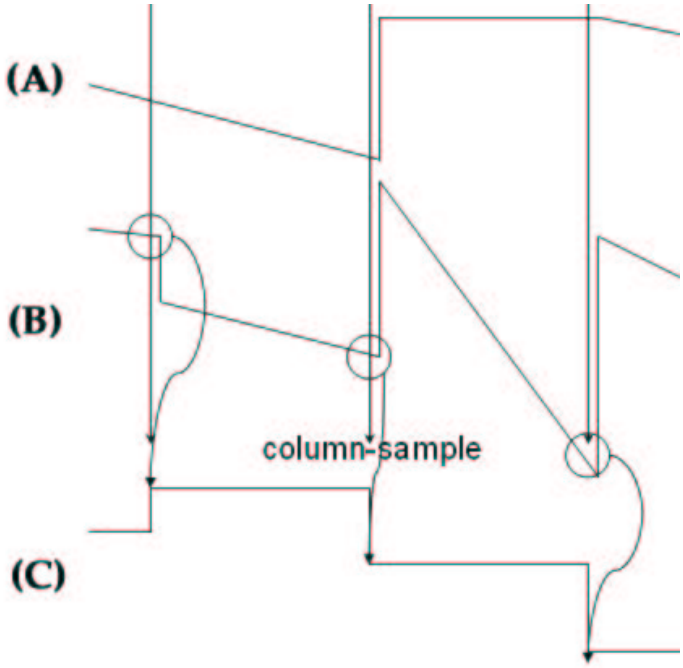

Figure 4.19: Expected system waveforms at points (A) - pixel output; (B) - column voltage; and (C) - sampled column voltage of Figure 4.18. Conceptually, (B) will track (A) depending on the one row-select signal that is asserted, and when the column-sample signal is pulsed, the value from (B) is sampled onto node (C).

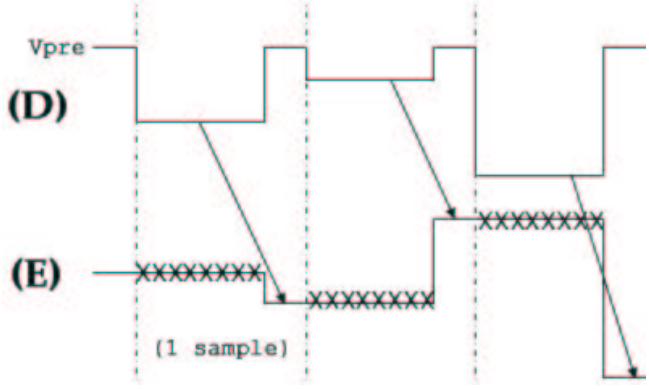

Figure 4.20: Expected system waveforms at points (D) - analog multieplexor output; and (E) - PGA output of Figure 4.18. Conceptually, the arrows link an analog output value from the analog mux to its amplified value out of the PGA. The “XXX” denotes a floating value at the output of the PGA that is meaningless.

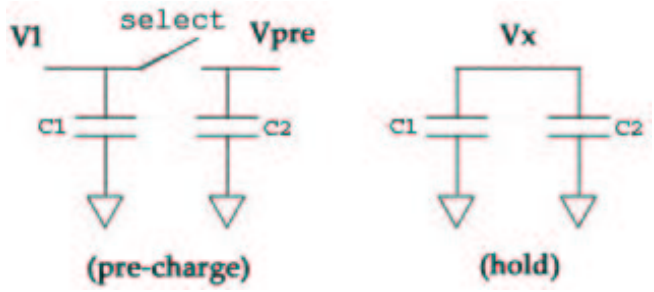

Figure 4.21: Illustration of charge redistribution and capacitive divider through analog multiplexor when using the charge-sharing avoidance scheme developed in this design. In this image, C_1 is the column capacitance, and C_2 the (parasitic) capacitance at the output of the analog multiplexor. On the left, during the sample's pre-charge period, C_1 is driven by an in-pixel source-follower to a voltage level V_1 proportional to a particular pixel's light intensity. At the same time, the analog mux output is pre-charged to V_{pre}. On the right, once the column is selected for read-out, the charge is redistributed over both capacitances, and the resulting voltage V_X read out by the circuit is given by $(C_1V_1 + C_2V_{pre})/(C_1 + C_2)$.

4.5.3 Total Integrated Noise

Based on the previous noise discussion and using equations 3.3; 3.5; and 3.9, 3.10, and 3.11 as derived for reset, integration, and total read-out noise, respectively, the following values are obtained for the output-referred rms noise in each mode of operation:

- **Reset Noise**: $v_{n,rst} = 224\mu$V rms
- **Integration Noise**: $v_{n,int} = 800\mu$V rms
- **Read-out Noise**: $v_{n,ro} = 82\mu$V rms
- Therefore, **Total Noise**: $v_{n,out} = \sqrt{v_{n,rst}^2 + v_{n,int}^2 + v_{n,ro}^2} = 835\mu$V rms

The parasitic capacitances used to calculate these figures are derived from the junction capacitance values given in the electrical design rules, giving $C_{photonode} \approx 40$fF

4.5.4 Figures of Merit

The following performance metrics are expected for the system as designed, and will be compared to the fabricated chip in testing:

- **Static Power Consumption (imager only)**: 51.6 μW
- **Analog Output Voltage Swing**: 412mV (88mV - 500mV)
- **PGA Output Swing**: 824mV (176mV - 1.0V)
- **Electronic Noise (at analog output of imager only)**: 835 μV rms
- **Optical dynamic range**: 20-50dB (depending on exposure time and the accuracy of leakage currents quoted in the electrical design rules)
- **Field of View**: 36°, assuming a 2mm-diameter lens with focal length of 1.5mm.

Chapter 5

Testing Results

This design is realized in National Semiconductor, Inc.'s 0.25-micron CMOS8 (2-poly, 5-metal) process. No imager-specific modifications were made to the process, so it is likely that the imager performance will suffer, as mentioned previously in the technology scaling discussion. Nevertheless, the chip should still be functional, and is tested in two phases: (1) a preliminary characterization phase, and (2) final system functionality testing. Initially, I hoped that one phase of testing would be adequate; however, some bugs were discovered during initial testing, necessitating another revision, tape-out, and ultimately a second round of testing. Therefore, the preliminary testing phase is only significant for the single pixel and basic array (column-level output) results, and system-level results are verified in the final round of testing. For both testing phases, the methodology flow diagram is given in Figure 5.1.

5.1 Preliminary Testing Phase

In the preliminary phase, two separate chips are fabricated. The first of these (hereafter known as TC1) contains small structures useful for pixel-level characterization and proof-of-concept. The second (hereafter known as TC2) contains the entire camera system; however, a layout wiring error – reversed control signals – was discovered to cause unsightly image lag. But nevertheless, TC2 is still useful in taking some initial measurements and was used to refine the testing flow.

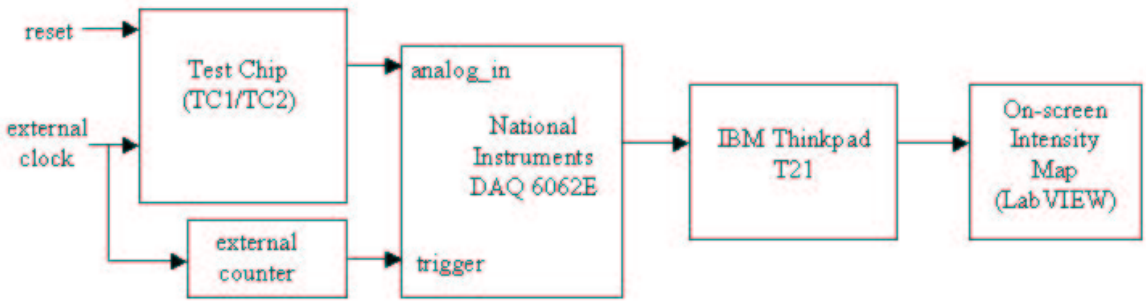

Figure 5.1: Flow diagram of the test setup. Note that external power supplies must also be provided and are not shown in this diagram. Also, the test chip output (plus trigger) can alternatively be sent to an oscilloscope instead of the DAQ card; the DAQ card is necessary only when processing voltage samples and converting them into a visual image. For the final testing phase, the external counters are not used, and the trigger signal is generated on-chip.

5.1.1 TC1

TC1 serves two purposes: (1) measurement of single-pixel characteristics, and (2) proof-of-concept of the rolling shutter architecture. Correspondingly, it contains two structures of interest: (1) a bare active pixel cell, whose photonode and output node are wired to pads for off-chip measurement, and (2) an 8×8 array of these pixels, which takes as inputs an external clock and a reset signal. On-chip digital circuitry generates the proper signals to traverse the rows of the 8×8 array, as described before. At the output, however, the array's eight columns are *not* passed through an analog multiplexor; instead, each column is buffered and wired to a pad for off-chip measurement. In all, there are four bare pixels, and two 8×8 arrays per TC1 die.

5.1.2 TC1 - Single Pixel

The following performance characteristics of a single pixel are measured:

- DC Biasing: Will the pixel turn on and provide the expected reset-level voltages?
- In-pixel Photodiode Characterization: how much photocurrent is generated under lighting conditions of interest, and what kind of responsivity does such a small photodiode have?
- Transient Response: Does the pixel behave as expected during alternating reset and expose cycles?

- Frequency Response: Does the in-pixel source-follower meet the bandwidth specification for which it was designed?
- Dynamic Range: What are the maximum and minimum signals that can be distinguished by a pixel?
- Power Consumption

DC Operating Point

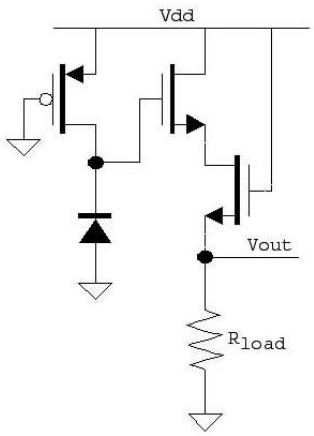

Figure 5.2: Schematic of test setup used to verify DC Operating Point and biasing. The situation of interest here is the pixel reset state, in which the reset switch is wired shut (figuratively). In addition, the row select transistor is also wired shut (and thus acts only as a resistance). Of this schematic, the photodiode and three transistors are included on-chip in the single pixel cell, but the resistor R_{load} is supplied externally for biasing.

To test the operating/bias point during pixel reset, a single pixel is configured as in Figure 5.2, with a 1.27MΩ external resistor serving as the load device. The pixel's reset levels are as expected: for a photonode voltage at 1.0V ($= V_{dd}$), the output is at 434mV.

In-Pixel Photodiode Dark Current

To measure the in-pixel photodiode's dark current, the pixel is connected as shown in Figure 5.3, and a piece of black electrical tape placed on the chip to block out light. By using Ohm's Law, the photocurrent generated can then be extracted by measuring the voltage across R_{test}. Doing so

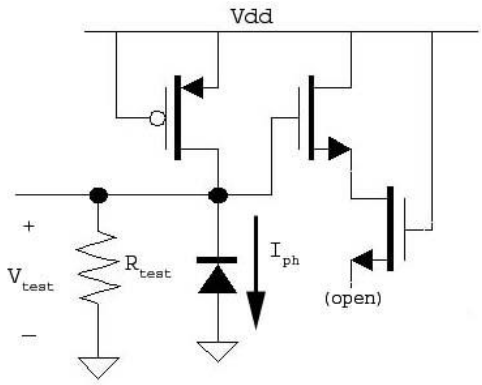

Figure 5.3: Schematic of test setup used to characterize the in-pixel photodiode. An external 2MΩ resistance R_{test} is connected in parallel to the photodiode, and the voltage across its terminals measured. From this, the photocurrent I_{ph} can be computed. The reset switch is kept open, and the source-follower disconnected to ensure that the photodiode behaves uninterrupted.

produces a dark (leakage) current of 1.4pA, which is orders of magnitude higher than the 26fA expected based on the process electrical design rules. However, it is possible that the pad diode connected to the photonode's layout pad is responsible for a large fraction of this leakage current. Considering that the pad is 17 times the area of an in-pixel photodiode, the dark current can be estimated at 1/18 of this, or 78fA.

In-Pixel Photodiode Transient Response

To verify proper behavior during exposure and reset states under various lighting conditions, the pixel is connected as in Figure 5.4. This is identical to the configuration of the DC operating point test, except that the reset device is driven by a clock, forcing the pixel to alternate between reset and expose states. Further, the distribution of the reset and expose time can be controlled by varying this clock's duty cycle.

As mentioned previously and shown in Figure 4.19, node (A), it is expected that the pixel output voltage will (1) hold steady at the DC bias value during reset, and (2) slew with a linear slope during exposure. The amount of slewing should be proportional to the incident light intensity – higher intensity should lead to more photocurrent, and therefore a higher slew rate. Test results in Figure 5.5 confirm that this is indeed the case.

In Figure 5.5, the three test conditions are:

- Darkness: the pixel is covered and the only current on the photosite should be the dark

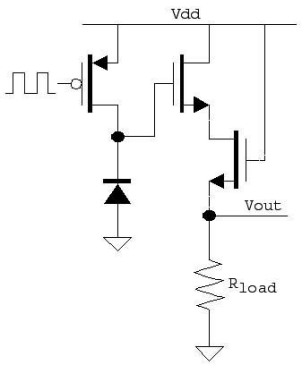

Figure 5.4: Schematic of test setup used to characterize a single pixel's transient response to light. The row-select device is forced on and acts like a resistor, but the reset signal is driven by a pulse between V_{dd} and ground, alternating between exposure and reset stages, respectively. The load resistor, R_{load} is an off-chip component connected to V_{out}. Everything else in this schematic is part of the unit pixel cell. For this testing condition, R_{load} is chosen at 1.27MΩ for a reset-state V_{out} of 434mV.

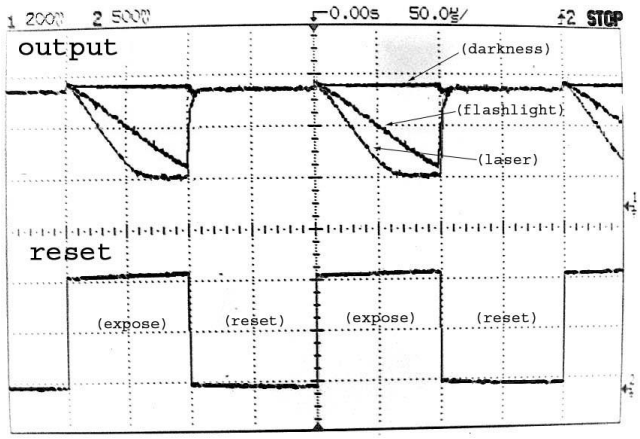

Figure 5.5: Pixel output voltage under varying incident light intensities, alternating between reset and exposure modes of operation. With a PMOS reset device, the pixel is reset by a low voltage and exposed when reset is high. As expected (Figre 4.19), the output voltage remains (relatively) constant in darkness, whereas in light, the slew rate is proportional to the light intensity. Also, this figure illustrates the effects of pixel saturation in the most intense incident light, as well as some charge injection.

current.

- Dim Flashlight: a flashlight is placed right against the chip should flood the bare pixel with light (no lens)
- Laser: a red laser (λ =630-680nm) is aimed onto the bare pixel (again, no lens)

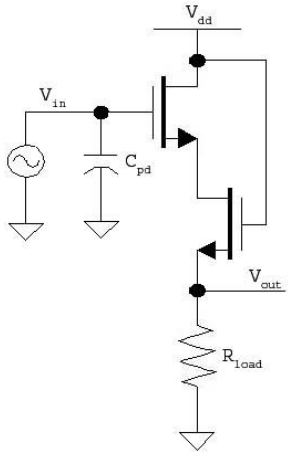

Figure 5.6: Schematic showing the effective circuit used to verify the transient response to sinusoidal excitation and frequency response. In this case, the reset signal is tied to V_{dd} to shut off the reset device. As before, R_{load} is a 1.27 MΩ off-chip resistor.

Transient & Frequency Response of In-Pixel Source Follower

Next, the in-pixel source-follower's transient response is verified by turning off the reset switch and driving the photonode with a sine wave, as illustrated in Figure 5.6. With this test, the output clipping limits can be obtained. For each of these tests, a sine wave at 1kHz is used since this is approximately the highest frequency at which a pixel in the 128$\times$128 array will be sampled, and the resulting waveforms are shown in Figures 5.7 and 5.8.

For even further verification, the in-pixel source-follower's AC response is measured and compared with the expected characteristic from hand calculations and simulation. The test setup shown in Figure 5.6 is re-used, except now, V_{in} is fixed at a constant 800mV $\pm$ 200mV (covering the range of interesting photonode voltage levels), and V_{out} is measured on an oscilloscope for magnitude and phase data at varying frequencies. The results are summarized in Figures 5.9 and 5.10 for magnitude and phase, respectively.

