Depth-Dependent Single-Event-Effect Analysis of a 3D-Integrated Imaging Chip

By

Michael D. Hu

Thesis

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

MASTER OF SCIENCE

in

ELECTRICAL ENGINEERING

December 18th, 2021

Nashville, TN

Approved:

Robert A. Reed, Ph.D.

Michael L. Alles, Ph.D.

ACKNOWLEDGEMENTS

Some research is performed in a vacuum, but none is performed alone. First and foremost, I would like to thank Dr. Robert Reed for his guidance, support, patience throughout this work, and ability to always provide a new perspective. Secondly and with no less gratitude, I would like to thank Dr. Michael Alles for his oversight, insight, equally large amount of patience, and for always giving me an extra push. I would like to thank Dr. Brian Sierawski and Dr. Andrew Sternberg for their invaluable insights and suggestions. To Dr. Kevin Warren, thank you for your feedback and assistance. To Dr. Enxia Zhang and Mr. Mike McCurdy, thank you for your experimental assistance and expertise. Finally, thanks are due to Dr. Ron Schrimpf, whose optimism was always welcome.

I would like to express my gratitude towards the Defense Threat Reduction Agency for their support through the basic research program and to JPL for providing the funding and enthusiasm to make this work possible.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

In the ever-evolving landscape of commercial off-the-shelf (COTS) integrated circuits (ICs), there are increasing efforts to create three-dimensionally (3D) integrated systems: devices where multiple structures, sometimes of different technologies, are intimately integrated. The radiation responses of such intimately integrated devices, particularly of cutting-edge imaging chips, is of interest in spacecraft designs. However, radiation responses of such 3D devices are largely unexplored. Due to the complexities of such systems and proprietary protections, devices are often little more than "black-boxes" without manufacturer collaboration. This black-box nature limits the viability of COTS parts for space environment applications, as little more than characterization can be done for imagers that fly on space missions.

This work presents a method to analyze the radiation response of such black-box COTS parts. The strength of this approach is that it is generalizable; it works just as well on black-box devices as with devices where structures and functions are well known. In both cases, the approach's effectiveness is maximized when coupled with the presented "depth-profiling" approach for analysis of the imager of interest. First, background on 3D structures, black-box type devices, and imagers is presented, followed by a brief description of the 3D integrated circuit (3D-IC) device under test (DUT) - the Sony IMX400 imager. Next, the depth profiling approach is described, and an in-depth discussion of the data analysis technique is presented. The remainder of this work details a case study of the DUT, including the experimental results and conclusions drawn from analysis of those results.

In the case study, transient experimental results from different functional layers are compared. It is found that the holistic transient response of the device is dominated by the transient response of the pixel layer, and that a region in the middle device layer is vulnerable to single-event latch-up (SEL). These results are found using a combination of the depth-profiling approach and an analysis technique that are both described in this work.

3D Integrated Circuits

A 3D integrated circuit is exactly that: an integrated circuit that not only utilizes in-plane integration, but also takes advantage of the z-direction to achieve better electrical and computational performance [1]. To do so, silicon wafers or dies are stacked vertically and interconnected, often via through-silicon vias (TSVs). This greatly shortens electrical connections between fabricated devices, allowing faster data transfer and reducing necessary power consumption. There are many different integration schemes, from interposer-based integration to stacked ICs (Fig. 1).

This work will focus on the latter - 3D stacked ICs (3D-SICs) – where ICs are fabricated on different wafers before stacking. Stacking is achieved is by die-to-die bonding, with thinning of either die (or both dies) performed before or after bonding [2]. Similarly, TSV creation is possible prior[1] to or after bonding [3]. Die-to-die bonding has the advantage of increased chip yield, as individual dies will not affect the viability of an entire wafer, only the dies that were bonded.
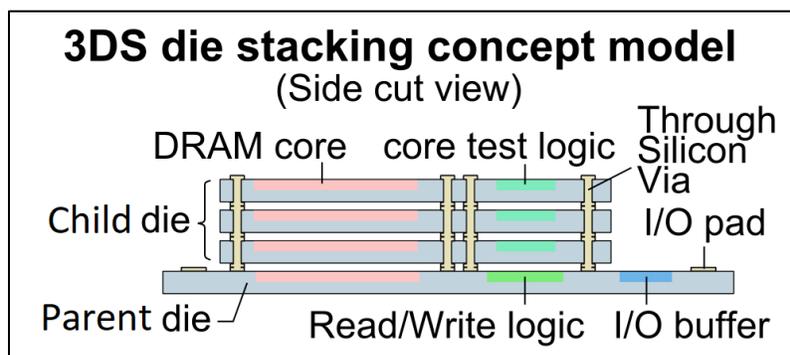
Fig. 1. A conceptual diagram of 3D IC implementation. [4]

---

[1] TSVs can be created on both dies, then connected when the dies are bonded. [3]

Black Box Devices

In electronic circuits theory, it is often necessary to design a circuit which has a desired function. However, there are several ways to create a circuit that fulfills the same purpose. A simple example is the use of two (ideal) 50-ohm resistors in place of one (ideal) 100-ohm resistor to regulate current. Thus, circuits can be abstracted as opaque (black) boxes whose function is known but contents cannot be determined. It is from this abstraction that a black box can be defined as not only a circuit, but a system of unknown components where only stimuli and responses can be observed [5]. (Any system can be treated as a black box, though it may not be sensible to do so.)

Knowledge of *what* the system should do allows one to estimate *how* the system might perform its task. This is particularly useful when a response to a given stimulus is known and estimation of response to a similar stimulus is necessary. This may also allow estimation of the effective operational lifetime of a device in a specific environment, e.g., an imager on a spacecraft.

This work focuses on a functionally black box device: an advanced CMOS imager. Its high-level purpose - *what* it does - is known: capture a photo or video of whatever may be in front of it. However, many details of this device are unknown, and indeed cannot be known without communication/interaction from the manufacturer. It is known that there is a pixel array with embedded micro-lens filters in a Bayer pattern. It is unknown what wavelengths to which the individual color filters are tuned. It is known that there is compression of the raw input, though it is unknown how many pixels are compressed (e.g., 4x4 pixel blocks or 8x8 pixel blocks). It is also unknown exactly *which* computations are being performed, and *where* the computations are physically taking place, be it on-chip or in an external processor.

Radiation Effects

Radiation is present in many environments, including space, nuclear reactors, and particle accelerator facilities. In these environments, particles such as protons, neutrons, and ions - from light ions like helium (i.e., alpha particles) to heavy ions like iron - commonly produce *soft errors*.

A soft error is an often-recoverable unintended change of state in a circuit. Soft errors are the manifestation of single-event upsets (SEUs), radiation-induced changes in a circuit's logic

state. SEUs are a type of single-event effect (SEE), which are defined as any effect on an integrated circuit resulting from a single incoming ionizing event. A single-event transient (SET) is a radiation-induced current (or voltage) at a given circuit node.

Incoming ionizing particles can cause SETs. As they enter a device, they transfer (deposit) some of their energy to the target device's constituent material. If the deposited energy is sufficient, the particle can liberate electrons from the target material, generating electron-hole pairs (EHPs) and effectively liberating charge [6]. These EHPs are new free charge carriers in the target material, contributing to increased current and as a result, generating an SET at a circuity node or even multiple nodes. A related concept is the sensitive volume, which is a volume within a device where deposition of energy can affect circuit operation. If a sensitive volume is struck by an ionizing particle, the device is likely to experience a SET.

Transistor size has decreased to the point where a single high-energy ionizing particle can pass through multiple transistors. While the particle itself may not pass through multiple transistors, its area of influence (i.e., region where it generated EHPs are collected) may encompass many neighboring transistors. This is problematic, as a single particle can now cause multiple soft errors [7].

An incoming ionizing particle with sufficient energy may also displace not only electrons, but also atoms from their lattice sites. Such vacancies, known more generally as displacement damage, introduce new allowed energy states in the semiconductor material's bandgap, changing the electrical performance of the device [8], [9].

Ion Stopping

Despite development for over a century, there is no exact ion stopping theory based on first principles of physics; to date, all theories normalize their calculations using experimental data [10]. This work does not attempt to change this fact. As an ion passes through a target material, interaction with the target's electrons and nuclei through Coulomb interaction gradually saps it energy and thus, the ion slows down. The particle's average energy loss per unit path length (dE/dx) is known as its stopping power. The energy deposited in the target material is then -dE/dx,

which can be normalized by target material density. The normalized energy deposited is known as the linear energy transfer (LET).

$$LET = -\frac{1}{\rho}\frac{dE}{dx} \qquad \text{(Equation 1)}$$

The ion loses most of its energy immediately before stopping, creating a peak in the LET vs unit path length diagram known as the Bragg peak (Fig. 2) [11].
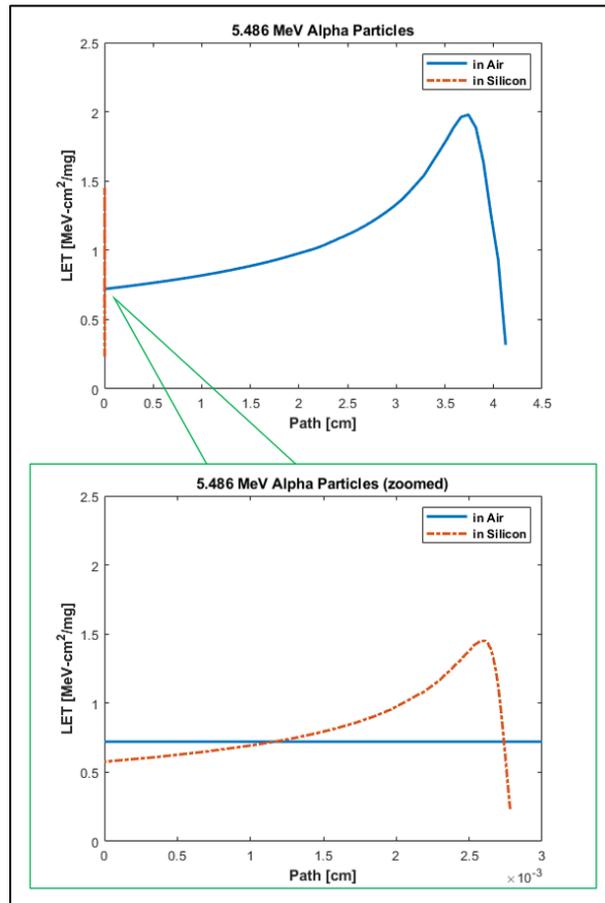
Fig. 2. LET vs path length of Alpha particles in air and silicon. The bottom is an expanded section of the top figure with the same vertical scale. The pronounced peaks are the Bragg peaks.