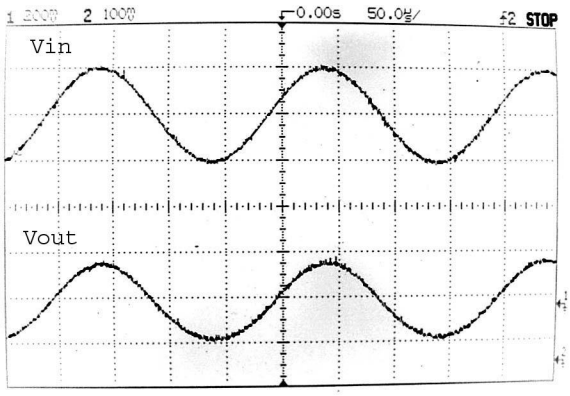

Figure 5.7: Relatively clean sine wave output through the in-pixel source follower. V_{in} is 840mV $\pm$ 215mV, and V_{out} is 140mV $\pm$ 90mV (though some distortion is evident).

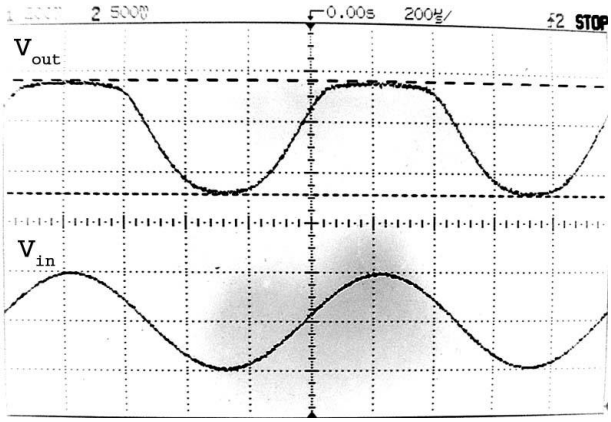

Figure 5.8: Oscilloscope output showing output clipping limits (in V_{out}). In this case, a V_{in} of 800mV $\pm$1V is used to drive the output to upper and lower limits of 450mV and 10mV, respectively.

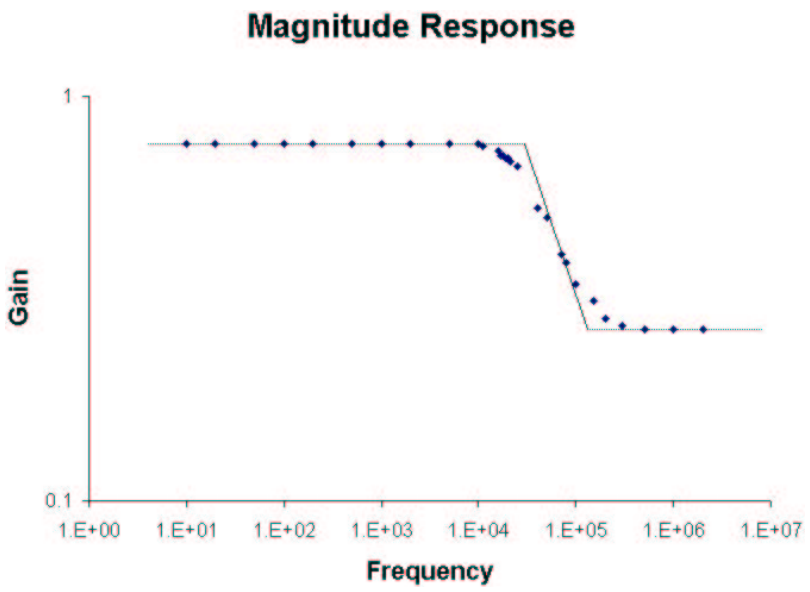

Figure 5.9: Bode magnitude plot of in-pixel source-follower gain, with straight-line approximations superimposed on the data. There appears to be a pole near 40kHz and a zero near 100kHz, with a low-frequency gain of 0.77.

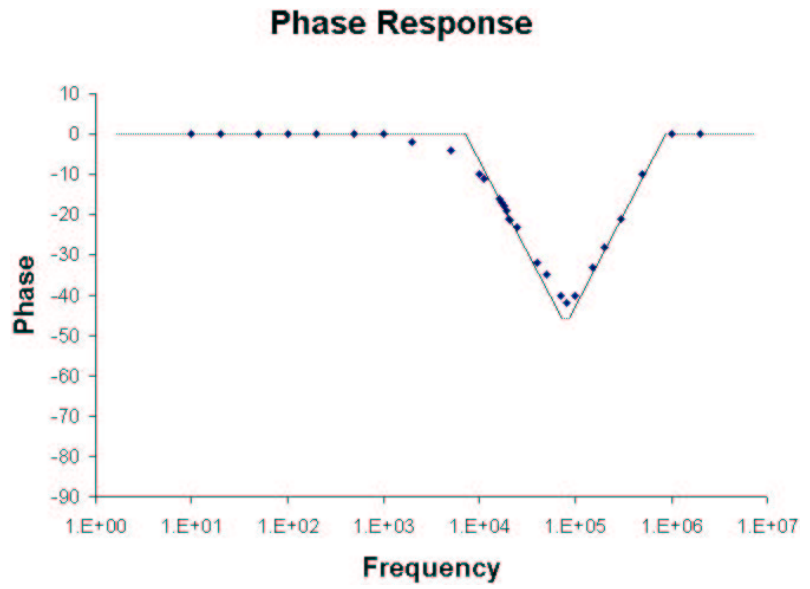

Figure 5.10: Bode phase plot of in-pixel source-follower gain, under the same conditions of the magnitude plot above. Straight-line approximations are also included on top of the individual measured data points, giving approximate pole and zero locations that agree with the estimated values from the magnitude plot.

Although low-frequency gain agrees rather well with hand-calculations and simulations (0.77 vs. 0.81), the dominant pole does not: hand-calculations of this design place the pole well above 10MHz, yet the measured pole is much worse than this. This difference can be explained by considering the extra output capacitance incurred during measurement – a 1-2pF (estimated) packaging capacitance combined with a 11pF scope probe capacitance contribute to the poor frequency response. Subsequent simulations that include a 12pF output capacitance produce the AC output shown in Figure 5.11, which agrees with the measured data.

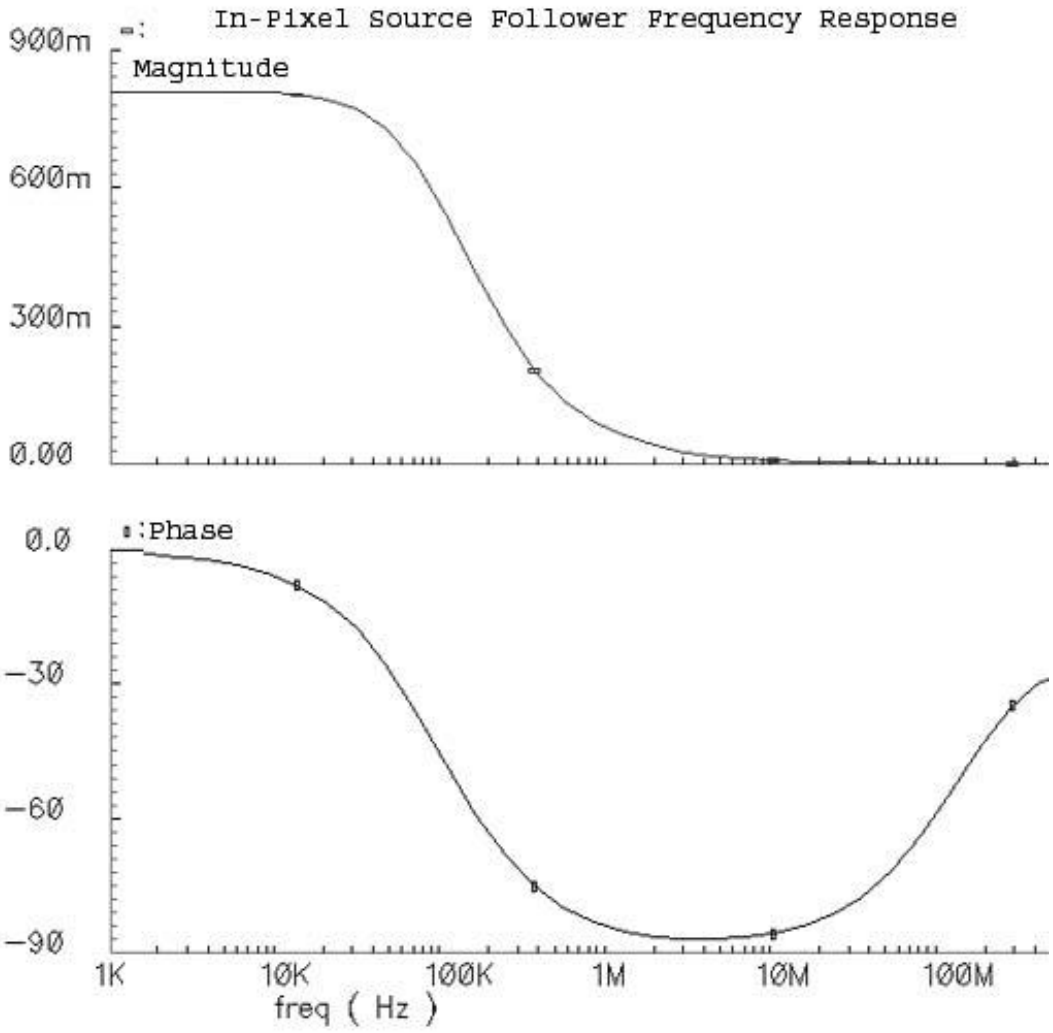

Figure 5.11: Simulated AC response of in-pixel source follower, with addition of 12pF output capacitance (to simulate the effects of a scope probe). Visually, the -3dB frequency is estimated around 50-70kHz, which agrees well (or at least better than initial calculations) with the measured data from TC1.

Optical Dynamic Range

With a dark current of 78fA and a maximum pixel output voltage of 450mV (as measured above), the optical dynamic range of a single pixel is found to be between 5dB (long exposure) and 45dB (short exposure). These numbers are rather poor, and will probably result in less-than-ideal image quality. However, this can be attributed to the measured dark current which is considerably higher than that of a comparable imager mentioned in the literature.

Single Pixel Static Power Consumption

The power consumed by a circuit is given by the product of its supply voltage and the total current being drawn ($P = I \times V$). One way to measure this current is to configure the circuit as shown in Figure 5.12. A large capacitance is charged to a voltage slightly higher than the supply voltage (in this case, 1.2V, instead of the 1.0V desired supply) while the circuit runs. Afterward, the voltage source is disconnected and the capacitor supplies voltage to the circuit, discharging while the circuit continues to run. As this occurs, the voltage across the capacitor is measured, and using the relation $i = C\frac{dv}{dt}$, the current drawn by the circuit can be found by evaluating its slope at the desired supply voltage, as illustrated in Figure 5.13.

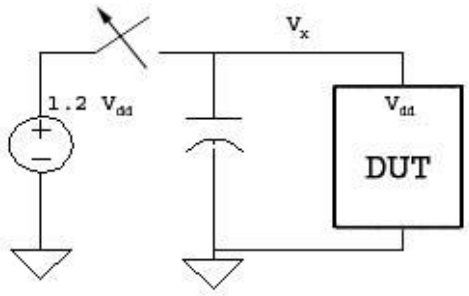

Figure 5.12: Schematic illustrating the power consumption measurement. In all cases, a 100μF capacitor is placed in parallel with a 1.2V supply, and the circuit run for a while with V_{dd}=1.2V. Then, the supply is disconnected, and the voltage across the capacitor measured while it discharges; the slope of this curve as it crosses the 1V level is then used to compute the total current drawn by the circuit.

From this method, the power consumed at V_{dd}=1.0V is measured to be 380nW for one pixel, which is near the 400nW target value.

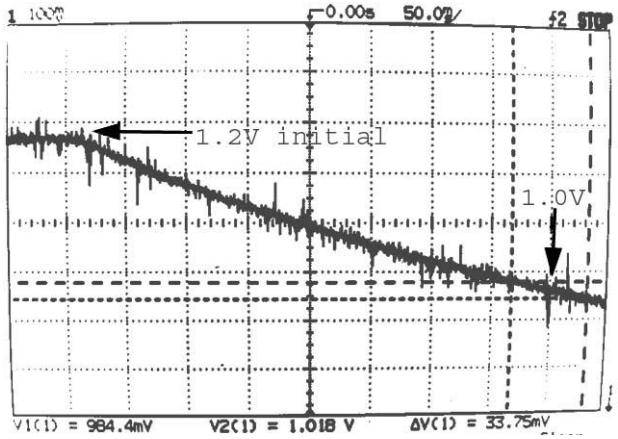

Figure 5.13: Example of measured capacitor voltage when employing the method shown in Figure 5.12. As can be seen, the capacitor is initially charged to 1.2V, and the capacitor voltage gradually decreases as it acts as the V_{dd} of the circuit. The slope at the point where the capacitor voltage is 1.0V gives an indirect measure of the current being supplied at V_{dd}=1.0V. *(The figure shows data for a 100μF capacitor hooked up to the full 128$\times$128 array running at 50kSample/second).*

Summary of Single-Pixel Test Results

Single-pixel test results are summarized in Table 5.1, comparing expected values and measured values of certain performance metrics. On the whole, the measured single-pixel parameters are close to their hand-calculated counterparts.

5.1.3 TC1 - 8$\times$8 Active Pixel Array

Following the verification of a single pixel, the array configuration and traversal, and the column voltages are verified – specifically, (1) whether the rolling-shutter architecture traverses the rows properly, and (2) if an image focused onto the array gets mapped correctly. With these results, the validity of the rolling-shutter architecture will be shown.

The 8$\times$8 array structure only contains control signals for the rows, and no column parsing is done. Instead, each of the eight columns is probed directly. A block diagram of the test structure and setup is shown in Figure 5.14, and the corresponding layout in Figure 5.15. For comparison purposes, the first column of the array is covered by a top-metal light shield, while the other seven are left exposed.

Results are verified both electrically (on an oscilloscope), and visually (with an intensity map).

Table 5.1: Summary of Single-Pixel Testing Results. For each metric, both the expected value (from hand calculations and simulation) and the measured value (from fabricated structure) are shown.

Performance Metric	**Expected**	**Measured**
Pixel Measurements		
$V_{out,DC}$ (reset)	411mV	434mV
$V_{out,max}$	500mV	450mV
Static Power	394nW	380nW
In-Pixel Source Follower		
Low-Frequency Gain	0.81	0.77
Dominant Pole Freq	$\approx$ 50kHz*	40kHz
Photodiode Characterization		
Dark Current	26fA	78fA
Dynamic Range	20-50dB	5-45dB

*NOTE: The expected pole frequency is adjusted to include the 12pF scope probe load.

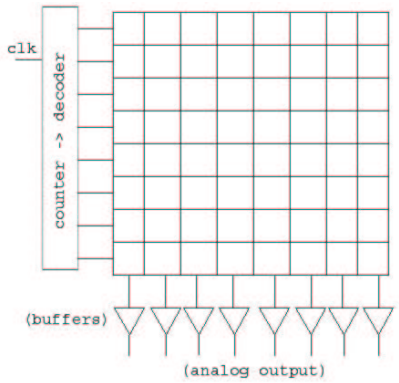

Figure 5.14: Block diagram of the TC1 8×8 pixel array. The only input is a single clock; on-chip counters generate the addresses that drive the decoder, forcing it to cycle through the rows of the array. Further, the signals are connected in a rolling-shutter configuration for proof of concept. Each column is then buffered (through a PMOS source-follower) before being wired to an output pad.

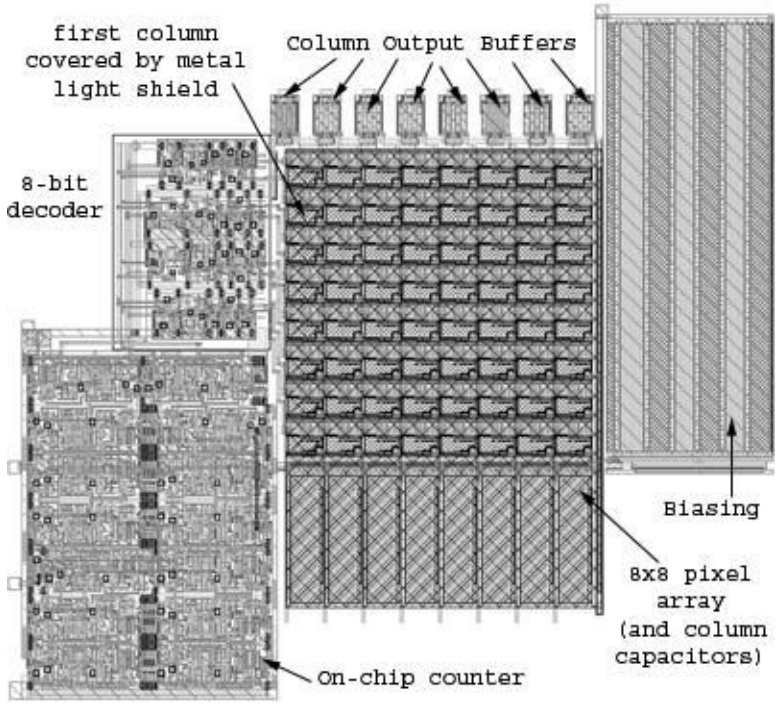

Figure 5.15: Layout of the 8×8 array system described in the block diagram of Figure 5.14. Note that the biasing consumes a substantial amount of space, as a large (1.3MΩ) on-chip resistor is used.

Electrically, the column-outputs of the pixel array behave as expected: the measured output of Figure 5.16 compares well to the expected waveforms of Figure 4.19, node (B). Visually, since it is difficult to focus a lens image onto a $64\mu\text{m}\times64\mu\text{m}$ area, intensity is verified (without a lens) by liberally sweeping a laser pointer over the array. One such intensity map is shown in Figure 5.17.

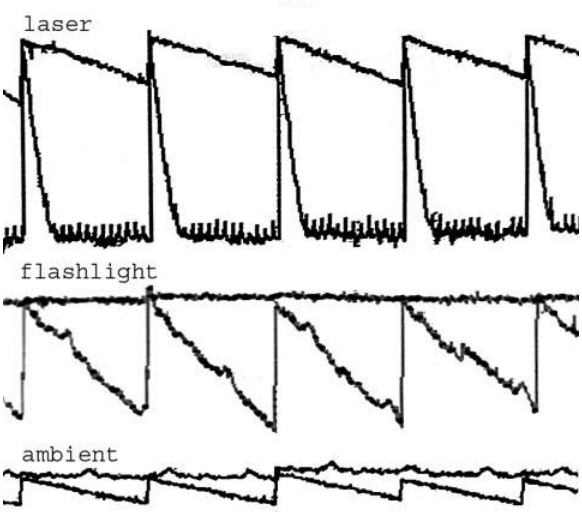

Figure 5.16: Oscilloscope output of column voltages under varying lighting conditions. In each of the three conditions shown (incident laser, incident flashlight, and ambient light), two curves are superimposed: one is the light-shielded column, and the second is the sixth column (chosen arbitrarily). As expected, (1) every column cycles through each of the eight rows' exposure periods (since no rows are ever connected to the columns during reset); and (2) higher light intensity results in a steeper slewing slope, to the extreme that direct laser exposure saturates the output.

It can be concluded from these figures that different light intensities are correctly detected. Imager sensitivity can be adjusted by changing the frequency of the system clock – a slower clock will result in a longer exposure time, which is conducive for low-light situations.

As mentioned previously, focusing an image onto an 8×8 pixel array is difficult, so detailed image (and locality mapping) tests are reserved for the 128×128 array.

5.1.4 TC2

The second test-chip (TC2) contains the entire system: the 128×128 pixel array, digital control logic, and column-output processing circuits (bias, capacitors, and analog multiplexor), the layout of which is shown in Figure 5.18. As with the 8×8 array of TC1, this structure needs only an

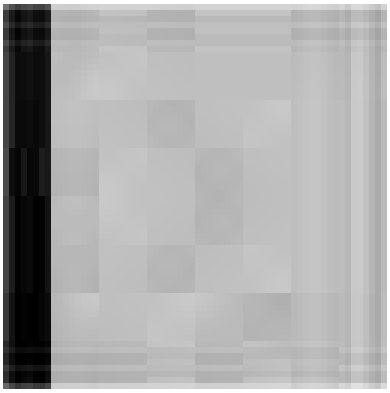

Figure 5.17: Image resulting from shining flashlight onto entire 8×8 array. In this case, no lens is used, so the intensity in all pixels should be the same except for the first column, which is covered by a metal light shield: that one should be darker. *(Image generated in LabVIEW based on 8 parallel analog inputs from the buffered columns of the 8×8 array on TC1)*

external clock and reset signal as inputs. However, unlike the TC1 array, TC2 generates only a single analog output.

During the testing process, a number of problems were discovered. First, due to a schematic error, the lower nine bits of the global 16-bit counter – which drive the analog multiplexor and auxiliary control logic – were inverted. As a result, the column-sample signal occurs at the beginning of the row-select period instead of the end as desired. In addition, the analog-mux pre-charge signal occurs at the wrong time and the control signals for the programmable-gain amplifier are reversed, nullifying the charge-sharing avoidance scheme developed earlier and rendering the PGA useless. In short, the signals of Figure 4.17 are mirrored (that is, they occur in reverse for a given row period) from what they are supposed to be.

Fortunately, the address lines to the row decoder are fine, allowing for *some* testing to be done. However, the sampling rate must be slowed considerably to allow for substantial exposure time, and the analog mux output pin is unbuffered, further exacerbating the charge-sharing problem.

Transient output

One of the problems of measuring an unbuffered analog output node (without the charge-sharing fix) on an oscilloscope with probes that add 11pF of capacitance is that voltage readings are lower than expected. Nevertheless, photo-response can be measured by observing the *relative* voltage differences (compared across lighting situations) at the system output, while discounting inaccura-

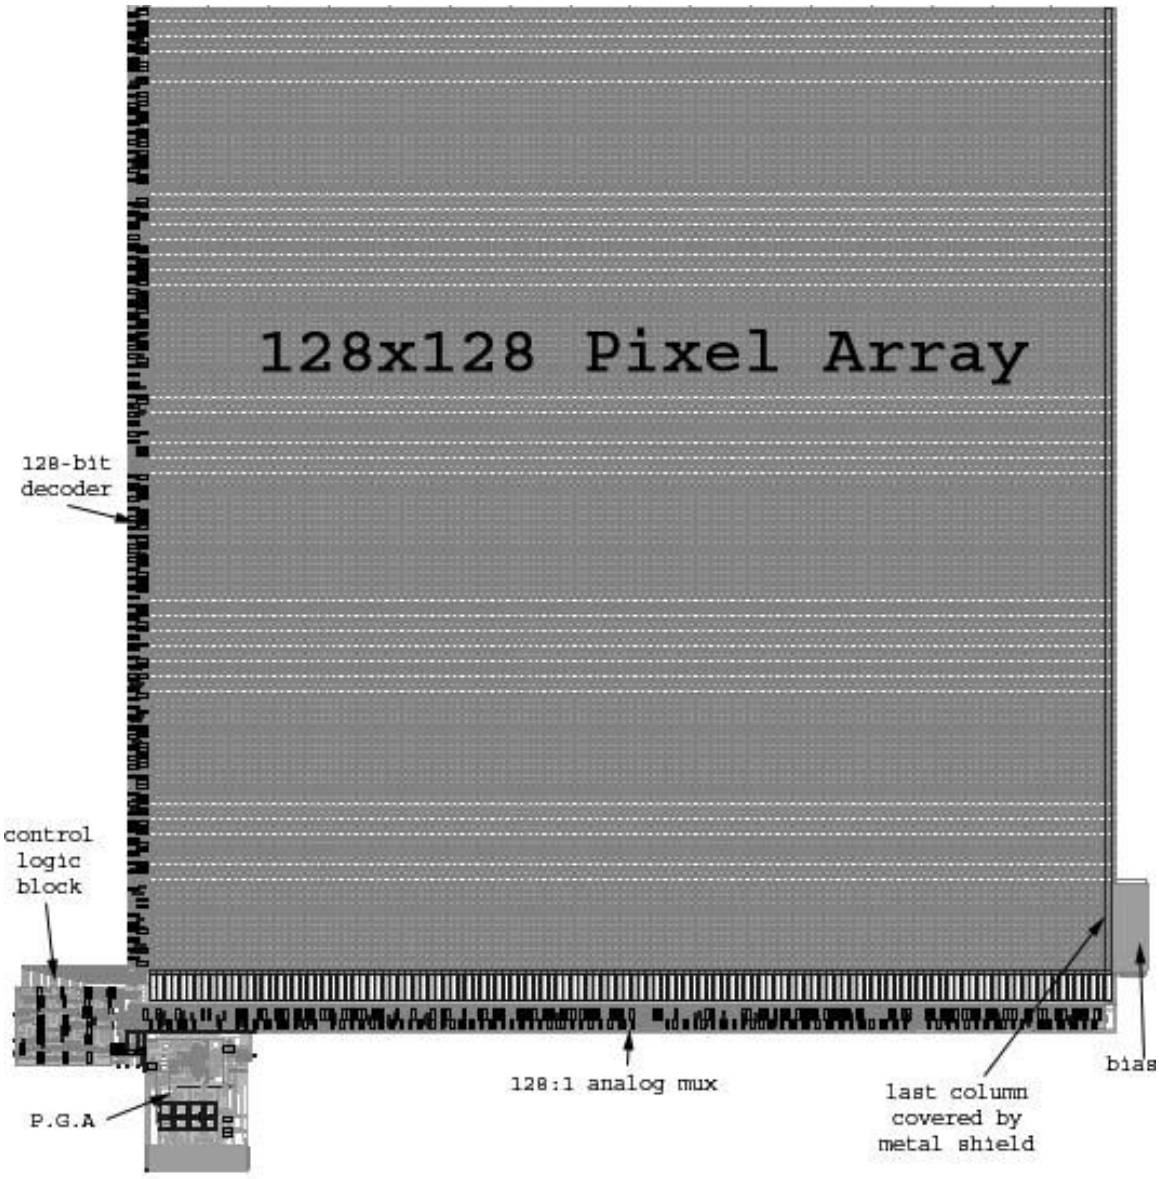

Figure 5.18: Layout of camera system (minus the ADC) included on TC2. The system area is dominated by the pixel array (which measures 1mm on a side)

cies in the absolute voltage level.

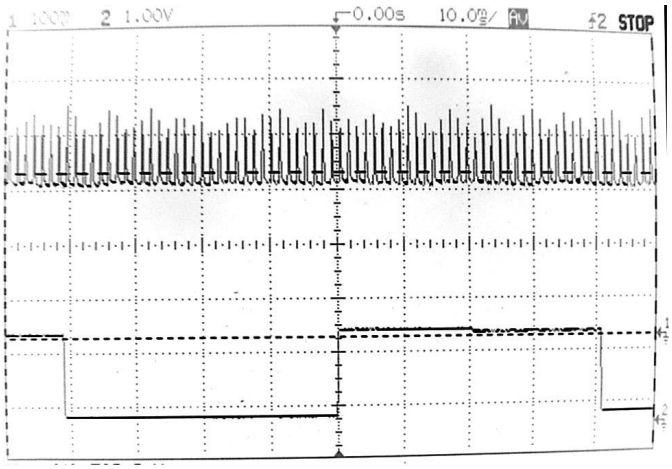

Figure 5.19: Oscilloscope waveforms of the (unbuffered analog) output of TC2. Shown are 64 row periods on top of a "half-trigger" signal used only for generating this image; if the usual trigger signal were used, the 128 row period outputs would have been impossible to see. For each of the 64 rows, the output appears to go to a high level (corresponding to the one column of the entire chip that is covered by a metal light shield) and then fall down to a "baseline" level proportional to the ambient light incident on the imager.

Figure 5.19 shows the analog output for the top 64 rows of the imager array under ambient light. For the sake of clarity, all 128 rows are not shown since it would be difficult to distinguish one row from another in an image of this size. By comparison, Figure 5.20 is a picture of the same signals, but with a spot of light focused on the top of the array. Here, it can be seen that for the rows that are affected by the light, there is a $\approx$100mV drop in the output signal level. For further verification, Figure 5.21 shows eight of these rows from near the center of Figure 5.20: it is clear that electrically, (1) a spot of light is detected correctly in its location, and (2) ambient light levels are the same for the pixels that are exposed only to ambient light.

So after verifying the electricals of the imager chip, let us see what this can translate to in terms of an image.

Visual Verification

A more reliable measure of TC2's performance comes from LabVIEW, since its data acquisition device does not load the analog output as much as a scope probe. In addition, as LabVIEW is

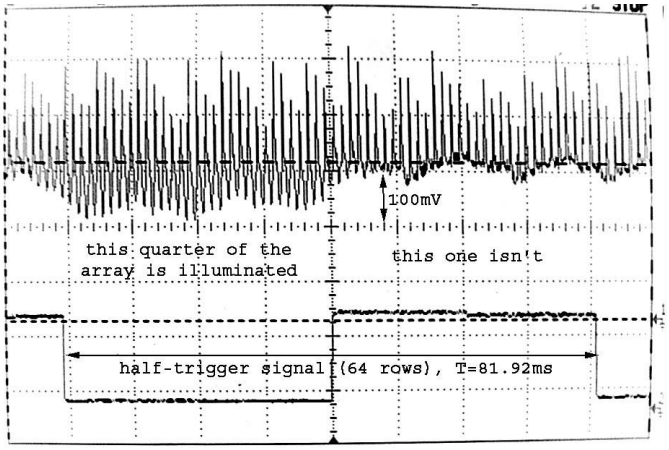

Figure 5.20: This is the same measurement as in Figure 5.19, except that a spot of light is focused on part of the array. From this array-level view, we can see that there are lower voltage readings corresponding to the rows on which light is focused.

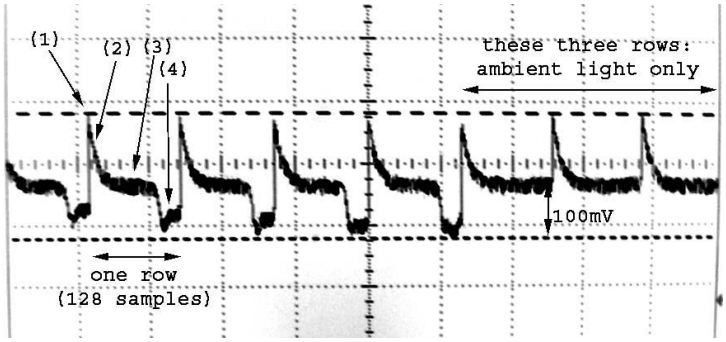

Figure 5.21: This is a "zoomed-in" version of the measurement taken in Figure 5.20. As annotated, **(1)** the end of each row period is denoted by the column covered by the metal light shield; **(2)** Since the charge-sharing-control mechanism is broken, there is a slight lag among samples as the voltage gradually falls to **(3)** the ambient level (which is measured on parts of the row where the spot of light is not focused; and **(4)** there is about a 100mV drop for the samples that the spot of light is shining on.

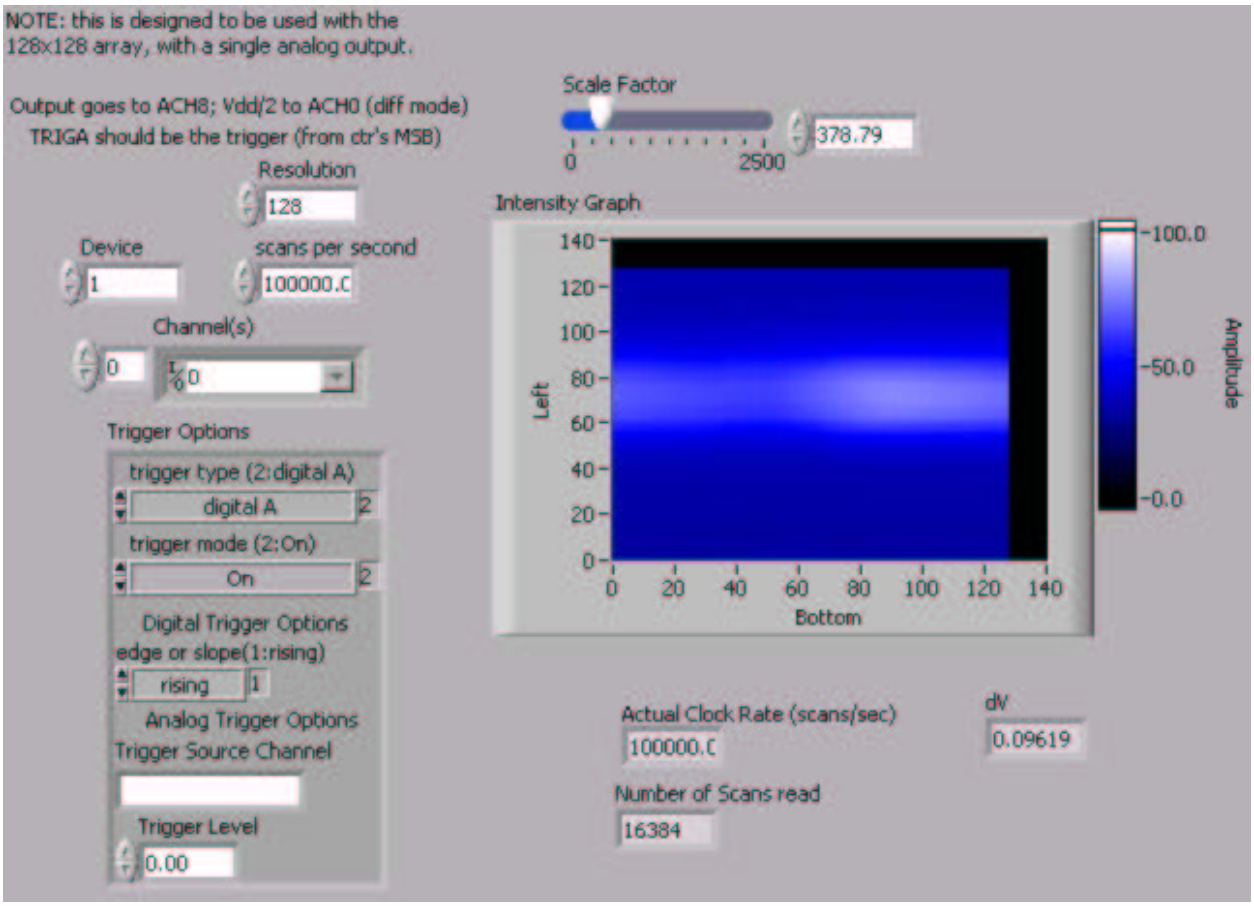

Figure 5.22: Screen-shot of LabVIEW interface to the camera, used to test the 128×128 imager. Full details of the interface are given in the Appendix.

a software test environment, it can be used to simulate the effects of the PGA while generating images based on light intensity (in the form of voltage levels). The interface to the LabVIEW program created for this serial-voltage-to-image-map conversion is shown in Figure 5.22, and explained fully in the Appendix.