Imagers

Imagers convert light reflected off objects into digital images. They do this by having an array of picture elements, called pixels, that sample the light reflected at different locations. The greater the intensity of light, the greater the digitized intensity value. Pixels in both imagers and digital displays both use 3 color elements: red, green, and blue. Unlike displays, imagers are laid out in a Bayer filter pattern (Fig. 3) [12], using 1/3 the color elements as a physical display. Through a process called demosaicing [13], physical pixels are assigned their missing color band values based on their neighbors. After demosaicing, each pixel has 3 values, one for red, green[2], and blue, and this pixel representation is now the "display pixel".
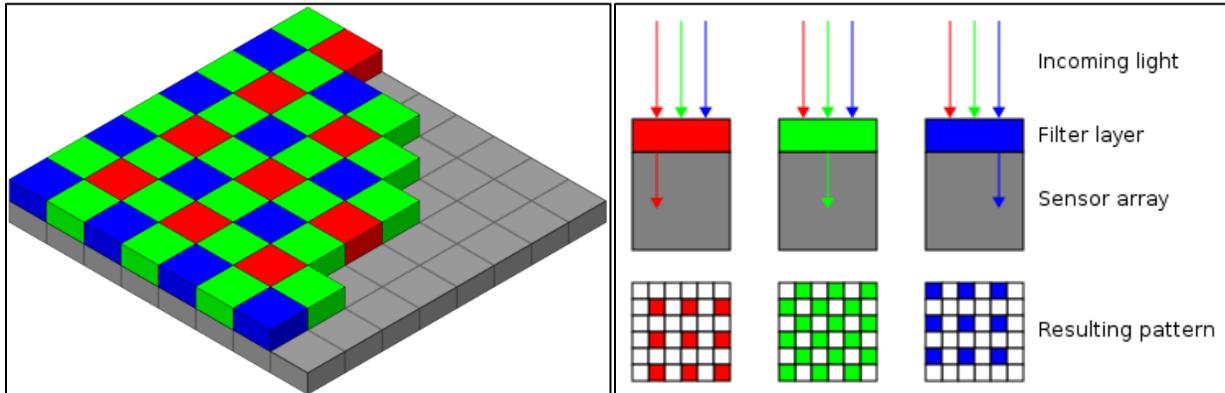
Fig. 3. The Bayer filter pattern. Display pixels have 3 bands, however, it is impractical to stack 3 photodiodes on top of each other. Instead, they are aranged in this pattern and color data from neighboring sensors is extrapolated and assigned to each detector's location, resulting in a 3-band image.

Typically, each physical pixel within an imager contains a photodiode [14] that generates an amount of charge proportional to the energy of the incoming photons. A physical color filter is placed on top, allowing only light from specific colors (wavelengths) to reach each detector. Through the demosaicing process, each physical pixel is assigned color values based on their neighbors' intensities. This produces high fidelity results when capturing photographs. However,

---

[2] There is a single value saved to represent the green band, but there are twice as many green detectors than red or blue such that the ratio of red:green:blue is 1:2:1. This is because the human eye is about twice as sensitive to the color green than either red or blue [31].

when the incoming particle exerts a region of influence on the order of multiple physical pixels, information is lost. Information about exactly which physical pixel was struck is further obstructed by demosaicing. Similarly, information about the size of the particle's region of influence may be lost, as multiple struck physical pixels may appear larger when their values are assigned to neighbors.

Image Files

Standard 24-bit true color image, often saved as ".png" files or ".jpg" files, store the data of an image in the form of a 3D matrix. This image matrix has the dimensions of row, column, and bands, of which are red (R), green (G), and blue (B) (Fig. 4).
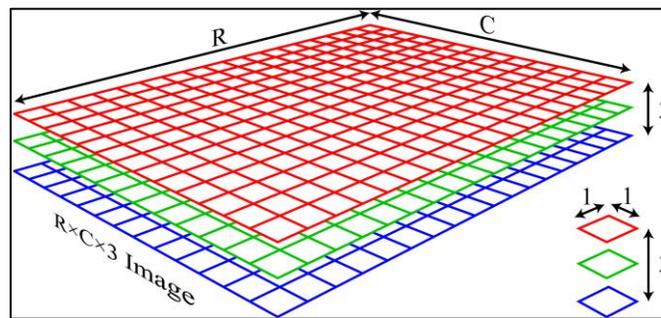
Fig. 4. An illustration of the typical data structure layout of RGB image files [15].

The planar (row and column) elements of this matrix are the display pixels. Digital displays use additive light to represent color, with the values of each band representing their relative weights. In a 24-bit true color image, each band has 8 bits, allowing valid values within the range of 0 to 255. A pixel with equal values of R, G, and B are grayscale, as equal intensity values additively create white, transitioning to black as the intensities are (equally) decreased. In general, an uncompressed image needs a minimum of $rows \times columns \times 24$ bits to be stored.

A JPEG image, the source of information in this work, compresses data from the raw image. It does so by performing a 2D Fourier transform in blocks, assigning new colors that appear to be the same as the original colors while losing an acceptable amount of information It is important to note that there may be lossy data compression of the image file itself, which may further obscure data. This is beyond the scope of this work.

CHAPTER III

DEVICE UNDER TEST: SONY IMX400

Hardware

The Sony IMX400 was one of the first smartphone camera chips to use a stacked implementation of complementary-metal-oxide-semiconductor (CMOS) device image sensors (CIS). This imaging chip's layers are shown in Fig. 5; it is a 3-layer stacked device that consists of a 90 nm active pixel sensor (APS) array layer, a 30 nm DRAM layer, and a 40 nm computational logic layer, each fabricated in their bulk CMOS nodes [16]. Fig. 6 displays a cross section of the device with bonding information. The DRAM layer is flip-bonded front-to-front with the logic layer, and its substrate is thinned. The front of the pixel layer is then bonded to the thinned backside of the DRAM layer. Each functional layer is on the order of 10 µm thick, with photodiodes in the pixel layer on the order of 5 µm, not including micro lenses or the color filter layers. The photodiodes are partially separated by relatively shallow trench oxides, as seen in Fig. 7.
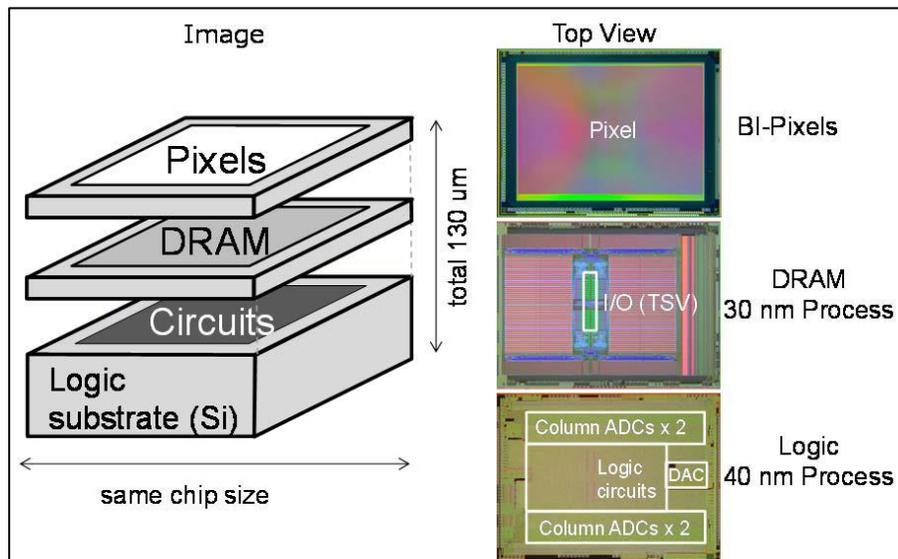
Fig. 5. Illustration of the 3D structure.
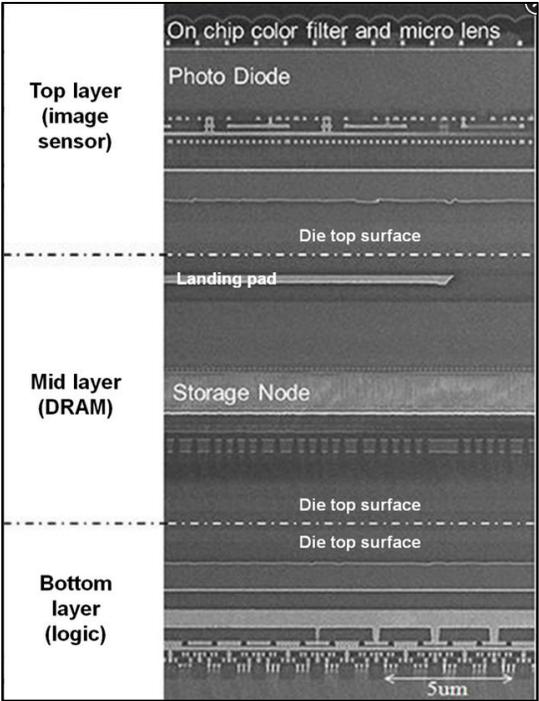
Fig. 6. A cross section of the imaging chip [15]. The cross section illustrates the flipped orientation of the top 2 layers. Micro lens layer visible at top of the image.
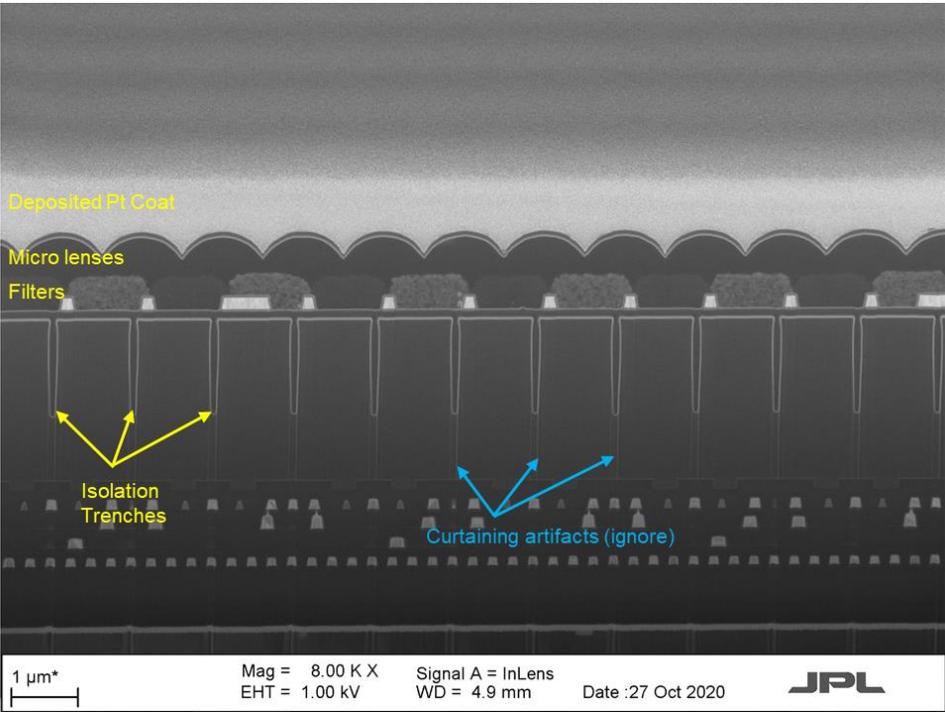
Fig. 7. An SEM of the DUT. Viewing angle of 54°. Courtesy of JPL's Analysis & Test Laboratory.