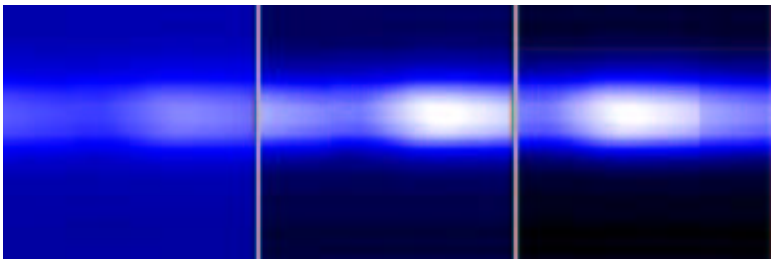

Figure 5.23: Image generated by focusing a laser through a 3mm lens onto one side of the 128×128 array. The image appears to have horizontal lag that wraps around the edge of the picture from one row to the next, which can be attributed to the charge-sharing problem. Three images are shown above: (1 - left) the raw output from LabVIEW; (2 - center) the same picture with some adjustments to output levels made in an image editor; and (3 - right) a spliced version of the image, to illustrate that the image smear indeed jumps from one row to the next (hence the discontinuity in the spot), giving the false impression that there are *two* distinct spots of light focused onto the array.

The LabVIEW-generated images clearly demonstrate locality (namely, that light aimed at one part of the array can be differentiated visually from all other parts of the array), as in Figure 5.23. Similar figures taken at varying sample rates (100kS/s, 50kS/s, and 10kS/s) show that reducing sample rate also reduces the effect of charge sharing (i.e., there is less horizontal lag). With this in mind, images taken at a much slower sample rate than desired (e.g., 5000 samples/second) should be relatively recognizable, so long as the lens is focused properly. However, none of these images are included here since the final-phase testing images are of more significance.

Power Consumption

Using the same power-measurement method as before, the current drawn by the running circuit at 100kSample/s is found to be 60μA at a 1V supply, which corresponds to 60μW power consumption. The difference between this number and the 50μW design specification can be attributed to a number of factors:

- Leakage during reset: during pixel reset, there is a direct path from V_{dd} to ground through the reset PMOS and the reverse-biased photodiode. In the "short-exposure" timing scheme (which is what is fabricated in TC2), only one row is being exposed at a time, so on average, there are over sixteen thousand pixels leaking current in the 127 other resetting rows.

- Biasing: the initial selection of 400nA bias current gives only *approximately* 50μA total current. Considering all 129 branches (columns and bias), each drawing 400nA, the total static current of the in-pixel amplifiers (even before considering leakage) is closer to 52μA.

After considering these two effects, a measurement of 60μA of current is quite reasonable. Further, at the full 100kSample/sec. rate, one frame of 16384 pixels is processed in 164ms, corresponding to an energy consumption of less than 10μJ per frame.

5.2 Final Testing

In the last iteration of design verification, the bugs found in TC2 were corrected, and exhaustive top-level simulations run before taping out. Both the short-exposure and long-exposure (previously described in the design of the digital control circuits, and Figures 4.12 and 4.13) versions of the 128×128 system were fabricated. These revised systems are exactly the same as TC2 systems, except:

- The analog multiplexor output is buffered.

- The inverted address lines have been corrected.

- The trigger signal needed to acquire the samples in LabVIEW is generated on-chip, instead of with external counters.

This system's testing methodology is the same one used for TC2 – first, measure the signals on an oscilloscope and compare them to the expected values presented earlier, and then interpret the signals to form an image in LabVIEW.

5.2.1 Electrical Characteristics

The buffered analog output of the final test-chip is measured on an oscilloscope, and the output shown in Figure 5.24 under ambient light, and Figure 5.25 under mixed illumination. With the

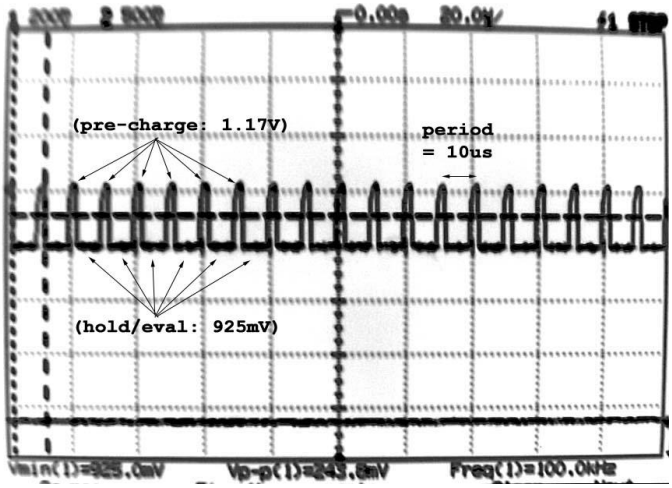

Figure 5.24: Analog output (buffered) of the long-exposure chip under ambient light, with pre-charging scheme in place. And as expected, there is a distinct sample period of 100kS/s, one fourth of which is spent in pre-charge, and the other three-fourths of which hold a voltage corresponding to a pixel's measured intensity. The buffered pre-charge value is about 1.2V, and the ambient "hold" value is about 200mV less.

pre-charge scheme in place (to eliminate image lag due to charge-sharing), the measured output behaves as expected from Figure 4.20, node (D): for one quarter of each sample period, the output resets to a pre-charge value; and during the remaining three-quarters of the period, the output voltage is a measure of the light intensity. Additional measurements of the analog output were taken with a lens placed over the chip (details to follow later) and a spot of light directed at various parts of the array with similarly-expected results. Therefore, it is expected that as long as the output is sampled by subsequent stages during the "hold" state, images should be able to be generated from the chip's output without too much difficulty.

Measured Voltage Levels

The measured output voltage ranges from 1.17V (during pre-charge) down to 728mV (under a laser) as shown in Figure 5.25, for a swing of 442mV. The expected analog output voltage mentioned previously is between 88mV and 500mV (swing of 412mV) and clearly, there is a relatively good match in output swing. But the difference in DC level can be attributed to the PMOS source-follower (Figure 5.26) used to buffer the signal. The external resistor is chosen to provide 100μA of bias current, and at this bias, $|V_{GS}|$ of the PMOS device (with a W/L ratio of 100 and using

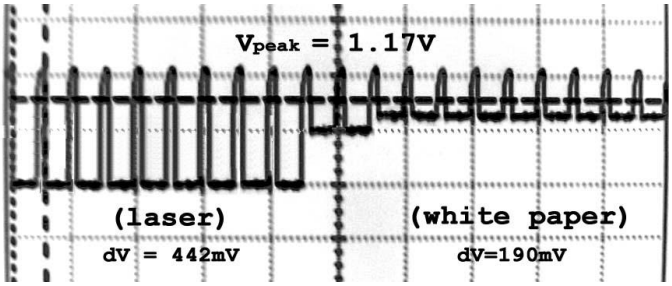

Figure 5.25: Analog output of the long-exposure chip with pre-charging scheme in place. In contrast to Figure 5.24, a sheet of white paper is placed over half of the array and a laser aimed over the entire array. Assuming that a laser causes pixel saturation, the "maximum" drop is almost 450mV.

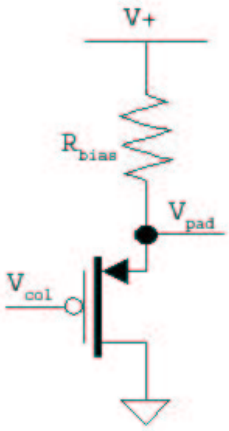

Figure 5.26: Schematic of Brian Leibowitz's PMOS source-folower used as an on-chip analog buffer. For the purposes of connecting this chip to a test board, a bias resistor R_{bias} and a large voltage V^+ can be chosen to provide any desired bias condition. In testing, R_{bias}=81.2kΩ, V^+=9.0V, and the resulting bias current I_{bias}=100μA.

k_p' and V_{Tp} values from the electrial design rules) is 700mV (to first order). Therefore, the measured output range of 728mV-1.17V corresponds roughly to an unbuffered range of 28mV-470mV, which agrees with the expected values. The difference (28mV vs. 88mV) at the lower-end of the voltage range can be attributed to imperfect modeling of the parasitic junction capacitances used in the calculation of the analog mux output capacitance; as it is, the capacitance-per-unit-area of a source/drain junction is just an approximation and probably does not reflect the exact performance of a fabricated device. Further disagreements in the DC level show up because second-order effects (e.g. channel-length modulation in the PMOS) are ignored in the hand calculation of $|V_{GS}|$.

Noise

To measure electronic noise, the chip is covered by black electrical tape and placed in the shadow of some lab equipment (to make it as dark as possible). Further, the charge-sharing-avoidance pre-charge scheme is disabled to eliminate any extra sampling noise that it might introduce; and because in this noise measurement (in complete darkness), image lag does not matter. Then, it is connected to an oscilloscope, the scope input is AC-coupled, and V_{rms} read off of the display (e.g., Figure 5.27).

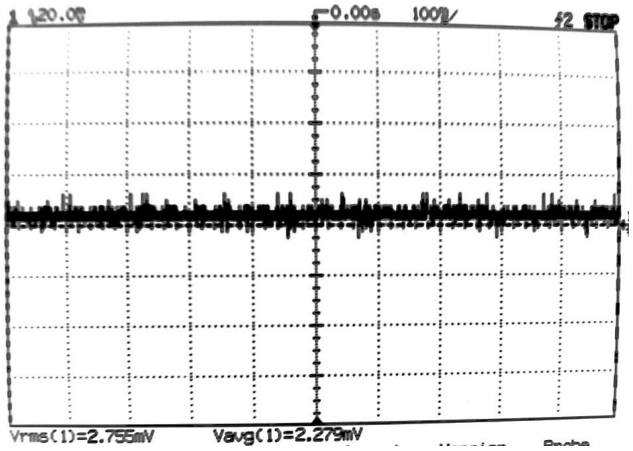

Figure 5.27: Oscilloscope output of noise measurement. Here, the charge-sharing avoidance pre-charge scheme is disabled, the imager covered with black electrical tape, and the chip otherwise left running. On the scope, the signal is AC coupled, and the V_{rms} voltage is taken as the noise measurement (2.755mV, as shown here).

In an attempt to eliminate the noise caused by extraneous sources (e.g., the scope probe), a second calibrating measurement is taken in which the probe's positive and negative terminals are

shorted together, and the same V_{rms} value measured. By taking the square root of the difference of the squares of these two, the noise due only to the circuit can be obtained.

$$V_{rms,ckt} = \sqrt{V_{rms,meas}^2 - V_{rms,cal}^2} \tag{5.1}$$

The voltages measured are $V_{rms,meas}$=2.755mV (shown in Figure 5.27) when the running circuit is connected to the probe, and $V_{rms,cal}$=2.572mV for the calibration measurement. Using equation 5.1, the effective total integrated noise of the imager is computed as 994μV.

By comparison, previous hand calculations estimate the rms noise at 835μV; but the 159μV difference can be attributed to noise from the PMOS source-follower output buffer. In accoradnce with the analysis presented in [45], noise from the external bias resistor and the PMOS are:

$$\overline{I_R^2} = \frac{4kT}{R}\Delta f \tag{5.2}$$

and

$$\overline{I_{PMOS}^2} = 4kT\gamma g_m \Delta f \tag{5.3}$$

respectively, where γ is approximately 2/3 for the PMOS. Using $R = 81.2\text{k}\Omega$ and a bias current of 100μA to compute g_m, and applying small-signal noise analysis, the output-referred mean-squared noise components are:

$$\frac{\overline{V_R^2}}{\Delta f} = 1.71 \times 10^{-19}\text{V}^2/\text{Hz} \tag{5.4}$$

and

$$\frac{\overline{V_{PMOS}^2}}{\Delta f} = 9.54 \times 10^{-18}\text{V}^2/\text{Hz} \tag{5.5}$$

for the bias resistor and PMOS, respectively. Adding these two produces a total output-referred mean-squared noise of 9.711×10^{-18} V^2/Hz. Approximating the common-drain buffer as a single-pole system, the total integrated output-referred mean-squared noise is given by [45] as:

$$\overline{V_{total}^2} = \overline{V_{output,dc}^2} \times B_n \tag{5.6}$$

where the noise bandwidth B_n is $\pi/2$ times the pole frequency (in this case, g_m/C_L). Putting this all together, the total integrated output noise of the PMOS source-follower V_{total} (to first-order) is found to be 125μV rms.

While this number does not provide an *exact* match between expected and measured values for rms output noise, it is nevertheless in good agreement, considering:

- The pixel's reset and integration noise depend on the value of the parasitic photodiode capacitance. Hand calculations only estimate this capacitance based on the electrical design rules, which already introduce some inaccuracies.

- In addition, other model parameters used in hand calculations (e.g., transistor properties) are also just quoted out of a design rule document and again, are only approximations of the fabricated silicon device behavior.

- The hand-calculated output-buffer noise relies on an *estimated* load capacitance C_L, and the approximation that the source-follower is a one-pole system. These will give good first-order numbers, but introduce error when considering exact system performance.

- Noise due to interference or other external sources that affect the oscilloscope reading are not perfectly canceled out by the calibration method mentioned above.

Other possible sources of noise include the CMOS pass-gates that act as resistors when "closed" (to disable the PGA and send the analog output to the buffer and off-chip), sampling noise from the column circuitry (even though the imager is dark for the noise measurement, it is still running, and the counter is still cycling through the array), and even supply noise (since a single 2.0V supply was used to generate three bias voltages: 2.0V for the trigger signal, 1.0V to power the array, and 500mV for the pre-charge scheme) as a result of the resistive-divider network employed.

All told, the measured value is easily within an order of magnitude of the expected value from hand calculations, and given the possible sources of inaccuracies, the agreement is good enough.

Summary of System-Level Electrical Figures of Merit

A summary of the hand-calculated expected output vs. the measured output is given in Table 5.2. As demonstrated throughout this chapter, the measured voltages at various points along the signal path agree with the expected values presented previously in Figures 4.19 and 4.20. It is no surprise then, that the figures of merit also agree.

5.2.2 Visual Verification

Since the previous electrical verification shows that the chip behaves as expected, it is reasonable to expect that images can be generated using the same LabVIEW interface as in the TC2 tests. The optics needed to focus images onto the chip are shown in Figures 5.28 (top-view) and 5.29

Table 5.2: System-Level Figures of Merit Summary

Metric	**Expected**	**Measured**
Power Consumption	51.6μW	60μW
Maximum Output Voltage	1.2V [1]	1.17V
Minimum Output Voltage	788mV [1]	728mV
Analog Output Swing	412mV	442mV
Output Noise	960μV rms [2]	994μV rms

[1] The minimum and maximum hand-calculated output voltages are adjusted: 700mV (hand-calculated $|V_{GS}|$) is added to account for the DC offset introduced by the PMOS source-follower.
[2] The expected total noise value is approximated as the sum of the expected total integrated noise from the imager and the hand-calculated total integrated noise of the PMOS source-follower.

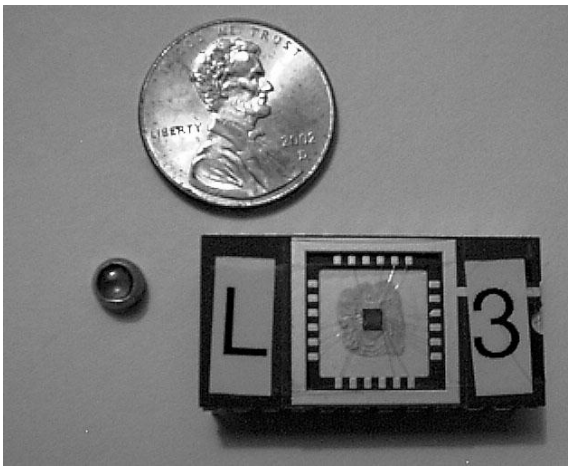

Figure 5.28: Relative sizes of the packaged chip and lens mount tube, compared to a penny. The 1.6mm$\times$1.6mm chip is packaged in a standard 24-pin package available in the UC Berkeley Microlab, and covered with a clear plastic coverslip. A 3mm-diameter, 3mm-focal-length lens is attached to one end of the tube (made out of a pipe washer from a hardware store) with epoxy and the tube's edges shaved down to position the lens at a desired distance above the imager chip.

(side-view). Since the chip itself is packaged in a standard 24-pin package (and not a specialized camera body), a "lens tube" is constructed to be placed over the chip, blocking out ambient light and simulating a camera body. For the purposes of this test, a 3mm-diameter lens with a 3mm focal length is affixed to one edge of the metallic tube. The tube is then shaved down at the other end to position the lens about 2-3mm above the imager when it sits on top of the chip package, as in Figure 5.29.

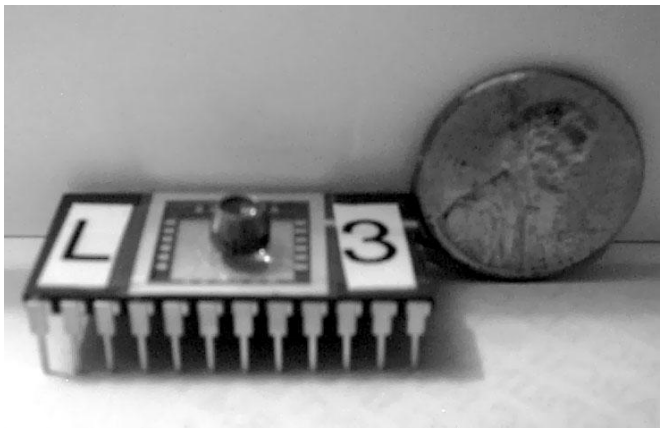

Figure 5.29: Side view of lens mount tube placed over the mounted chip (the same components as shown in Figure 5.28, compared to a penny. This configuration ideally will focus images onto the chip while blocking out significant amounts of ambient light.

With this setup, spots of light (from a flashlight, held about 1m above the lens) can be directed to distinct parts of the array and representative images captured, as in Figure 5.33. However, these spots are *unfocused*, and after a number of days spent shaving down the edges of a metallic tube in unsuccessful attempts to improve the focusing, this method was abandoned in favor of a more time-efficient means.

Artificial Optics Setup

To improve the ease of focusing, an artificial optics environment is created with a drill-press stand, optical lens holders from Edmund's Laboratory Optics Kit, and a 3mm-diameter, 3mm focal-length lens identical to the one used in the lens mount tube described above, and is shown in Figures 5.30 and 5.32. Objects for these configurations are placed at 5cm and 15cm above the lens, respectively.

In each scenario, the test board (containing the imager) is placed on the base of a drill press stand, and the lens holder clamped to the stand and suspended over the chip (as shown in Figure

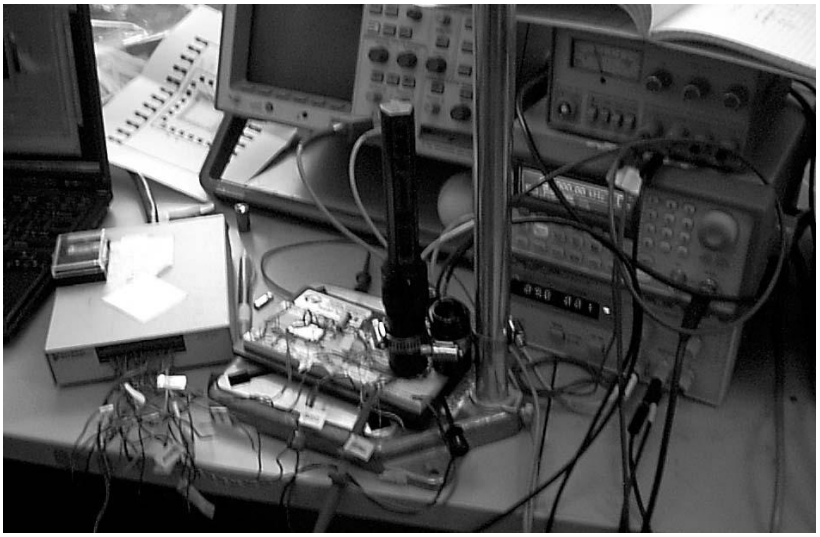

Figure 5.30: Test setup used to image the test patterns – the flashlight and checkerboard – at an object distance of 5cm. Shown are the test board, external signal generators (power supply and function generator), and output receivers (LabVIEW DAQ interface, laptop computer, and oscilloscope). The lens assembly is held in place using holders from a laboratory optics kit, and pipe clamps that bind it to the drill-press used as a base station for the test board. The flashlight is placed *directly* on top of the 5cm lens holder.

Figure 5.31: Close-up picture of the lens mounted over the imager chip. The lens is stationed approximately 3mm from the imager itself (though the exact distance is unknown).

5.31). From there, the focus can be adjusted simply by turning the lens holder, moving the lens incrementally closer to or farther from the imager in a process that is quicker and easier than shaving the edges of a metallic tube.

However, as can be seen in Figure 5.31, there is a slight gap between the lens holder and the top of the packaged imager through which ambient light can add noise to the image projected by the lens. But since the lens holder itself is much wider than the imager and casts a shadow over the area surrounding the chip, and the distance between the chip and lens holder is relatively small, this should not be a major problem. Preliminary voltage measurements of the long-exposure chip show that compared to complete darkness (black electrical tape placed over imager), the artificial lens setup produces a voltage reading that is 8mV less. It is further calculated that an 8mV drop at the analog output corresponds to an additional 2.5fA of (dark) photocurrent, or *less than 1dB* reduction in optical dynamic range. Therefore, it is concluded that the extra ambient light introduced by using this artifical optics setup does not contribute significant noise, and that an image can be obtained from this configuration that is as good (visually) as an image from the lens mount tube that potentially blocks out more ambient light.

With this in mind, and knowing that recognizable images should be attainable even through the contrived optics, visual verification then follows in two stages: first, some test patterns are imaged onto the chip for focus purposes and to prove functionality; second, real images are captured.

5.2.3 Test Patterns

The two test patterns used are (1) a flashlight spot – to prove that there is no image lag, and (2) a checkerboard pattern for focusing the lens. In these tests, all patterns are imaged by the short-exposure chip at 100kS/s (input clock is at 400kHz).

Unfocused Spots of Light With No Lag!

First, a flashlight is turned on and held about two feet above the lens, and the resulting images are shown in Figure 5.33. In stark contrast to the image lag of TC2 shown in Figure 5.23, the bright spots are contained in a circular area (corresponding to the shape of the flashlight), and there is no residual lag of brightness. From this, we can conclude that the charge sharing problem has been overcome in the final test chip.

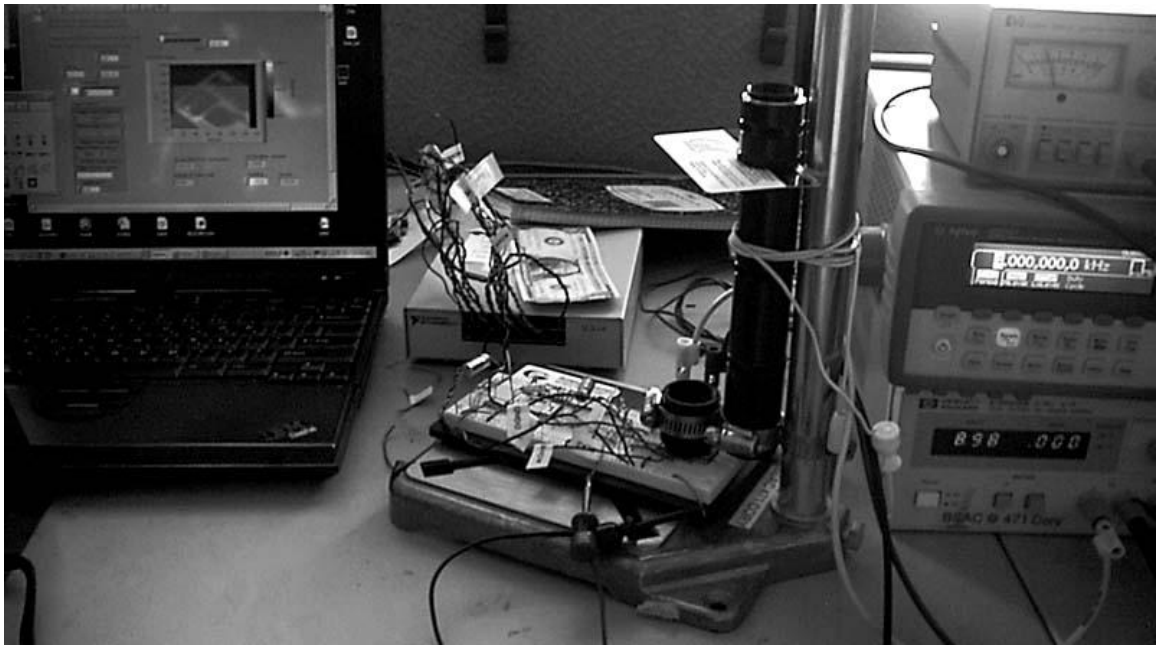

Figure 5.32: Test setup used to capture images of objects placed 15cm above the lens. Shown are the lens assembly – consisting of the lens, the drill-press base, and the black cylindrical holders, the object (in this case, a UC Berkeley Student ID card), the circuit board, and the output devices: LabVIEW DAQ card and laptop computer. A close examination of the screen of the laptop computer will show the LabVIEW interface (previously described by Figure 5.22), and an image of the ID card in the image map.

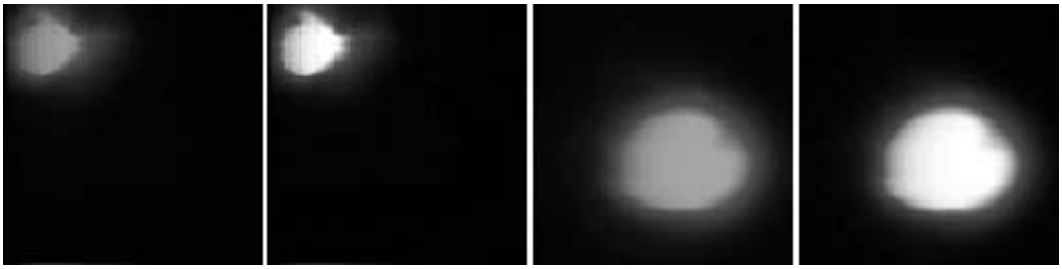

Figure 5.33: These captured images demonstrate that the image lag problem encountered in TC2 (see Figure 5.23) has been fixed in the final chip, since the charge-sharing-control pre-charge scheme works. From left to right, these four images are: (1) raw image from LabVIEW showing a localized spot of light that does not streak horizontally; (2) the same image, but with some contrast adjustments made in image-editing software; (3) a raw image of a spot of light localized in another part of the array; (4) an adjusted version from image-editing software to boost the contrast.

Focused Flashlight & Checkerboard Patterns

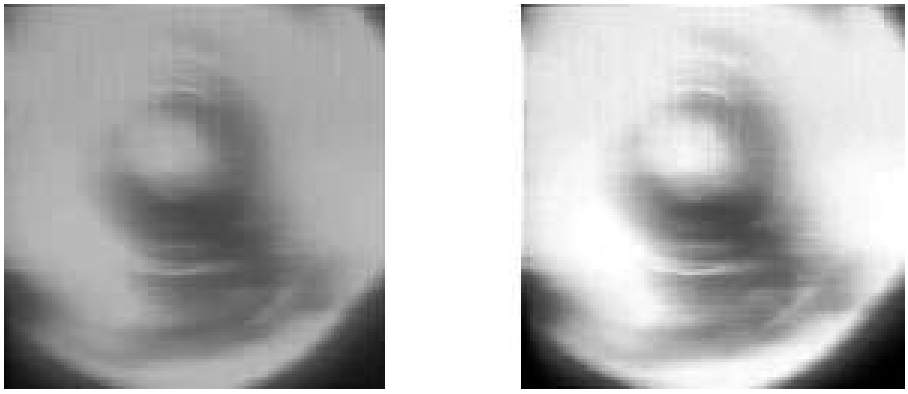

Figure 5.34: From the test setup shown in Figure 5.30 (in which a flashlight is perched directly on top of a lens holder), this is a focused image as detected by the chip and interpreted by LabVIEW. It is easily identifiable as a flashlight, with its characteristic bright bulb in the center, and rings around it caused by the reflective surfaces in the bulb holder. Two versions of the same image are shown here: on the left is the raw data from LabVIEW; on the right is the result of tweaking the output levels in an image editor to improve contrast.

Second, this same flashlight is placed directly on top of the lens holder (Figure 5.30), and the lens adjusted until the imager is focused for objects at this 5cm distance. The resulting image is shown in Figure 5.34. In addition to focusing the lens for the imager chip, this picture gives an estimate of the field of view with the 3mm-diameter lens used. Since the 3cm flashlight diameter gets mapped almost across the imager's diagonal, the horizontal field of view is determined to be 29° for all test patterns using this setup, which is close to the intended 36° field of view for the final button-camera system.

Next, checkerboard patterns are focused onto the chip from the same 5cm distance as the flashlight mentioned above. These patterns were generated on a computer and printed onto a sheet of paper, and this paper placed between the flashlight and the lens holder, producing the images in Figure 5.35 (for 10-pixel and 6-pixel pitch squares). In addition to the two patterns shown here, another pattern with four-pixel-wide squares was successfully imaged; however, a pattern with two-pixel wide squares was not identifiable. From this, we can conclude that the imager is in good enough focus to identify objects that are as small as four pixels wide.

The remaining blurriness can be attributed to cross-talk and focusing. Cross-talk occurs when incident photons penetrate beyond the photodiode junction and generate charge in the p-substrate

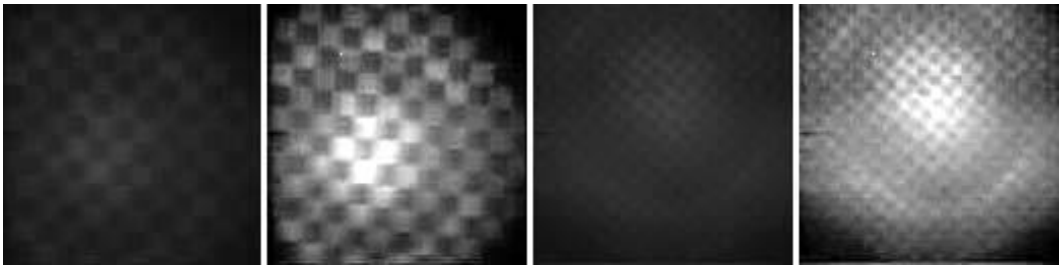

Figure 5.35: Four images of two unique checkerboard patterns used primarily to focus the lens assembly. As with other images captured from LabVIEW, both the raw data version and an image-editor-adjusted version are shown. The two images on the left show a checkerboard pattern where each square is 10 pixels on a side; the two on the right show a pattern where the squares are 6 pixels on a side. Though the focusing is not perfect, the squares can clearly be identified.

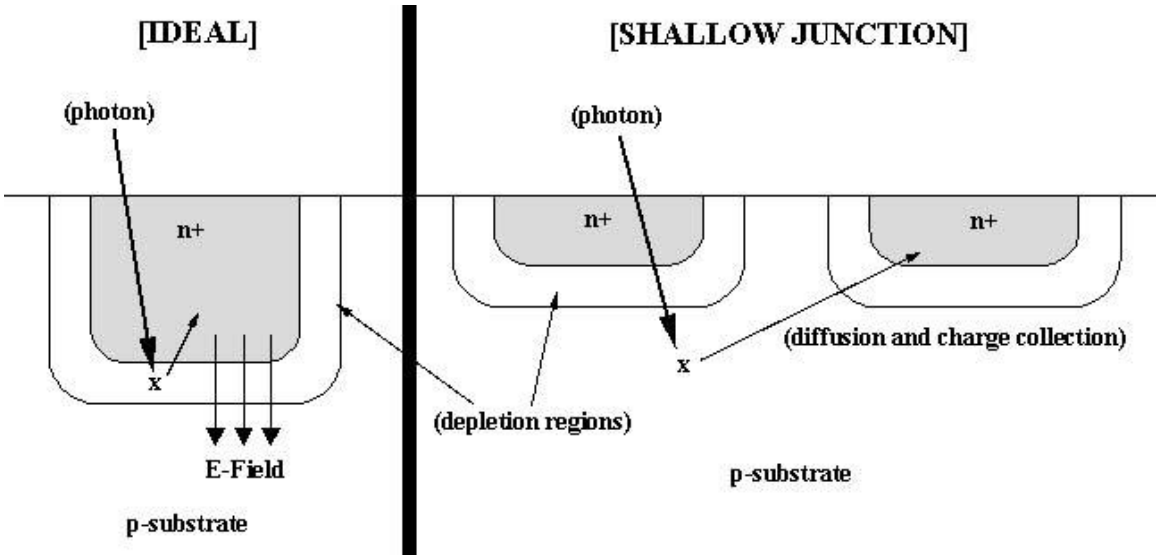

Figure 5.36: Illustration of "ideal" photo-generation of charges (left), and cross-talk (right) for two different junction depths with incident light of the same wavelength (and absorption depth). In the ideal case, the photon is absorbed in the depletion region, and the resulting charges swept by the induced electric field to the quasi-neutral (n^+ and p-substrate) regions of the pixel in question producing photocurrent. On the other hand, in 0.25μm processes, visible light gets absorbed in the quasi-neutral p-substrate (because the junctions are so shallow) and generates photocurrent via carrier diffusion. Cross-talk occurs when the charge diffuses to and is collected by a neighboring pixel, as shown above on the right, and can result in image blurring.

[34]. Since the high electric field of a diode depletion region is not present in the quasi-neutral substrate, it is possible that this charge will be collected by a neighboring pixel instead of the one through which the photon arrived, as shown in Figure 5.36. In the CMOS8 process, n^+-diffusion junction depths are less than 0.15μm; however, as shown in [34], the light absorption depth of most visible light (specifically, for λ >450nm) in silicon is greater than 0.25μm, resulting in some substrate-generated charge for all images of interest. And while the addition of a substrate contact in each pixel helps reduce this effect, it may not eliminate cross-talk completely. With regard to focusing, finer resolution could be achieved with a more detailed focusing job; however, that is not the focus of this project (no pun intended) and therefore it is not explored here.