Operation

      In a traditional CIS, the pixel data is sequentially read out to the logic layer, which pre-processes and sends the data off-chip to external devices, such as a cell phone's main CPU. The serial interface to such external devices is significantly slower and thus is the limiting factor [14], so the pixel layer does not capture data during the read-out process. The IMX400 introduces a DRAM layer, which serves as a data buffer for the pixels. No longer limited by the readout speed from the logic layer, the pixel layer can capture data more often.

      In Fig. 8, a block diagram of standard operation is shown. The pixel layer captures an image that is sent to the pre-processing block, where some operations, such as black-level compensation and digital gain, are performed. The image is then sent to the DRAM for temporary storage. From the DRAM, the image is sent to the main processing block where other operations - such as defect correction, noise reduction, and scaling - are performed. Once main processing is complete, the image is queued for read out, returning to the DRAM if necessary. The packaged imager, shown in Fig. 9, communicates via the serial MIPI D-PHY standard [17] to some of Sony's Xperia-model smartphones via a ribbon cable. The imager can be controlled through Sony's built-in camera app.
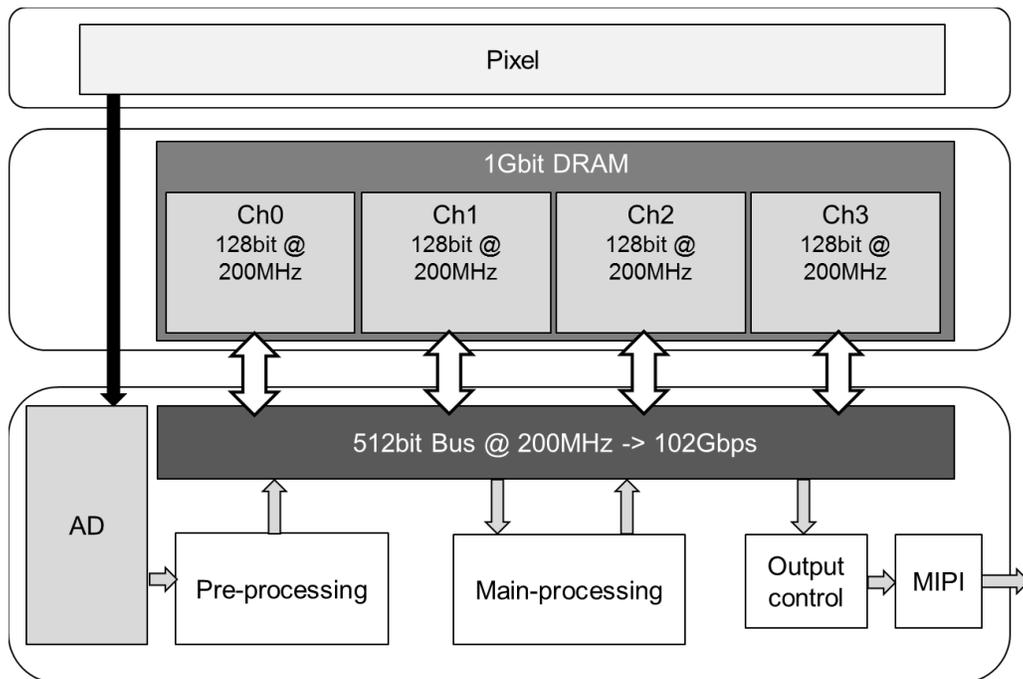
Fig. 8. IMX400 operational black diagram. [18]

Fig. 9. The packaged IMX400 chip. Photo taken by the author.

While this much information is known about the device under test (DUT), it is very much a functionally black box device. Specifics of the device, such as electrical layouts of individual pixels, DRAM cells, readouts, and logic layer are all unknown. Similarly, specifics of its image processing are unknown; it is unknown exactly which functions are performed and where they are performed. For example, it is unclear if image data is converted to a mpeg-compatible format on-chip or in the smartphone's main CPU. The amount of known information (and lack thereof) makes this an appealing device of interest for this work.

CHAPTER IV

DEPTH PROFILING AND THE DATA ANALYSIS TECHNIQUE

This section is presented as follows: First, an introduction to the idea of *depth profiling* is presented. Next, the analysis technique is presented in 3 parts: Intensity Thresholding, The Splitting of Images, and lastly, Classification of Events.

In this chapter, there will be references to "transient events". An "event" is defined as a grouping of adjacent pixels that exceed the defined threshold, and a "transient event" is defined as an event that lasts no more than 1 frame. In a sufficiently dark environment, all observed events can be assumed to be radiation-induced.

This analysis technique can be used to report statistics for a given data set. It is generalizable but must be coupled with an appropriate experiment for useful conclusions to be drawn. For example, when coupled with a depth profiling experiment, this technique reports results like: "percent of all recorded events of size Row x Column containing number N lit pixels". Furthermore, this technique may help to formulate the following example statement: "A majority of all transient events were observed when alpha particles of known energy deposited energy up to a certain depth from the surface of the DUT at a 0° angle".

This section will present an analysis technique that was originally created to analyze data generated by a CMOS imager, output in the form of an MP4 video file. However, this technique can be modified to work with almost any data recorded as matrices.

 Unless otherwise stated, the word "pixel" in this section refers to a digitally *displayed* pixel, as opposed to the physical pixel used during image capture. A "lit pixel" refers to a displayed pixel that meets or exceeds its defined threshold value.

Depth Profiling

A planar integrated circuit (IC) is simpler than a 3D-IC; the planar IC is typically a chip with one high-level purpose. For example, a planar DRAM consisting of addressing drivers, logic refreshes, read/write circuitry, etc., performs one dedicated high-level function: read and write digital data quickly. A DRAM may report erroneous data when subjected to radiation, or it may draw more power to perform the same action.

In a 3D-IC, things get more complicated: there are multiple intimately integrated functional layers, commonly performing different high-level tasks. With more material, there are more things that could happen. For example, a single incoming particle may deposit energy in multiple functional layers, causing SEEs in each layer. In a planar IC, this would impact one high-level function, but in a 3D-IC it may affect the operation of multiple interdependent layers.