5.2.4 Images

More intriguing results come from trying to generate an image of an object that is not back-lit by a flashlight. For this, the following setup (shown in Figure 5.32) is used:

- **Object Distance:** 15cm
- **Object to be imaged:** my UC Berkeley Student ID Card (Figure 5.37).
- **Illumination:** a fluorescent lamp placed 20cm away from the object (on the table)
- **Field of view:** 22°

Figure 5.37: This student ID card is used as the test object for the imager (test setup has been previously depicted in Figure 5.32). Generated images are shown in Figures 5.38 and 5.39 for varying exposure times.

The Effect of Exposure Time

The first few attempts to capture an image with the short-exposure chip at the highest sampling rate (100kS/s) were unsuccessful: the image on the screen was nothing more than a blob bearing only a remote resemblance to a rectangular card. Apparently, the full-speed exposure time is not long enough to allow for the features to be differentiated. With this in mind, the sampling rate is reduced to 50kS/s, and then 20kS/s, and although images at these rates are more discernible than at 100kS/s, they are still a mess. In addition, the imperfect lens focusing contributes quite a bit to the lack of image clarity.

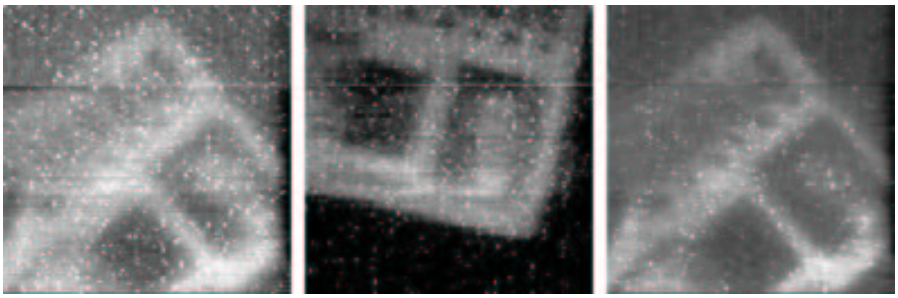

Figure 5.38: Three imager-generated pictures of the Student ID card of Figure 5.37 at sample rates of 10kS/s (leftmost, and center images) and 5kS/s (right image). On the whole, the pictures are blurry; however, the features of the card can be distinguished – the face and picture in the lower right and the title at the top of the card. Additionally, the 5kS/s sample appears to have slightly less noise (random specks of stuff) than the 10kS/s images.

However, at sampling rates of 10kS/s and 5kS/s (12.8ms and 25.6ms exposure times, respectively), the on-screen image generated by LabVIEW starts to resemble the ID card, as shown in Figure 5.38. For these figures, the LabVIEW images are converted to shades of grey (instead of blue), the aspect ratio changed to produce a 128×128 resolution image, and the output levels modified slightly to enhance the contrast. Otherwise, no processing was done. The images in Figure 5.38 are rather noisy – there are random specks of brightness in the image; however, there is less noise in the 5kS/s image than in either of the 10kS/s ones.

Next, further increasing the exposure time of the image can be achieved in two ways: (1) run the short-exposure chip at 1kS/s (128ms exposure time), or (2) just use the long-exposure chip at 100kS/s (162ms exposure time). While both scenarios have similar exposure times, the short-exposure chip would take 100 times as long to generate the image; so the long-exposure chip is

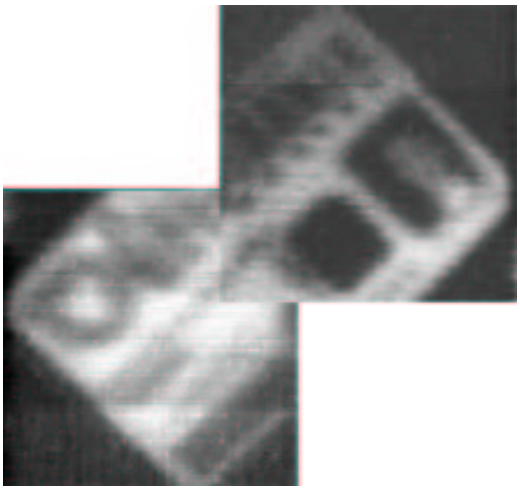

Figure 5.39: Composite image of my Student ID Card generated by stitching together two separate images, both of which were captured by the long-exposure chip at 100kS/s. It appears that this longer exposure time (compared to the images captured in Figure 5.38) produces a cleaner image (though the images are all slightly blurry).

used instead.

These images are much cleaner, as the specks of noise are drastically reduced. (Of course, it is also possible that the long-exposure chip just happens to be less noisy than the short-exposure, or that something else in the environment changed while modifying the test-bench, but these will not be considered.) Two separate images of different parts of the ID card were taken and stitched together in an image editor to produce Figure 5.39. Here, the features of the ID card (such as the the picture in the bottom right, the title, the University of California logo in the top left, and the bar code) are easily identifiable, despite the image blurriness.

Finally, in keeping with the tradition of image sensors, a $1 bill is captured (again, with the long-exposure chip at 100kS/s) and shown in Figure 5.40. As with all the other pictures generated by this test setup, it is blurry, yet its features are identifiable – the shape of George Washington's head, the banner across the bottom that should have "Washington" on it, and the location of the "ONE DOLLAR" text underneath.

5.2.5 Post Processing

Two post-processing elements are intended for the final camera system: (1) a programmable-gain amplifier (PGA) designed by Al Molnar, and an ultra-low-power 8-bit ADC designed by Michael

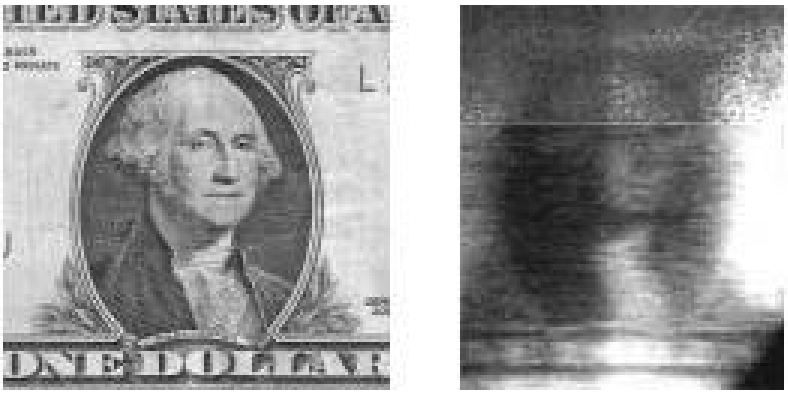

Figure 5.40: The image on the right is the imager-generated version of the dollar-bill shown to the left. This was done on the long-exposure chip at a full 100kS/s, and as with the other pictures, though there is considerable blurriness, the features can nevertheless be seen.

Scott [38]. As both of these elements are designed and tested by others, they are not emphasized in this project; however, some preliminary tests are done to verify whether the system will behave (electrically) as expected.

PGA

As mentioned previously, the final test-chip includes an on-chip PGA that can be selectively enabled. All the prevous results mentioned in this chapter have been taken without the PGA, but it is useful to at least observe its behavior and verify the on-chip system-integration controls.

The PGA is designed specifically to drive the ultra-low-power ADC referred above. Its gain is set by two control bits, which in the test chip, are supposed to be hard-wired for a gain of 2. In testing, however, it is discovered that the gain specification was misunderstood, resulting in a PGA gain of 4 ($= 2^2$ – the control bits set the *power* of the gain, and not the gain itself) instead. Double-checking the schematic in software confirms that this is the case. Nevertheless, the PGA appears to work and be timed properly by the digital signals generated on-chip.

In testing, the (long-exposure) imager chip is first exposed to the overhead room lighting (and any ambient lighting due to the sunlight coming in from the window, etc.), producing the waveforms in Figures 5.41 and 5.42. Additionally, the system is clocked at one-tenth of full speed (10kS/s) in order to generate longer peak values that are easier to read off an oscilloscope (though this will also increase the exposure time and change voltage readings accordingly). Nevertheless, Figure 5.41 shows distinct voltage levels for the one pixel in each row that is covered by a top-metal

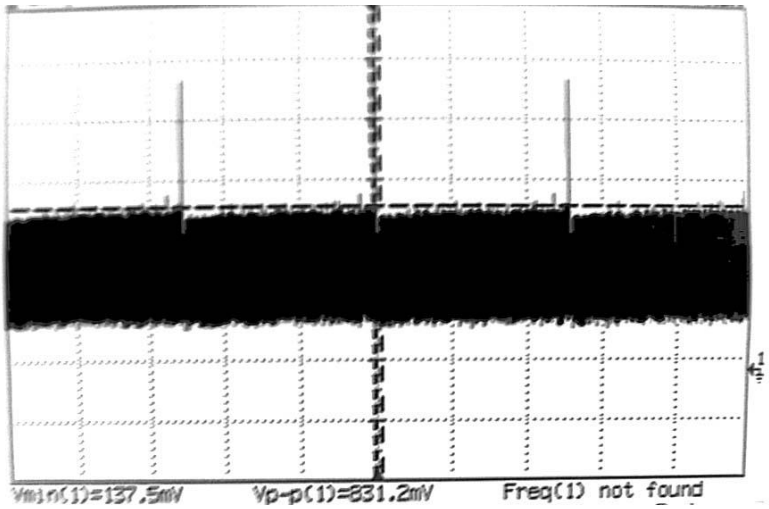

Figure 5.41: Oscilloscope output of PGA under ambient light. In this condition, the long-exposure imager is placed (without lens) under room lighting. 3 rows are shown here, as denoted by the single spikes of voltage (the reading from the one column that is covered by a top-metal light shield in layout). Further details are given in Figure 5.42.

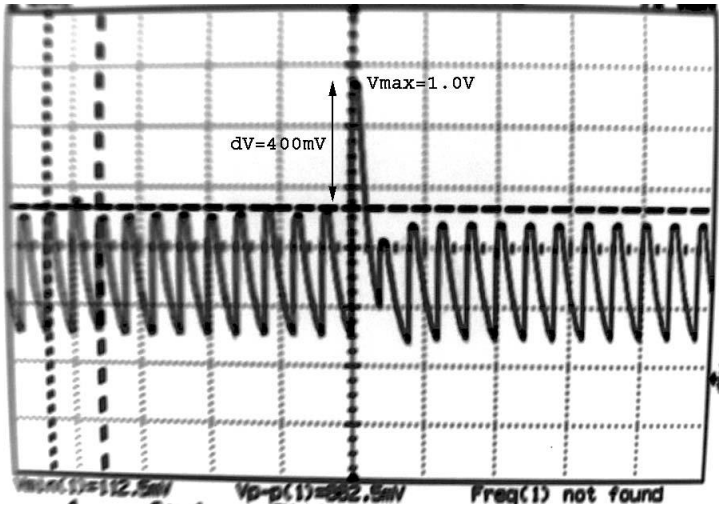

Figure 5.42: Oscilloscope output of PGA under ambient light (same as Figure 5.41, but zoomed in to individual samples). As shown, the output is held at one value for one quarter of the sample period, and appears to slew once the output is left to float for the remaining three quarters of the period. The “spike” (of one sample being higher than the others) corresponds to the one column per row that is covered by a metal light shield. There is about a 400mV drop between the shielded (dark) sample and the pixels exposed to ambient light.

light shield compared to the other 127 pixels per row (which should all produce the same voltage reading). Closer inspection (Figure 5.42) verifies that the PGA output behaves as expected (and shown in Figure 4.20, node (E)): an amplified voltage is held constant for one quarter of a sample period, and then is left to float for the other three-quarters. However, an unexpected phenomenon that is not predicted by simulation is the voltage drop once the node is left to float. This is probably the effect of parasitic diodes (including the giant diode on the output pad) on the output node acting as weak photodiodes discharging the output value. This is supported by the fact that when a laser is aimed onto the array, these voltage drops become steeper.

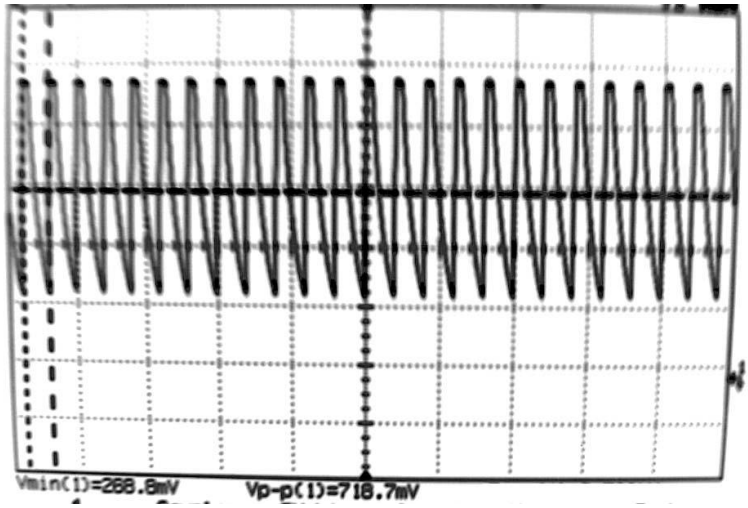

Figure 5.43: Oscilloscope output of PGA in darkness. For this case, a sheet of white paper was placed over the chip. All samples now go up to 1V, but as mentioned before, since the gain is erroneously set to 4, this could map to any array-output voltage greater than 250mV.

For comparison purposes, Figure 5.43 shows the PGA output when the imager is covered by a sheet of white copier paper. As is expected, each sample voltage is higher (less photocurrent) than it is under ambient light. And for verification that the PGA output indeed corresponds to the series of analog voltage samples coming out of the array, Figure 5.44 shows the output at a row-boundary when the right half of the array is covered by a sheet of white paper (and the left half left under ambient light): the end of a row gives lower voltages corresponding to the ambient light intensity, while the covered beginning of a row gives high voltage readings. Had the gain not been set incorrectly, some differentiation between the “dark” (metal light shield) pixel and the pixels under only a sheet of white paper would probably be seen.

Finally, it is interesting to note that when light is directed onto the imager, the “peak” values

do not go much lower than the ambient ones ($V_{min,peak}$ was only about 500mV), most likely due to pixel saturation occuring at lower light intensities because of the slower sample rate.

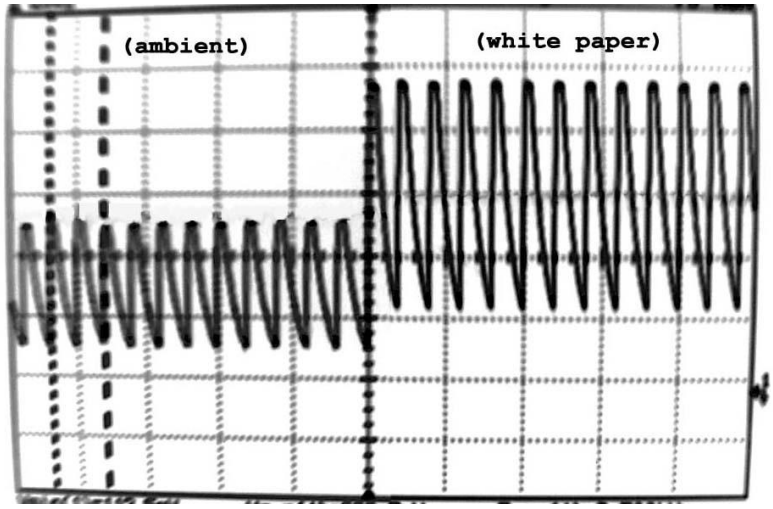

Figure 5.44: Oscilloscope output of PGA when half the imager is covered by a sheet of white paper. As expected (intuitively), the half exposed to ambient light gives lower voltage readings than the half that is covered by paper.

Ultimately though, because of the erroneous gain setting (and the subsequent reduction in output range where any analog array output between 250mV and 500mV gets mapped to 1.0V at the output of the PGA), the PGA output signals are not connected to LabVIEW for image processing.

ADC

The desired 8-bit charge-redistribution ADC for this system is not included on the test-chip, and is not tested. It has been previously shown to work in silicon [38]; and a number of applications have independently been designed and simulated extensively showing that the PGA described above will drive the ADC properly to produce eight parallel bits per sample. Furthermore, since this is a circuit block that has been designed by someone else and tested and characterized extensively by the designer, I am confident that an eventual camera system with this ADC at the output will function properly.

Chapter 6

Future Work and Conclusions

As demonstrated, the designed CMOS imager chip works as expected, though its performance is slightly lacking in that it produces noisy, blurry images. However, the lack of performance may be attributed to a number of factors:

- An unmodified "standard" 0.25μm CMOS process was used, with none of the imager-specific optimizations to junction depths, diffusion doping levels, or transistor oxide thickness that are widely used in today's commercially-available CMOS imagers. Research has shown that without these optimizations, an imager is subject to higher dark currents (and subsequently, reduced optical dynamic range), and is more susceptible to cross-talk/blooming effects (which may manifest as blurriness) [34], [42].

- No attempt at on-chip noise cancellation was made (but according to the literature, it should be quite simple to add correlated double sampling to this system)

- The aggressive low-power design may have traded off too much performance.

Future implementations of this imager should address these issues to improve image quality. First, as is mentioned in the literature, a CMOS imager fabricated in a 0.25μm (or smaller) process without imager-specific optimizations is very unlikely in practice for high-quality imaging [34], [43]. In fact, some have even argued that process modifications are mandatory to maintain image quality comparable to CCD sensors [42]. If possible, a future fabrication of this camera system should use a slightly modified process that includes some of the following properties currently used in 0.25μm and 0.18μm CMOS imagers:

- unsilicided deep-junction photodiodes optimized to reduce capacitance, reduce cross-talk,

and increase the quantum efficiency in the visible spectrum, resulting in improved photosensitivity and dynamic range [34], [42]

- SiON (instead of SiO_2) oxide layers to increase light transmission [42]
- adjusted threshold voltages to reduce leakage (high V_T for reset transistor) and increase voltage swing (low V_T for source-follower) within a pixel [43], [42]
- using thick-oxide transistors inside each pixel to reduce leakage [34], [42]

Of these options, the last two (thick-oxide switches and multi-V_T implementation) are the easiest next-step improvements as they are currently available in many unmodified standard CMOS processes. Second, correlated double sampling (CDS) should be included in the system's readout circuits (and the appropriate control signals added) to reduce noise, as described in [14] (among others). CDS is a widely-used technique that was not included in this project for reasons of imager control complexity. Third, a higher supply voltage can be used to increase the output range of a single pixel. While it may not seem like much, it is estimated that increasing the output range of a pixel by 500mV can improve its dynamic range by 5-10dB (depending on the specific pixel architecture).

In practice, however, there are two classes of work to be done in future versions of this chip: (1) short-term improvements, and (2) optional feature enhancements.

6.1 Short-Term Improvements

First, these short-term improvements should not affect the camera performance in any way:

- System integration: correct the gain-setting error on the PGA, and then test it with the ADC, analyzing only an 8-bit digital output. Currently, this would involve bringing the ADC onto the final test chip (as of now, the final chip only has a PGA, and even though TC2 had an ADC on it, it was never tested since errors in the imager array were discovered).
- After that, bring the ADC timing signals on-chip (using digital logic to generate them). Currently, the three signals that drive the ADC (clock, SAR, and RESET) are generated externally. This will save die area (three less pads), and bring the system one step closer to being fully autonomous.

- Bring the analog-mux pre-charge bias generation on-chip. Currently, it is a DC voltage originating off-chip (so that it can be adjusted manually for optimal performance). Ideally, it will be generated on-chip as a fraction of the supply voltage, consuming minimal additional static power.

Second, a number of essential design modifications to improve the camera's image quality could be implemented easily in future revisions:

- Use double-gate-oxide transistors for the in-pixel reset and row-select switches, as this is the only available process modification of those mentioned in [34] and [42] that is available in the current process. Alternatively, another CMOS process with imaging optimizations can be used.

- Add a correlated double sampling readout scheme to the column circuits. This effectively samples each pixel twice – once in reset, and once after exposure – and sends only the difference of these signals to be processed, reducing fixed-pattern and flicker noise. Minor changes to the design will have to be made, and additional control-signal complexity added, but the performance gain will be well worth it.

6.2 Longer-Term Enhancements

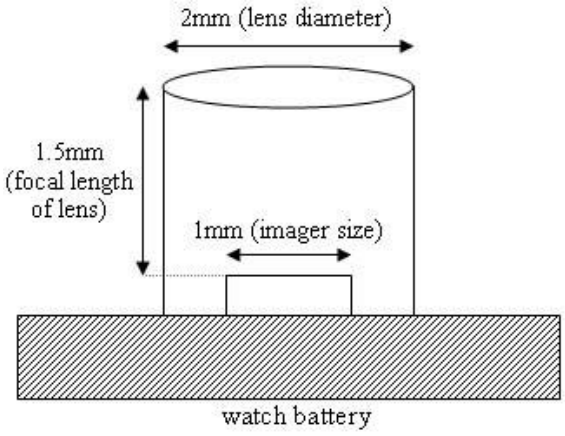

Figure 6.1: Cross section of final (button-sized) camera system; mock-up is shown in Figure 6.2. Eventually, far off into the future, the chip containing the imager will also contain a communication module to be used for receiving instructions and transmitting image data off-chip.

Finally, a number of "bells and whistles" can be added further down the road as this camera system nears its realization as a networked sensor:

- Incorporate micro-lenses on top of each pixel to improve the effective fill factor. This will help when capturing images in relatively low light conditions.
- Add row/address registers and a more complex timing scheme to make sub-array sampling possible, for the times when one is not interested in a full 128×128 resolution or is interested only in a small portion of the field of view. In theory, it should be possible to pre-load address values into the row decoder and output analog mux that will limit its read-out to a small subsection of the array.
- A wireless communication module (i.e., a radio) can be added on the same chip as the imager to allow for the autonomous network organization and communication between the imager and whoever is interested in its data. This is the final piece of the high-level "camera on a button" idea that motivated this project.

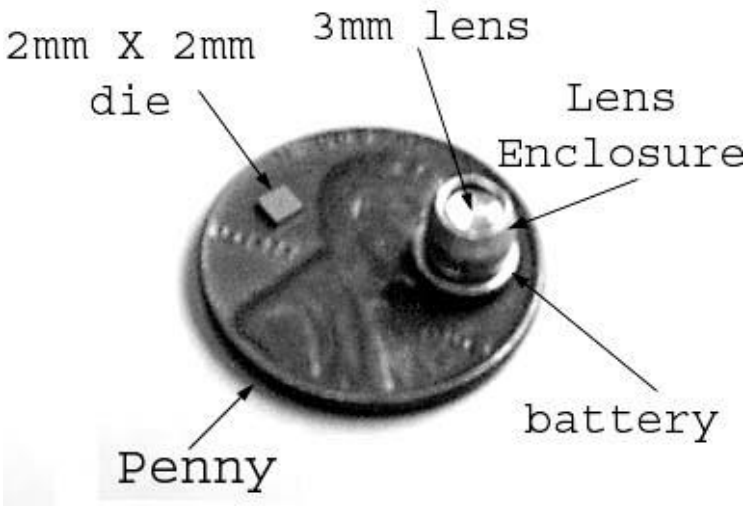

Figure 6.2: Mock-up of the final camera system as envisioned, in comparison to the size of a penny. While this figure shows a 3mm lens casing for convenience, eventually, a 2mm lens (with a much shorter focal length) and casing can be used in order to further reduce size and increase field-of-view.

In the end, the camera is envisioned to be part of a button-sized autonomous system, as shown in Figures 6.1 (in cross-section) and 6.2 (macro-level view). Though the imager chip presented

in this report does not implement this entirely, a necessary first step toward this goal has been achieved: a low-power CMOS system has been shown to be possible, and with a few modifications and additional work, this idea can definitely come to fruition.

6.3 Conclusions

A 10μJ-per-frame 128×128 CMOS active-pixel imager in a standard 0.25μm process is presented in this paper. This system is designed for low-power operation at the expense of image quality. And as demonstrated, considerable image quality has been given up; but this can mostly be attributed to the use of an unaltered process. Nevertheless, the images that it generates are recognizable, and this level of image quality can be tolerated as-is; or perhaps even improved if the data is sent to a DSP for post-processing. Using an imager-optimized process in the future can further improve the image quality.

For comparison purposes, a commercially-available "ultra-low power" CMOS imager on the market today (as of October 2003) consumes on the order of 20mW of static power in capturing images. It uses a supply voltage of 2.8V, with a maximum frame-rate of 30fps, using a rolling-shutter architecture and column-parallel ADCs on a 3.3mm lens [49]. Considering that the imager designed here consumes two orders of magnitude less power at its full sampling rate and still produces recognizable still images, the future is very bright! And while this camera will not win any awards for crisp photography, it is definitely adequate for less stringent applications that only require object identification or recognition and are most concerned with power consumption and portability.

Appendix - Using the Imager Chip

This information is most useful for members of Kris Pister's research group (Berkeley Sensor and Actuator Center, 471 Cory) who would be interested in (1) improving or re-using this design in another chip, and/or (2) packaging and using the additional fabricated chips that were not used in testing. First, the test board configuration is given, followed by a detailed functional description of the chip's pinouts. Second, the LabVIEW interface is explained.

A.1 Board Configuration

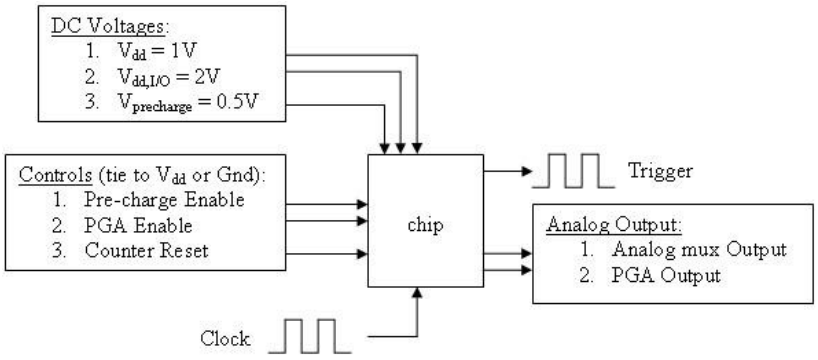

Figure A.1: Flow chart illustrating the connections to the imager chip from external electronic components as used in testing. The only input driver to this chip is the global clock.

In the final testing phase of this project, the chip was placed in a standard 24-pin package available in the Microlab and covered with a clear plastic coverslip. It was then put into a breadboard and connected to the following external components:

- **[INPUT] Power Supplies**: Three separate V_{dd} pins are on the chip (more will be explained later); however, they all can be connected to the same 1.0V supply. Additionally, a 2.0V $V_{dd,I/O}$ supply is needed for the trigger signal.

- **[INPUT] Global Clock**: A function generator was used to generate a square wave between 0-1V with a duty cycle of 50%. As mentioned earlier, this clock's frequency should be four times the desired imager sampling rate (e.g., a 400kHz global clock will produce a 100kS/s sampling rate).

- **[INPUT] Control Signals**: For debugging purposes, a number of control signals are connected to output pads on the chip, allowing me to selectively enable/disable the charge-sharing-avoidance scheme, the PGA, and the entire imager. These signals just need to be wired to either V_{dd} or ground depending on what configuration you want. More details will follow in the Chip Pinouts section.

- **[INPUT] Analog Multiplexor Pre-charge Level (bias)**: If the charge-sharing-avoidance scheme is enabled, then this voltage must be supplied. This is the voltage level at which the output of the analog mux gets pre-charged in each sample period. During read-out, the pixel-sampled voltage will discharge this node from its pre-charged value, so in theory, this voltage can be anything greater than the maximum pixel output voltage. For the tests in this report, 500mV was used *(the pixel's $V_{out,max} = 450mV$)* mainly because I was able to use a resistive divider network to provide all three of the 2.0V, 1.0V, and 500mV bias voltages from a single 2.0V power supply.

- **[INPUT] Output Buffer Biasing**: Brian Leibowitz's `pfollow_strong_open` (available in previous jupiter run directories) PMOS source-follower is used to buffer the analog multiplexor output. Referring to the schematic in Figure A.2, V_{pad} is the on-chip pin, and a bias resistor R_{bias} must be connected between this pin and a large supply V^+. In the reported tests, V^+ was chosen at 9.0V, and R_{bias} at 81.2kΩ, leading to an output swing between 750mV and 1.2V.

- **[OUTPUT] Digital Trigger**: This is effectively a clock signal between 0-2.0V with a period equal to the imager's frame-rate. For instance, at 100kS/s, this will be a square wave at a frequency of 6.1Hz. This signal is meant to be used as a trigger for an oscilloscope reading (rising-edge-trigger) and for LabVIEW's data acquisition (falling-edge-trigger).

- **[OUTPUT] Analog Output**: If the PGA is disabled, then this pin provides the output signal – it will contain a series of 16384 output voltage levels (though they may not be perfectly

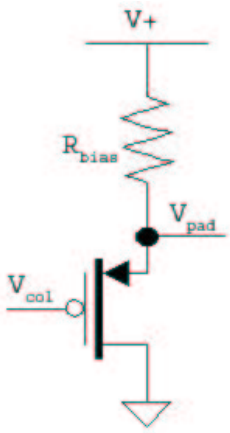

Figure A.2: Schematic of Brian Leibowitz's PMOS source-folower used as an on-chip analog buffer. For the purposes of connecting this chip to a test board, a bias resistor R_{bias} and a large voltage V^+ can be chosen to provide any desired bias condition.

level when viewed on an oscilloscope) corresponding to the pixel-sampled intensities. If the PGA is enabled, then this output pin should be ignored.

- **[OUTPUT] Amplified Analog Output**: If the PGA is enabled, then this pin will contain the series of output voltage levels corresponding to the 16384 pixel intensities. One thing to keep in mind is that these signals are delayed relative to the unamplified output by 3/4 of a sample period, as shown in Figure 4.17, and might throw off the timing of the trigger signal.

A high-level diagram of these connections is shown in Figure A.1. More detailed circuit set-up information can be found in the simulation schematics, located on the BSAC file-server at `~cadence/national/.../j33_jchoy/*sim`.

A.1.1 Chip Pinouts

The top-level layout of the final test chip is shown in Figure A.3, and the pins (most of which have been described above) are numbered as follows:

1. **Pre-Charge Enable**: Connect to V_{dd} to enable the charge-sharing-avoidance pre-charge scheme (strongly recommended). Connect to ground to disable it.

2. **Global Clock**

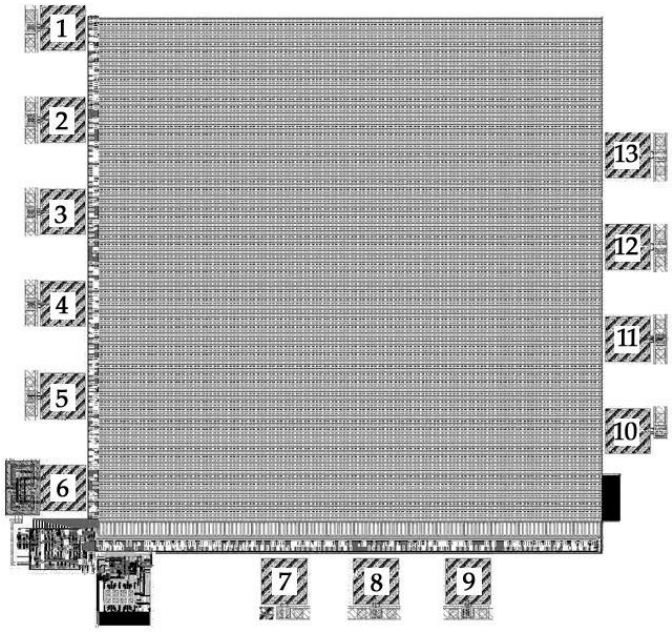

Figure A.3: Top-level layout of the imager chip used in final testing, with the pads/pins labeled for reference.

3. **Global Reset**: Connect to V_{dd} to disable the on-chip counter (and *not* generate any image data); connect to ground to enable the counter (recommended).

4. **PGA Enable**: Connect to V_{dd} to enable the PGA, disable the unamplified analog mux output, and read out the PGA Output signal; connect to ground to disable the PGA and use the buffered analog mux output.

5. $\mathbf{V}_{dd,I/O}$: Should be connected to a 2.0V supply, or whatever rail voltage you want for the trigger signal. *(This can even be the same 1.0V used for other V_{dd} pins, if so desired.*

6. **Digital Trigger**

7. **PGA Output**: Amplified analog output. Only valid if **PGA Enable** is high.

8. **Analog Output**: Buffered output from analog multiplexor. Both the pfollow bias resistor and the output probe should be connected to this node. Only valid of **PGA Enable** is low.

9. **Pre-Charge Bias**: Voltage level at which to pre-charge the output of the analog multiplexor during charge-sharing-avoidance in each sample period. I had this at 500mV for my testing.

10. $V_{dda,1}$: Should be connected to 1.0V; this V_{dd} is connected only to the PGA.

11. **Ground**

12. V_{ddd}: Should be connected to 1.0V; this is the digital supply voltage, and is connected only to the on-chip counter, control logic, decoder, and analog multiplexor.

13. $V_{dda,0}$: Should be connected to 1.0V; this is the analog supply voltage for the pixel array and biasing. Separate supplies were used to facilitate separate power measurements, but really, all three 1.0V V_{dd} signals can be connected to the same node.

A.2 LabVIEW Interface

Once the test board is configured as described above, its outputs can be connected to a computer for image generation. One way to do this is with LabVIEW. In order for this to work, you will need the following:

- A Laptop Computer running a 32-bit Windows operating system (available for University of California EECS students in 395 Cory)
- LabVIEW 6 Software (copy available in 471 Cory)
- LabVIEW DAQ-6062E PCMCIA interface (usually in 484 Cory)
- The wired "breaker box" that connects to the DAQ-6062E (usually in 484 Cory)
- The Camera.vi LabVIEW program that I used, or a variant thereof (the original – which is detailed in this Appendix – is available in compressed format on-line at http://www.ocf.berkeley.edu/~jonc/project/camera.vi.gz).

A.2.1 Data Acquisition Hardware

The LabVIEW breaker box should be configured as follows:

- **Digital Trigger A (TRIGA)** should be connected to the **Digital Trigger** pin of the chip.
- The associated digital ground **(DGND)** should be connected to ground.
- Pick an analog input channel **(ACH*X*)**. Connect its positive wire to either the buffered **Analog Output** (Pin #8 above), or the **PGA Output** (Pin #7 above) on the chip, depending on **PGA Enable**.