Depth profiling is the creation a functional mapping of depth (and when necessary, planar location) of a given sensitive volume to its output. This can be achieved by varying energies, types of radiation (Fig. 10), incoming beam direction, and the particle/beam's angle-of-approach to the DUT (Fig. 11). This is similar to the Variable Depth Bragg Peak (VDBP) method [19], in which an ion beam's energy can be degraded to determine the location of a sensitive volume - or in the case of 3D-ICs, the *most* sensitive volume(s) - of any given device. In contrast, depth profiling maps the change in output to the change in particle range. In fact, the two are not mutually exclusive: they can be combined to produce a mapping of where the sensitive volumes are located within the device and how their radiation responses may affect the device's aggregate response.
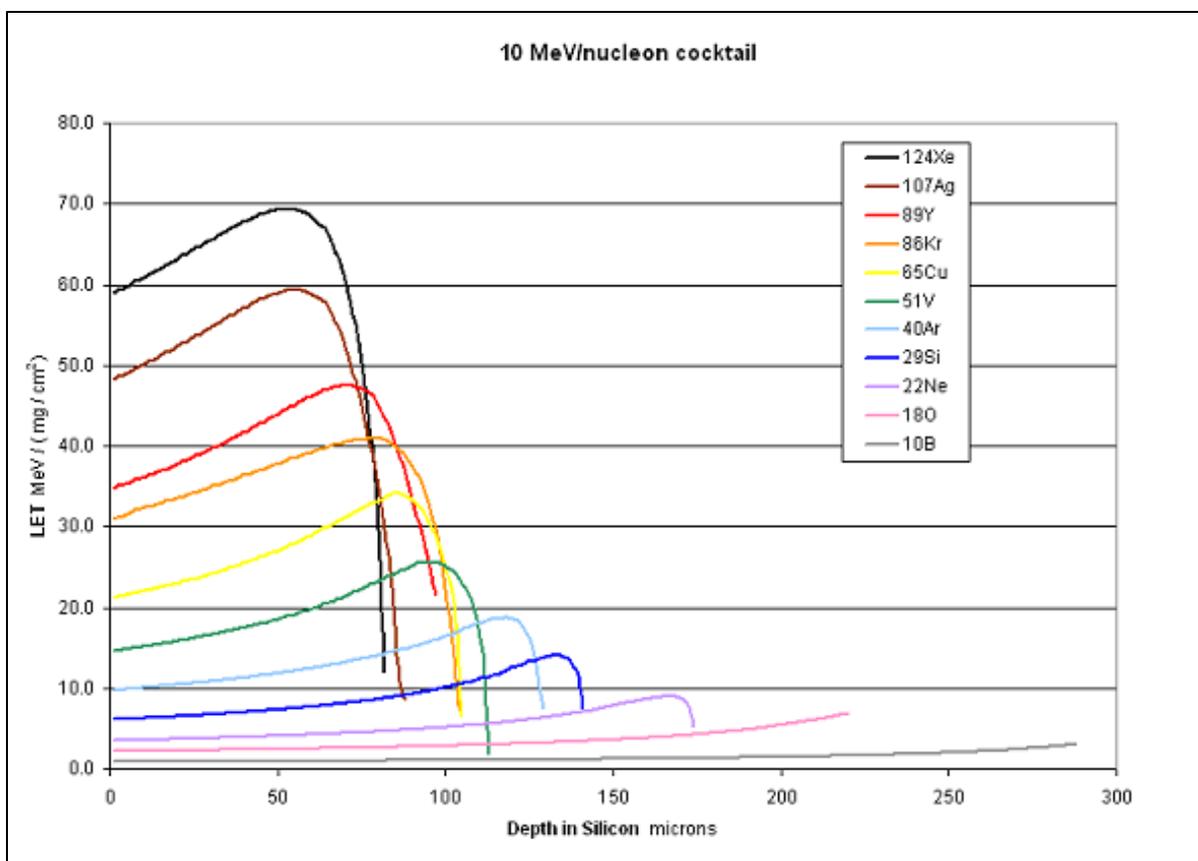
Fig. 10. LET vs Penetration Depth in Silicon for various ion species. The Bragg peak of a given energy and species of ion allows specific ion cocktails to be chosen to meet depth needs [14].
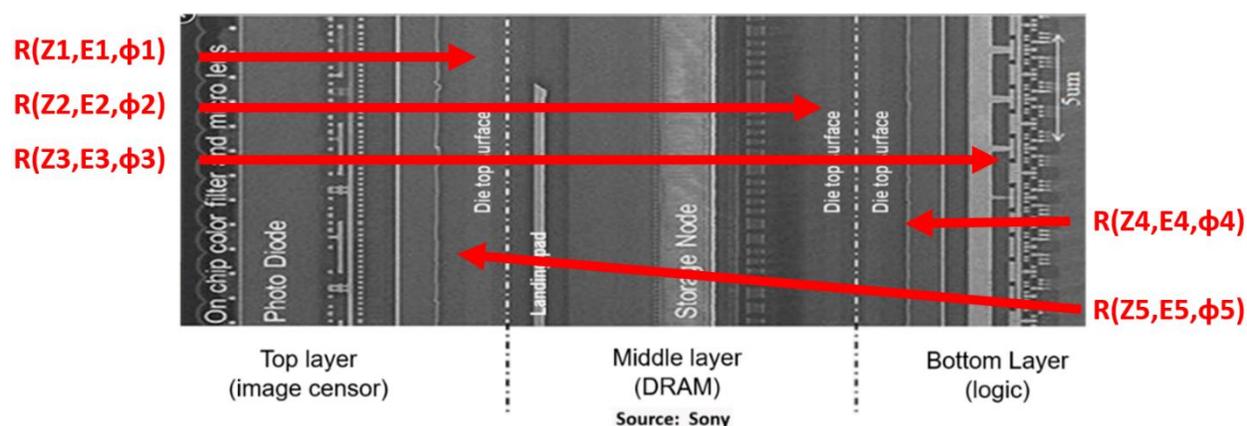
Fig. 11. Examples of different beam orientations and energies that can be used during depth profiling.

To perform depth profiling, it is useful to know something about the materials and geometries in order to narrow that parametric space for ions and energies; the DUT can be a black box in terms of its functionality. Structural information can be obtained by a variety of methods, such as scanning electron microscope (SEM) images or by utilizing the VDBP method. With this information, radiation sources can be tailored for desired tests. X-rays, for example, will penetrate an entire (unshielded) device, providing insight into how the radiation response of all layers simultaneously being struck might affect operation. With ions, energy and species can be chosen to penetrate only a subset of the device before stopping, providing insight into how aggregate operation is affected. Fig. 11 illustrates a few example orientations of incoming particles. Experiments could be designed for particles to enter 1 layer, or multiple, from the "front" or "back" of the device. In such a scenario, SEEs will occur before stopping and displacement damage as the particle comes to a stop.

Intensity Thresholding

As stated in the Background section, standard 24-bit true color images store their data in a matrix that has the dimensions of rows R, columns C, and color bands B, in which "intensity" values are stored (Fig. 4). The key input to this technique is the intensity reported by those bands. Intensity is the relative saturation[3] level of the imager pixel's corresponding detectors (typically standard photodiodes [14]). A 24-bit image's bands are 8-bits each, allowing a maximum intensity of 255 per band. Therefore, a detector would output a value near 0 if no light is detected, or 255 if there is sufficient light to fully saturate the detector element.

A probability distribution function (PDF) can then be computed for an image. An image's PDF is equivalent to its histogram normalized by the number of pixels in the image [15]. This work will refer to such a PDF as an image PDF (iPDF). It follows that during irradiation of the imager, any pixel in which a color band has an intensity that is higher than a significant percentage of the image is likely to be caused by an SEE[4]. The iPDF is typically Gaussian normal, whose

---

[3] Not to be confused with saturation in the context of HSV images [29]. HSV is an alternate way of digitally representing images in which H - the hue, S - the saturation, and V - the average value of color information from the detectors are stored. See section II.5

[4] Depending on where the particle deposits energy, SETs and SEUs are both possible. For example, SETs are more likely in the pixel layer, while SEUs are more likely in the DRAM layer.

properties can then be used to set a threshold, T, based on its mean, μ, and standard deviation, σ. A threshold determined in this way allows for quick calculations, but it is not a robust threshold. This is because for sufficiently large imagers, both μ and σ are both very small values. Thus, random noise fluctuations will typically rise above threshold, resulting in false positives. Similarly, setting a threshold too high will guarantee that any super-threshold events are true positives, but a high threshold loses true positives that are below it.

A better way to set a threshold is to instead calculate the per-pixel PDF, where there is not a single threshold value for an image, but there are instead $N = R{\times}C{\times}B$ separate thresholds per image. This is significantly more computationally expensive, but it provides sufficiently robust thresholds that are near dark field measurements such that few true positives are lost. A per-pixel PDF is calculated by tallying all values for pixel (1, 1), then pixel (1, 2), and so on for all (R, C). Thus, per-pixel thresholds may be calculated using $T_{R,C} = \mu_{R,C} + m{\times}\sigma_{R,C}$, where m is an integer.

A mock data set is shown in Fig. 12 consisting of 9 mock radiation-induced events across 10 frames for a $10{\times}10$ pixel array. This example illustrates the difference between an iPDF and a per-pixel PDF. Frames 2-10 would all have the same iPDF and therefore the same threshold. However, the PDFs for each pixel may be dramatically different. Fig. 13 shows the distributions for the top left pixel (location 0, 0) and the middle pixel (location 5, 5). As implied by these figures, an adequate threshold for one pixel may result in the loss of events in another pixel.
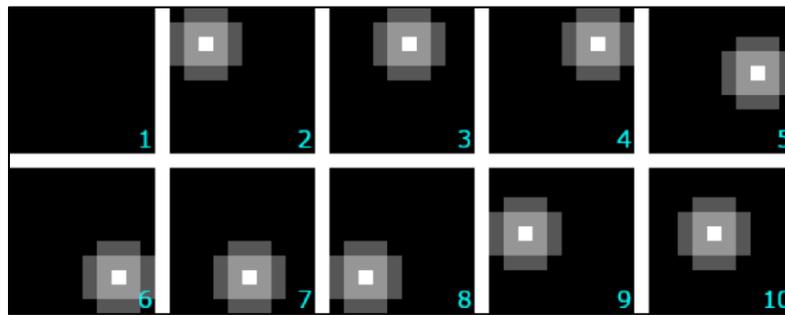
Fig. 12. Ten generated example frames. For simplicity, they are grayscale[5], where red=green=blue.