- Connect the associated ground **(AIGND)** wire to ground. *NOTE: currently, the AIGND wire for each channel is not wired to the AIGND port inside the box, but instead, to the channel's associated negative input terminal. For instance, the "GND" for ACH3 is really connected to ACH11, and the "GND" for ACH6 connected to ACH14. This is done so that the default Analog Input Config mode in LabVIEW would not have to be changed from differential to single-ended mode.*

All other connections in the breaker box can be ignored.

A.2.2 Camera.vi

Once everything is connected properly and LabVIEW is running, open the `Camera.vi` program. You should see a screen that looks like Figure A.4. The controls and indicators in each region are

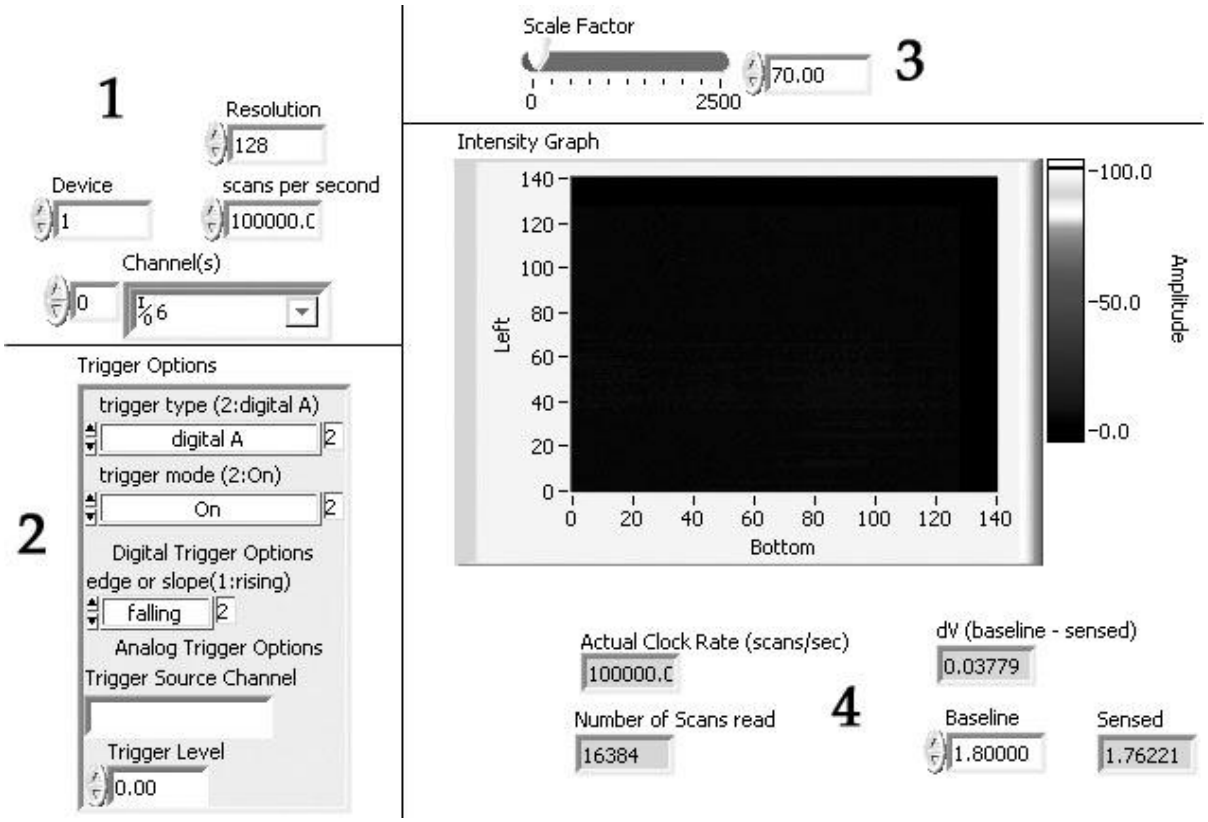

Figure A.4: Annotated version of Figure 5.22, used to explain the controls and knobs that can be adjusted in the Camera.vi LabVIEW program.

as follows:

- **Resolution [1]**: This value determines the number of data points for LabVIEW to sample when it sees a falling edge on its trigger signal. A value of n translates to n^2 total samples.

- **scans per second [1]**: This tells LabVIEW how long to wait between samples, and should be set to the chip's sample rate ($1/4$ of the global clock frequency).

- **Device/Channel [1]**: The Device should be left alone, unless you have multiple data acquisition devices connected to your computer. The Channel setting corresponds to the analog input channel that is connected to the chip's analog output. Make sure this is consistent with the analog input channel chosen when the breaker box was connected to the breadboard.

- **trigger type [2]**: Set to digital A, since the chip's trigger signal is connected to TRIGA of the breaker box.

- **trigger mode [2]**: Set this to "On" to enable triggering.

- **edge or slope [2]**: Set this to falling edge. At the trigger signal's falling edge, the on-chip counter is back to zero, and the imager's pixel array is back at row 0 and column 0; it's at this point that you want LabVIEW to start recording the next 16384 samples at the sample rate.

- **Analog Trigger Options [2]**: leave this alone; the data acquisition is triggered on a digital signal, not an analog one.

- **Baseline [4]**: This is a hack to try to simulate CDS in LabVIEW. Instead of processing the analog voltage that is received from the chip, LabVIEW will take the difference between the baseline *(think of it as a guess of what the pixel's reset level would produce)* and the sampled voltage *(from the chip)* and display intensity based on this derived value.

- **Scale Factor[3]**: This is a crude attempt at contrast adjustment, and simulates a PGA: the difference derived from the baseline and sampled value is multiplied by this scale factor before being translated into an intensity.

- **Actual Clock Rate, Number of Scans read, Sensed Voltage [4]**: These are indicators to show you a summary of what happened during LabVIEW's last analog data acquisition. They're more of a sanity check than anything else.

On another note, there are two possble ways to run this program: (1) single step-through, which is the equivalent of capturing an image by pressing a button on a camera, or (2) continuous run mode, which can act like a video recorder (but at much slower rates, of course). Either should work fine so long as the lens is focused properly and there is adequate lighting. Finally, the underlying details of this LabVIEW interface is not explained here, since most of it can be understood by using LabVIEW's help features (hit `Ctl-E` to pop open the details, and then `Ctl-H` to enable the help).

A.3 Additional Help

If you have any questions about this, you can reach me via e-mail at `jchoy@cal.berkeley.edu` and I can try to answer any questions you might have (assuming I still remember what I did).

Bibliography

[1] S. Kleinfelder, S. Lim, X. Liu, and A. El Gamal, "A 10000 Frames/s CMOS Digital Pixel Sensor," *IEEE J. Solid-State Circuits*, vol. 36, pp. 2049-2059, December 2001.

[2] S. Mendis, S. Kennedy, and E. Fossum, "A 128 x 128 CMOS Active Pixel Image Sensor for Highly Integrated Imaging Systems," *IEEE IEDM Tech. Dig.*, 1993, pp. 583-586.

[3] D. Yang, B. Fowler, and A. El Gamal, "A 128x128 CMOS Area Image Sensor with Multiplexed Pixel Level A/D Conversion," *IEEE 1996 Custom Integrated Circuits Conference*, San Diego, CA, May 1996.

[4] C. H. Aw and B. Wooley, "A 128x128-Pixel Standard-CMOS Image Sensor with Electronic Shutter," *IEEE J. Solid State Circuits*, vol. 31, pp. 1922-1930, December 1996.

[5] R. H. Nixon, S. E. Kennedy, B. Pain, C.O. Staller, and E. R. Fossum, "256x256 Active Pixel Sensor Camera-on-a-Chip," *IEEE J. Solid-State Circuits*, vol. 31, pp. 2046-2050, December 1996.

[6] I. Fujimori, C. Wang, and C. Sodini, "A 256x256 CMOS Differential Passive Pixel Imager with FPN Reduction Techniques," *IEEE J. Solid-State Circuits*, vol. 35, pp. 2031-2037, December 2000.

[7] K. Cho, A. Krymski, and E. Fossum, "A 3-Pin 1.5V 550μW 176x144 Self-Clocked CMOS Active Pixel Image Sensor," *International Symposium on Low Power Electronics Design*, Huntington Beach, CA, August 2001.

[8] M. Kyomasu, "A New MOS Imager Using Photodiode as Current Source," *IEEE J. Solid-State Circuits*, vol. 26, pp. 1116-1122, August 1991.

[9] N. Ricquier and B. Dierickx, "Active Pixel CMOS Image Sensor with On-Chip Non-Uniformity Correction," *IEEE Workshop on Charge-Coupled Devices and Advanced Image Sensors*, Dana Point, CA, April 1995.

[10] H. Tian, X. Liu, S. Kim, S. Kleinfelder, and A. El Gamal, "Active Pixel Sensors Fabricated in a Standard 0.18 μm CMOS Technology," *Proceedings of SPIE*, vol. 4306 (2001), pp. 441-449.

[11] H. Tian, B. Fowler, and A. El Gamal, "Analysis of Temporal Noise in CMOS Photodiode Active Pixel Sensor," *IEEE J. Solid-State Circuits*, vol. 36, pp. 92-101, January 2001.

[12] M. Gottardi and W. Yang, "A CCD/CMOS Image Sensor Array with Integrated A/D Conversion," *1997 IEEE International Symposium on Circuits and Systems*, Hong Kong, June 1997.

[13] S. Mendis, S. Kemeny, R. Gee, B. Pain, C. Staller, Q. Kim, and E. Fossum, "CMOS Active Pixel Image Sensors for Highly Integrated Imaging Systems," *IEEE J. Solid-State Circuits*, vol. 32, pp. 187-197, February 1997.

[14] E. Fossum, "CMOS Image Sensors: Electronic Camera-On-A-Chip," *IEEE Transactions on Electron Devices*, vol. 44, pp. 1689-1698, October 1997.

[15] B. Carlson, "Comparison of Modern CCD and CMOS Image Sensor Technologies and Systems for Low Resolution Imaging," *Proceedings of IEEE Sensors 2002, Piscataway, NJ*, vol. 1, pp. 171-176.

[16] J. Goy, B. Courtois, J. M. Karam, and F. Pressecq, "Design and Test of an Active Pixel Sensor (APS) for Space Applications," *Proceedings of SPIE*, vol. 4306 (2001), pp. 93-99.

[17] Y. Audet and G. Chapman, "Design of a Self-Correcting Active Pixel Sensor," *Proceedings of the 2001 IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems*, San Francisco, CA, October 2001.

[18] E. Fossum, "Digital Camera System on a Chip," *IEEE Micro*, pp. 8-14, May-June 1998.

[19] M. Hillebrand, N. Stevanovic, B. J. Hosticka, J. E. Santos Conde, A. Teuner, and M. Schwarz, "High Speed Camera System Using a CMOS Image Sensor," *Proceedings of the IEEE Intelligent Vehicles Symposium 2000*, Dearborn, MI, October 2000.

[20] W. C. Dash and R. Newman, “Intrinsic Optical Absorption in Single-Crystal Germanium and Silicon at 77K and 300K,” *Physical Review*, vol. 99, August 1955, pp. 1151-1155.

[21] E. Chou, A. Budrys, and K. Cham, “Low Power Salient Integration Mode Image Sensor with a Low voltage Mixed-Signal Readout Architecture,” *Proceedings of the 1998 International Symposium on Low Power Electronics and Design, New York*, pp. 42-47.

[22] B. Pain, G. Ying, B. Olson, T. Shaw, M. Ortiz, J. Heynssens, C. Wrigley, and C. Ho, “A Low-Power Digital Camera-on-a-Chip Implemented in CMOS Active Pixel Approach,” *Proceedings of the 2001 International Conference on VLSI Design, Los Alamitos, CA*, pp. 26-31.

[23] S. Tanner, S. Lauxtermann, M. Waeny, M. Willemin, N. Blanc, J. Grupp, R. Dinger, and E. Doering, “Low-Power Digital Image Sensor for Still Picture Image Acquisition,” *Proceedings of SPIE*, vol. 4036 (2001), pp. 358-365.

[24] S. Kemeny, R. Panicacci, B. Pain, L. Matthies, and E. Fossum, “Multiresolution Image Sensor,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 4, pp. 575-583, August 1997.

[25] I. Scherback and O. Yadid-Pecht, “Photoresponse Analysis and Pixel Shape Optimization for CMOS Active Pixel Sensors,” *IEEE Transactions on Electron Devices*, vol. 50, pp. 12-18, January 2003.

[26] R. Ooi, T. Hamamoto, T. Naemura, and K. Aizawa, “Pixel Independent Random Access Image Sensor for Real Time Image-Based Rendering System,” *Proceedings of the 2001 International Conference in Image Processing, Piscataway, NJ*, vol. 2, pp. 193-196.

[27] O. Yadid-Pecht, R. Ginosar, and Y. S. Diamand, “A Random Access Photodiode Array for Intelligent Image Capture,” *IEEE Transactions on Electron Devices*, vol. 38, pp. 1772-1780, August 1991.

[28] D. Scheffer, B. Dierickx, and G. Meynants, “Random Addressable 2048x2048 Active Pixel Image Sensor,” *IEEE Transactions on Electron Devices*, vol. 44, pp. 1716-1720, October 1997.

[29] B. Dierickx, D. Scheffer, G. Meynants, W. Ogiers, and J. Vlummens, “Random Addressable Active Pixel Image Sensors,” *Proceedings of SPIE*, vol. 2950 (1996), pp. 2-7.

[30] M. A. Abdallah, E. Dubaric, H.E. Nilsson, C. Frojdh, and C. S. Petersson, “A Scintillator-coated Phototransistor Pixel Sensor with Dark Current Cancellation,”

[31] I. Koren, G. Chapman, and Z. Koren, “A Self-Correcting Active Pixel Camera,” *Proceedings of the 2000 IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems, Los Alamitos, CA*, pp. 55-64.

[32] S. Ma and L. Chen, “A Single Chip CMOS APS Camera with Direct Frame Difference Output,” *IEEE J. Solid-State Circuits*, vol. 34, pp. 1415-1418, October 1999.

[33] M. Chi, “Technologies for High Performance CMOS Active Pixel Imaging System-on-a-chip,” *Proceedings of the 1998 International Conference on Solid-State and Integrated Circuit Technology, Piscataway, NJ*, pp. 180-183.

[34] H. Wong, “Technology and Device Scaling Considerations for CMOS Imagers,” *IEEE Transactions on Electron Devices*, vol. 43, pp. 2131-2142, December 1996.

[35] C. Xu and M. Chan, “The Approach to Rail-to-rail CMOS Active Pixel Sensor for Portable Applications,” *Proceedings of IEEE Region 10 International Conference on Electrical and Electronic Technology*, vol. 2, pp. 834-837.

[36] O. Yadid-Pecht and E. Fossum, “Wide Intrascene Dynamic Range CMOS APS Using Dual Sampling,” *IEEE Transactions on Electron Devices*, vol. 44, pp. 1721-1723, October 1997.

[37] J. Solhusvik, C. Cavadore, F. X. Audoux, N. Verdier, J. Farre, O. Saint-Pe, R. Davancens, and J. P. David, “Recent Experimental Results from a CMOS Active Pixel Image Sensor with Photodiode and Photogate Pixels,” *Proceedings of SPIE - the International Society for Optical Engineering*, vol. 2950 (1996), pp. 18-24.

[38] M. Scott, B. Boser, and K. Pister, “An Ultra-Low Power ADC for Distributed Sensor Networks,” *Proceedings of the 28th European Solid-State Circuits Conference (2002), Florence, Italy*, pp. 255-258.

[39] C. Koch, S.-B. Park, T.J. Ellis, and A. Georgiadis, “Illumination technique for optical dynamic range compression and offset reduction,” *British Machine Vision Conference*, 2001.

[40] E. Fossum, “CMOS Active Pixel Image Sensors,” *IEEE Trans. Electron Devices*, October 1997.

[41] G. Yang, O. Yadid-Pecht, C. Wrigley, and B. Pain, “A Snap-Shot CMOS Active Pixel Imager for Low-Noise, High-Speed Imaging,” *Proceedings of IEDM, San Francisco, CA*, December 6-9, 1998.

[42] A. El Gamal, “Trends in Image Sensor Technology and Design,” *Proceedings of IEDM 2002*, December 2002.

[43] E. Roca, S. Espejo, R. Dominguez-Castro, and A. Rodriguez-Vazquez, “Light-Sensitive Devices in CMOS,” *Towards the Visual Microprocessor: VLSI Design and the Use of Cellular Neural Network Universal Machines*, John Wiley & Sons, January 2002, Chapter 5.

[44] G. Holst, *CCD Arrays, Cameras, and Displays*, JCD Publishing & SPIE Optical Engineering Press, 1998.

[45] B. Razavi, *Analysis and Design of CMOS Analog Integrated Circuits*, McGraw-Hill, 2001.

[46] D. Johns and K. Martin, *Analog Integrated Circuit Design*, John Wiley & Sons, 1997.

[47] P. Gray, P. Hurst, S. Lewis, and R. Meyer, *Analysis and Design of Analog Integrated Circuits*, John Wiley & Sons, 2001.

[48] The SmartDust Project Webpage, http://robotics.eecs.berkeley.edu/~pister/SmartDust/

[49] Micrion Technology, Inc. CMOS Image Sensors, http://www.micron.com/products/imaging.