---

[5] Standard grayscale conversions are calculated based on the "luminance" of a color and assign different weights for red, green, and blue before ultimately setting them all equal. [23]
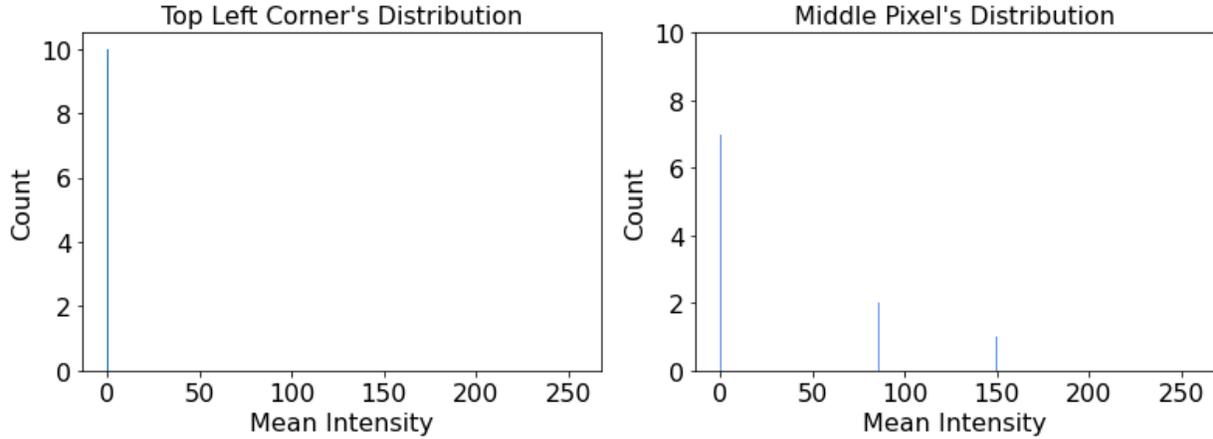
Fig. 13. Per-pixel distributions for the top left (left) and the middle pixel (right).

Splitting of Images

If a pixel has an intensity value that is greater than its respective threshold, it can be assumed to be the result of an ionizing particle event. If a group of adjacent pixels are also above their thresholds, their excitation can be assumed to be the result of the same ionizing particle event[6]. With these assumptions, a splitting algorithm was developed to systematically separate the image (or a subset of the image) into smaller slices containing potential events.

The splitting algorithm, at a very high level, must first determine which pixels exceed their threshold, T, be it a threshold recalculated for each frame or on a per-pixel basis. It must then generate a logical array that maps those pixels' locations; i.e., store a 1 where the pixels are above T, and a 0 where they are not. Both the row sum and columns sum of the logical array must be computed, noting that row sums are counterintuitively stored as a column vector, Y, while column sums are stored as a row vector, X. Vectors X and Y are then scanned, looking for "gaps". Because X and Y are calculated from the logical array, the "gaps" consist of 0's. Each gap's location (in both X and Y) is recorded, and the current frame is split along the location of the gaps. Each split is then searched to find events, and splits with events are saved as new files for subsequent processing.

---

[6] Assuming that the flux of incoming particles is sufficiently low.

A diagram summarizing the above steps is provided below in Fig. 14. In this example, the logical array is inverted, i.e., black squares represent an excited pixel that meets the event criteria. The left diagram illustrates computing of the row and column sums, finding that there are two "gaps" in the *X* vector. The middle diagram indicates where the splitting lines will be drawn. The right diagram illustrates the new saved splits. Note that there are multiple events per split in this example. This is because the splitting algorithm (correctly) found no gaps in the *Y* vector for the corresponding location in the *X* vector. This is not a significant issue if the particle flux is held low enough during data capture, as it is unlikely that two particles would strike close enough in both space and time to cause this example scenario. The data presented in this work were captured with a sufficiently low flux that this did not happen, though it is important to note that it is possible.
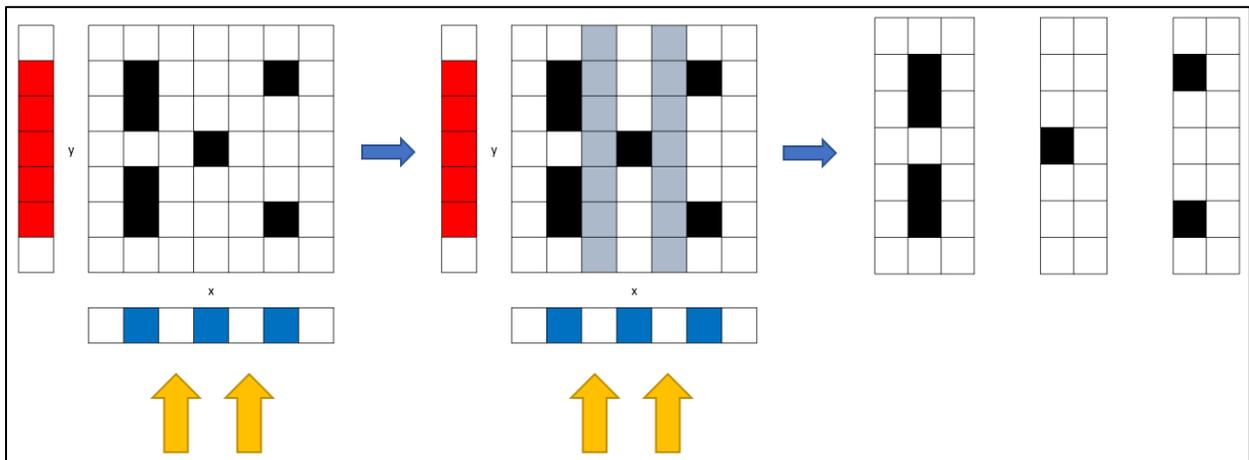
Fig. 14. A diagram detailing operation of the Splitting algorithm.

An example input to the splitting algorithm is provided below in Fig. 15. In this randomly generated example, the events are no longer grayscale, but are colored. Such vibrantly colored data are unlikely to occur during testing, and this example is solely for demonstration purposes. However, it should be noted non-grayscale data, if present, will be properly processed by the splitting algorithm. This example input is a 10×10 pixel array image in which colored pixels exceed their thresholds, while the gray and black pixels do not.
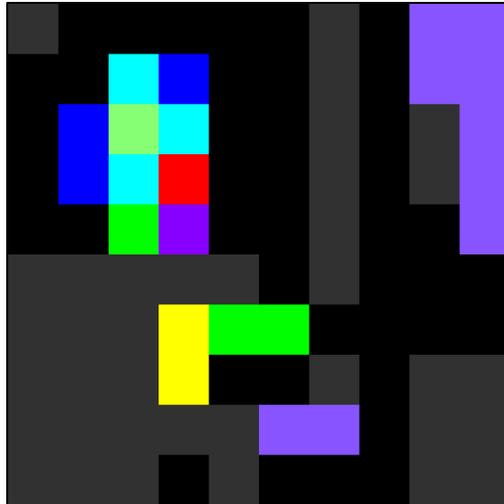
Fig. 15. An example input to the splitting algorithm.

Fig. 16 shows the output of the splitting algorithm for the above example input. As previously stated, the gray pixels do not meet their respective thresholds. Therefore, the gray pixels do not meet the criteria for an event and thus are not considered a part of any events; they are an example representative of dark current. Note that, like the explanatory diagram found in Fig. 14, the middle output contains multiple events in one subframe split. This is due to that event appearing continuous in both the *X* and *Y* vectors.
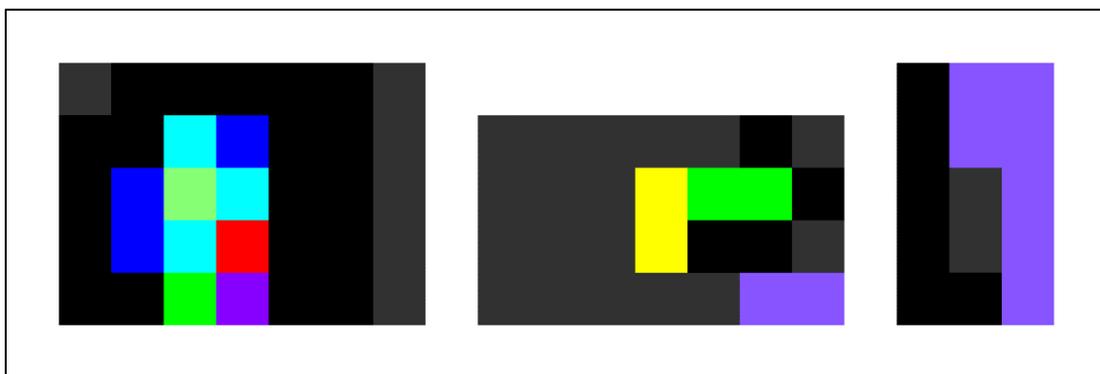
Fig. 16. The output of the splitting algorithm for Fig. 15's example data.

Classification of Events

For a spot of lit pixels to be classified as an event, it must once again meet the intensity threshold. However, there are a different set of criteria that are important to record: "size" and "shape". We define "shape" to be the number of consecutive rows and consecutive columns for which pixels are above their respective thresholds, while "size" is analogous to the number of lit pixels within the event's shape. It is important to record both size and shape, as an event may manifest as different numbers of lit pixels within an areal span. Experimental data can be probed to determine if there are different trends between events that occur in different layers of the DUT.

At a high level, the classifier algorithm must first load in the subframes output from the splitting algorithm that contain events. In these subframes, locations of pixels that exceed the threshold must be determined and a logical array mapping those locations must be obtained. Like the splitting algorithm, row and column sums of the logical array must be computed. Recall that row sums are stored as a column vector, Y, while column sums are stored as a row vector, X. This time however, the number of consecutive non-zero values are recorded, as opposed to the gaps. The number of consecutive non-zero values in X is the size of the event in X and similarly for Y. Then the sum of the entire logical array is computed. Because a logical array consists of only 0s and 1s, this sum will be equal to the number of super-threshold pixels within the event.

Because the classification algorithm uses a similar calculation method to the splitting algorithm, it is also sensitive to multiple events per split. Just as before, with sufficiently low particle flux, it is highly unlikely for such events to occur within a single row and single column of one another. Additionally, similarities between the algorithms act as a pseudo-recursion of the splitting algorithm, effectively providing a level of redundancy by performing a 2nd "splitting". The result is a reduction in multiple threshold-meeting pixel groupings incorrectly appearing to originate from 1 particle event.

The data flow for the example presented in Fig. 15 is now the following: data in Fig. 15 -> threshold calculation -> splitting algorithm -> data in Fig. 16 -> classification algorithm. An example of the output of this flow for the data in Fig. 15 is provided in Fig. 17. Note that once again the gray pixels in Fig. 16 do not meet their respective thresholds and therefore do not contribute to the size nor shape of the events. Note that it is possible for different events to have

the same areal span but different numbers of lit pixels, as illustrated by Fig. 17. By inspection of Fig. 16, it is easy to see that the left and middle splits do indeed both span a 4x3 shape.
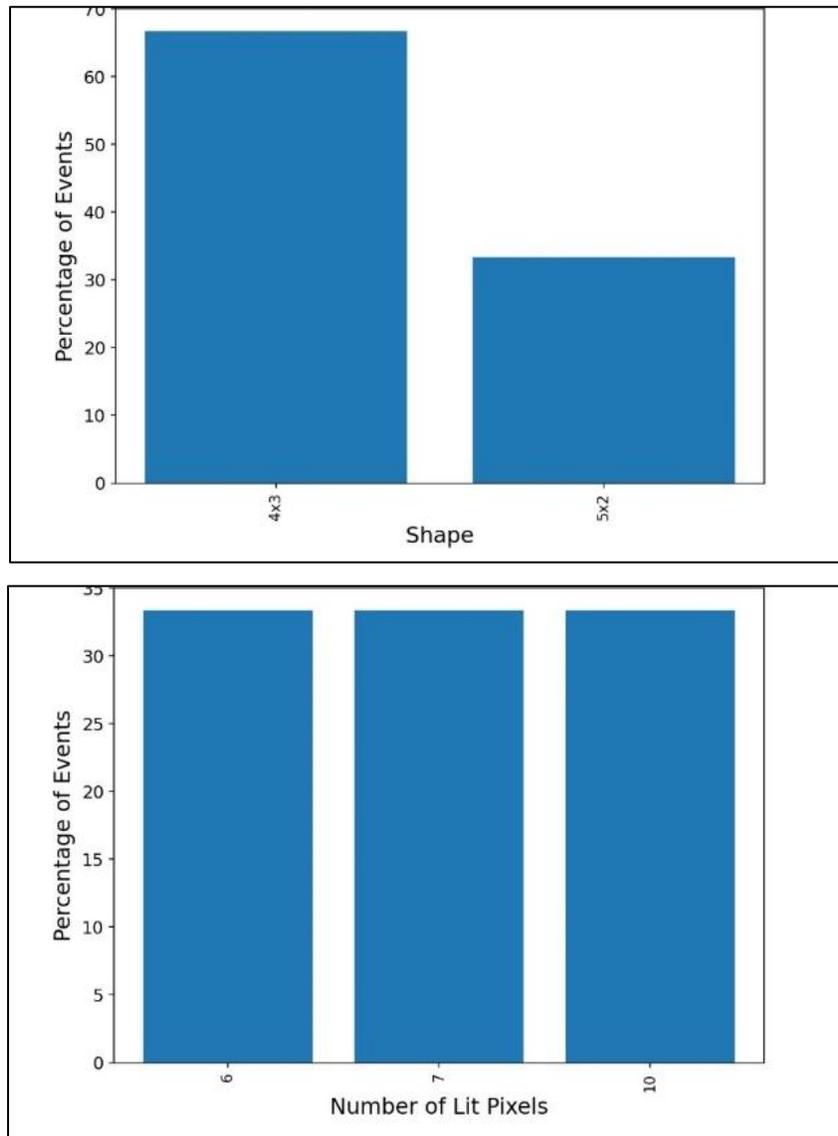
Fig. 17. An example visualization of Fig. 15's statistics, with shape on top and size on bottom.

CHAPTER V

CASE STUDY: ALPHA PARTICLE EXPERIMENTAL SETUP AND RESULTS

Experimental Setup and Test Procedure

Few studies, if any, have reported on radiation effects in this type of 3D-integrated device. However, radiation effects in traditional planar imagers have been reported extensively, such as in papers [20] – [25]. In [20] and [21], heavy ions are found to cause transient groupings of pixels exceeding their dark field measurements while devices remained fully functional during exposure. This section will present the results of an alpha particle test consistent with [20] and [21].

For all experiments presented in this work, commercial off-the-shelf (COTS) versions[7] of the DUT were used. In the COTS packaging, a metal housing with a mechanical lens was attached, as well as an integrated ribbon cable. The ribbon cable provides output to a specific connector used in Sony Xperia smartphones, and such a smartphone was used as the test vehicle for this imager. All imagers used in this work had their mechanical lenses removed (see Fig. 9 in DUT section). All experiments were performed in dark environments where no light could reach the imager.

The imager was mounted under a sealed thin-film Am-241 button source, distanced such that the energy of the particles would be attenuated to different energies (Table 1). Two different imagers from the same lot were used: Imager 01 at 0 mm, and Imager 02 at 18.5 mm. The alphas were calculated to attenuate to ~4.0 MeV and ~1.7 MeV at the surface of the imager from 0 mm and 18.5 mm, respectively. The experiment is set up such that the 1.7 MeV particles deposit energy and come to a stop within only the topmost pixel layer, while the 4.0 MeV particles deposit energy in both the pixel and DRAM layers before stopping in the DRAM layer.

---

[7] Prior to product release, Sony produced evaluation models of this imager. They were no longer available by the time of this work.

Table 1. Distances and Corresponding Energies

Imager used	Air spacing used	Corresponding energy
01	18.53 mm	1.70 MeV
02	0 mm	4.05 MeV

In this experiment, video recordings were captured at both listed energies with four combinations of ISO and shutter speed[8] (SS), analogous to gain[9] and integration time, respectively. ISO was available from 50 to 3200, while SS was available from 0.1 ms to 100 ms. These limits and the combinations used are provided below in Table 2. These conditions were captured to better understand the impact of ISO and SS on the imager's response. The "shape", or the areal span, of the imager's response to a particle was the primary metric of this study. As noted in the previous chapter, shape provides a general idea of the DUT's response but not the entire story.

Table 2. ISO/Shutter Speed Combinations

Available Settings		Combinations used	
ISO	Shutter Speed	50, 0.1ms	50, 100ms
50	0.1ms		
3200	100ms	3200, 0.1ms	3200, 100ms

The DUT was attached to a smartphone via ribbon cable, which was in turn attached to a laptop using a USB cable. An explanatory diagram is included below in Fig. 18. The DUT, in blue, was in contact with aluminum sheet ~0.06 mm thick with a pinhole ~1.0 mm in diameter. The alpha particle button source, in orange, was supported at two distances (Table 1) above the DUT with spacers. A rigid piece of plastic, shown in green, was cut to provide structural support at the

---

[8] The convention for "shutter speed" is misleading, as it refers to the duration for which the sensors capture charge.
[9] ISO is analogous to amplifier gain when shutter speed and aperture size are independently configurable. ISO values convey information about the relative gain [11]. (e.g., ISO-3200 has 64x the gain of ISO-50.)

same height as the smartphone to support the spacers. The smartphone was controlled remotely using a USB cable, and video in the form of 4K UHD mp4 using the h.264 standard video codec was captured using the phone's built-in Sony camera app. The videos were copied and converted into PNG format images, which were analyzed using the technique described in Chapter IV. A labeled photo of the setup is shown in Fig. 19.
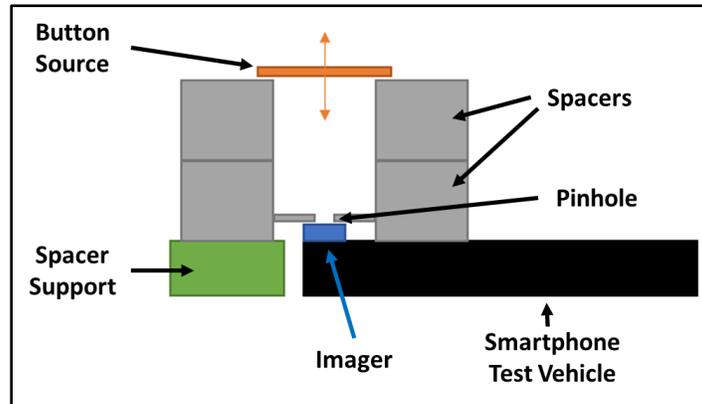
Fig. 18. Experimental test setup explanatory diagram.
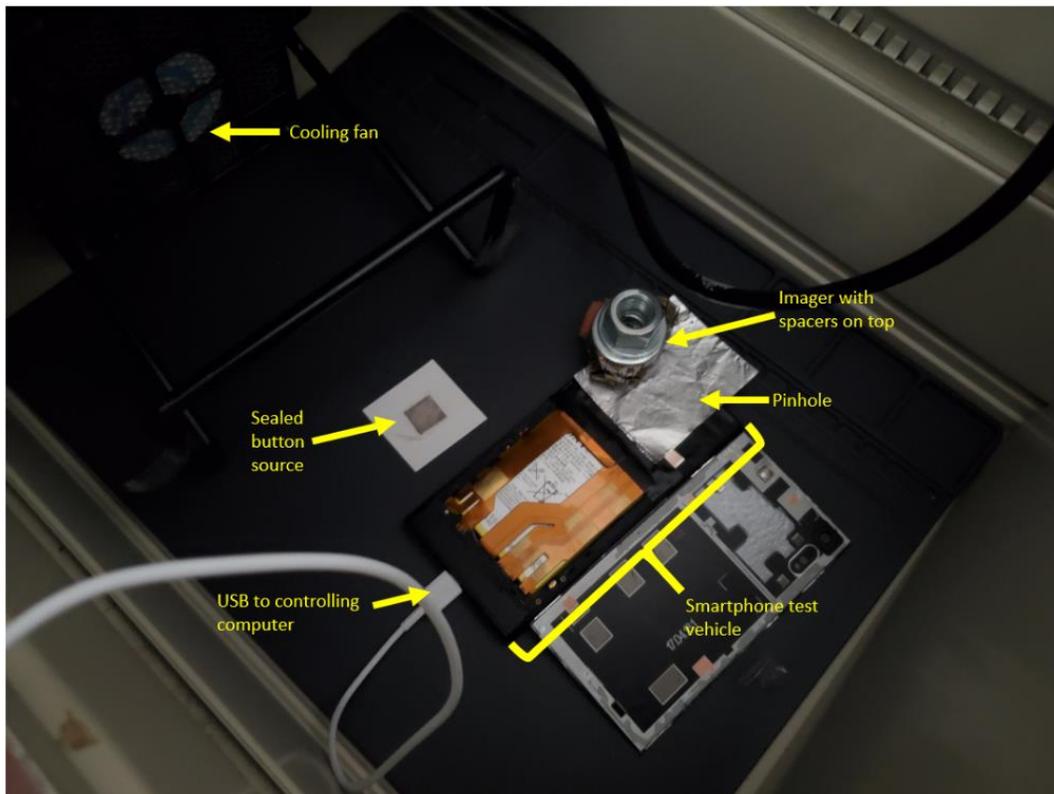
Fig. 19. The experimental test setup.

The data sets were captured such that the fluence of alpha particles through the pinhole was consistent at both energies. However, the "apparent" fluences of results in this section are *not* comparable. This is because of the geometry problem arising from using a button source: the particles are not collimated and do not originate from the same point. Rather, the particles can be thought to originate from a randomly located point source at random angles of incidence. Under these conditions, the particles that can make it through the pinhole have different maximum incidence angles. The region of influence for the button source then changes geometry in relation to its distance from the pinhole. At longer distances, the region of the imager directly underneath the pinhole will occupy a larger percentage of the button source's region of influence, while the opposite is true for close distances. This results in a wider hot spot of alpha particle strikes in the 4.0 MeV data, while having a significantly tighter hot spot for the 1.7 MeV data. An explanatory diagram is provided in Fig. 20 and Fig. 21. In Fig. 20, the orange button source emits particles at random incidences from random locations. Some of these particles make it through the gray pinhole to strike the black DUT. In Fig. 21, the shaded region represents the region in which the particles of most normal incidence strike the DUT, a region whose cross-sectional area decreases with increased distance between button source and pinhole.

A region of 780×780 square pixels, or roughly 1 mm^2, was chosen for the data analysis. This region was selected in the most densely populated center of all events in order to analyze data directly underneath the pinhole. This selected area occupied a large portion of all thru-pinhole alphas in the 1.7 MeV data. Conversely, this selected area occupies a small portion of all thru-pinhole alphas in the 4.0 MeV data.
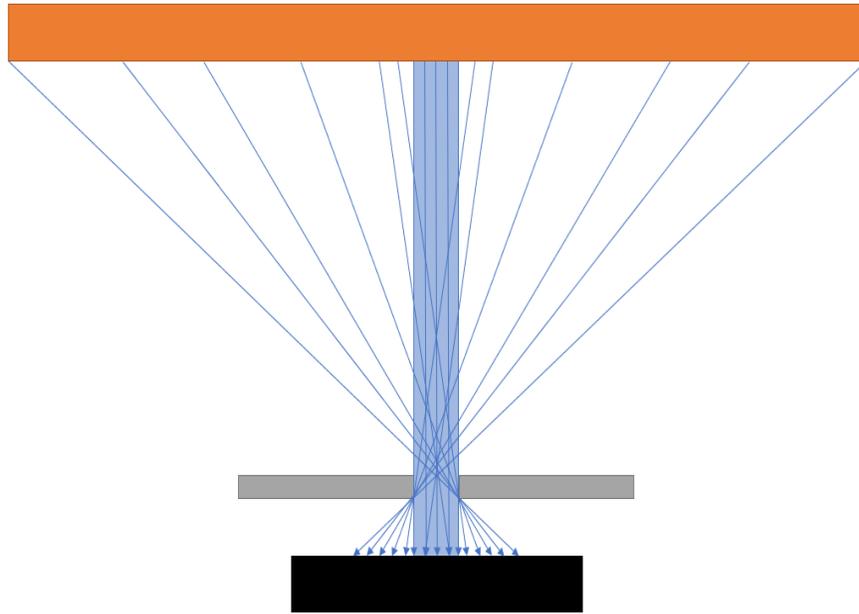
Fig. 20. This diagram illustrates the thru-pinhole fluence. The button source is represented in orange, the pinhole in gray, and the DUT in black. The region shaded in blue is the cross-section of the pinhole.
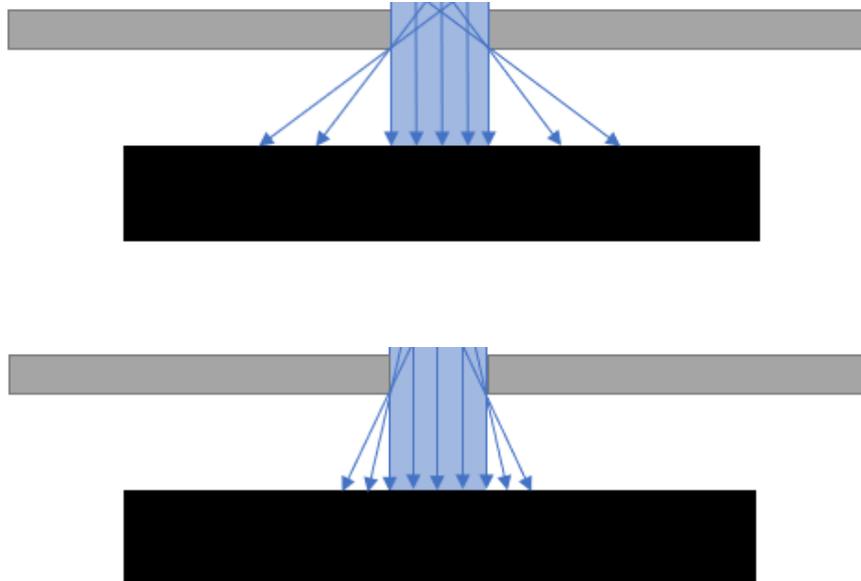
Fig. 21. The effect of source distance on fluence. While the same number of particles passes through the pinhole to strike the imager, the area directly underneath the pinhole occupies a smaller percentage of the source's apparent region of influence. The top diagram above illustrates a source close to the imager, while the bottom diagram illustrates a source far from the imager.

Experimental Results

  Results were determined by first converting the frames in the output videos to PNG files. As PNG files are lossless [22], they retain the same information as each frame in the video[10]. Data at each energy were recorded for different total exposure times that resulted in comparable thru-pinhole fluences[11]. Dark field values[12] of each pixel were analyzed and found to follow unique Gaussian normal distributions[13]. From these dark field measurements, a per-pixel mean, μ, and per-pixel standard deviation, σ, were computed. A corresponding threshold of

$$T = \mu + 8\sigma$$

was calculated for each pixel and stored into a threshold matrix. These thresholds were found to yield event counts comparable to expected apparent fluences. The threshold matrices were used to generate superimposed images of all data sets, allowing verification that the pinhole aperture was properly aligned. The data sets were analyzed for radiation-induced transient events. For this experiment, a transient event is defined as a group of adjacent pixels that have any color band value greater than their respective thresholds. That is, if a pixel has a red band value larger than its red band threshold, the entire pixel was considered super-threshold regardless of green or blue band information. Note that because of the DUT's properties, SEUs and SETs both manifest as single-frame transient events.

  The figures in this section present statistics of the "shape" of events. The shape of an event refers to its areal span, presented as *Row* x *Column*. All bar graphs are normalized to arbitrary cross-sections, and their distributions can be thought of as 3-dimensional PDFs.

---

[10] Data from the imager was compressed as it was saved as a video. The PNG conversions did not add any further data loss (from file compression).

[11] "Apparent fluences", however, were different. See explanation on previous page.

[12] Data was recorded in a dark chamber with the "beam turned off" (no incoming particles). The resulting intensity values are the "dark field values".

[13] I.e., pixel (1,1) has a different μ, σ than pixel (1,2), etc.

1.7 MeV Experimental Results

Cross sections (in arbitrary units) of the 1.7 MeV data set are plotted against row and column span (Fig. 22). A common feature of the 1.7 MeV data is that the distributions are somewhat bimodal, with a common peak centered near 2x2 and another near 6x6. The same button source and imager were used under the same conditions for all 1.7 MeV data, therefore, it is reasonable to assume that each of these plotted distributions should appear similar. In fact, it is not unreasonable to assume they are the same distribution with appearances skewed by device operation, as threshold is influenced by operation settings (impact to be explored in Discussion section). While all data sets captured are exposed to the same fluence of particles, the 100 ms shutter speed setting reports more events because the effective duty cycle is ~3 orders of magnitude longer.
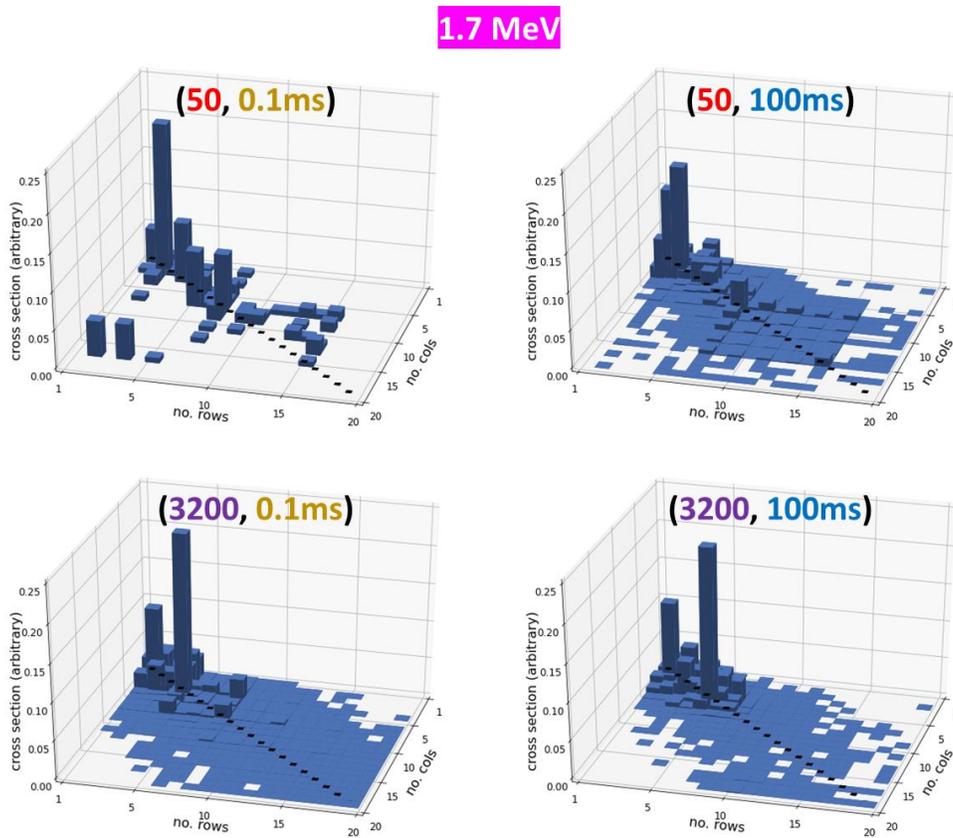
Fig. 22. ISO/SS effects on 1.7 MeV event detection. Threshold $= \mu + 8\sigma$.

4.0 MeV Experimental Results

The cross sections of the 4.0 MeV data are reproduced below in Fig. 23. Like the 1.7 MeV data, most of the distributions appear somewhat bimodal, with a peak appearing centered around 2x2 and another within the 5x5 to 8x8 range. Again, it is believed that these are the same distribution of particles, but their apparent distribution is influenced by device operation conditions. The obvious outlier to this trend is the (3200, 0.1 ms) operating condition, which has a dramatically different distribution. This operating condition (and particle energy) uniquely had repeated fatal errors during data capture, suggesting the occurrence of single-event latch-ups (SELs). This is the most readily apparent result from the using the depth profiling technique and will be further explored in the Discussion Section.
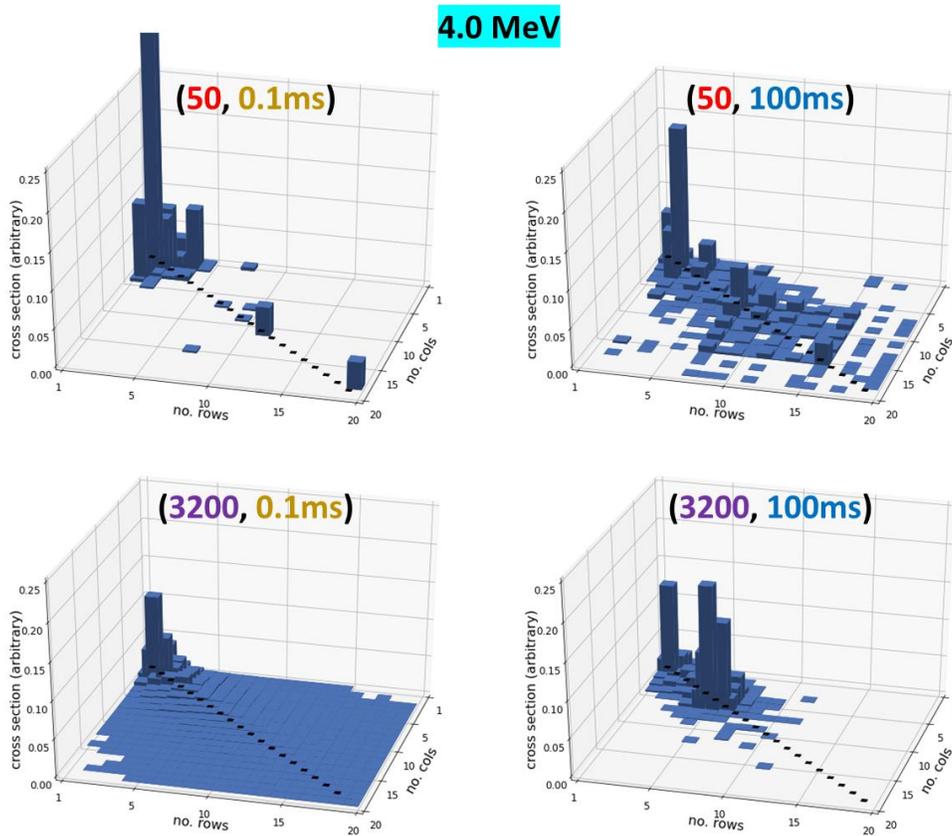
Fig. 23. ISO/SS effects on 4.0 MeV event detection. Threshold = $\mu+8\sigma$.

Comparison of Cross Sections Over Particle Energy

Side-by-side comparisons of the results for 1.7 MeV and 4.0 MeV for each operating condition are reproduced below in Fig. 24 - Fig. 27. Three of the four operating conditions yield similar apparent distributions for both energies; the (3200, 0.1 ms) data in Fig. 26 will be expanded upon in the Discussion Section. At ISO 50, both shutter speed settings (Fig. 24, Fig. 25) show an apparent skew toward the smaller peak at around 2x2 for the 4.0 MeV particles. This result is as expected: the 4.0 MeV particles have a lower LET while in the pixel layer and are thus less likely to share charge between pixels. The lower LET similarly causes the 4.0 MeV alphas to favor somewhat smaller groupings per event at (3200, 100 ms) (Fig. 27).

Somewhat unique to the (50, 100 ms) operating condition (Fig. 25) was the appearance of unusually large events, spanning extremely areas up to 780x4 (not shown in these plots). These disproportionately large events appear at both energies and are seemingly caused by transients in row/column readout circuitry in the pixel layer. These will be explored more closely in the "Discussion" section.
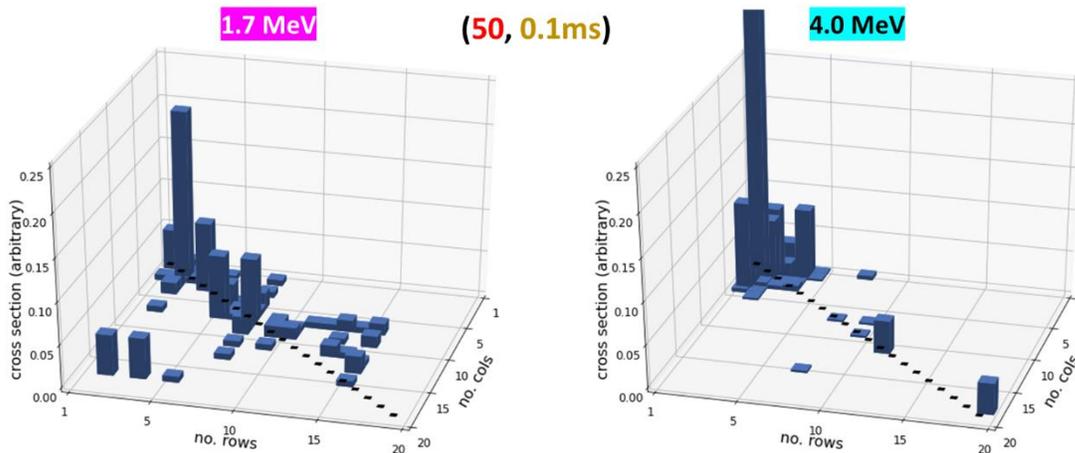
Fig. 24. Cross Section vs Row/Col Dimensions comparison for (50, 0.1 ms).
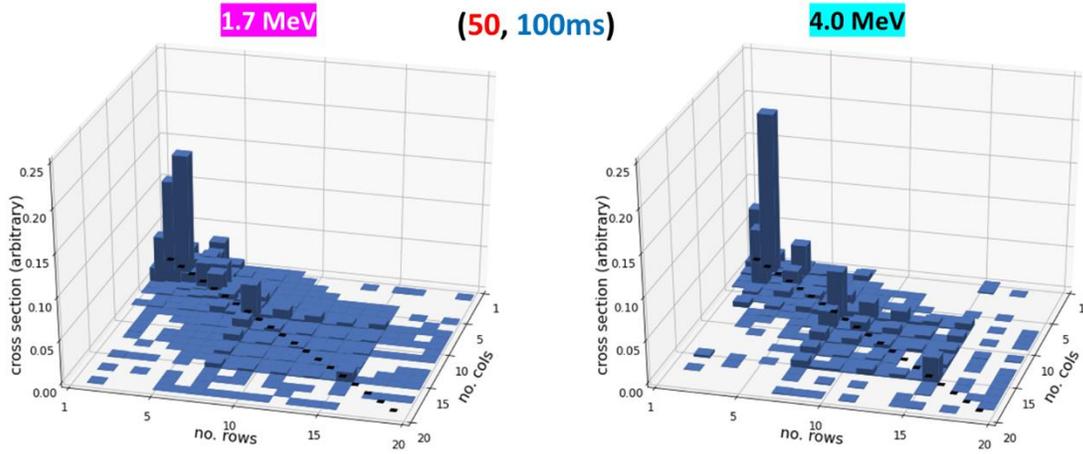
Fig. 25. Cross Section vs Row/Col Dimensions comparison for (50, 100 ms).
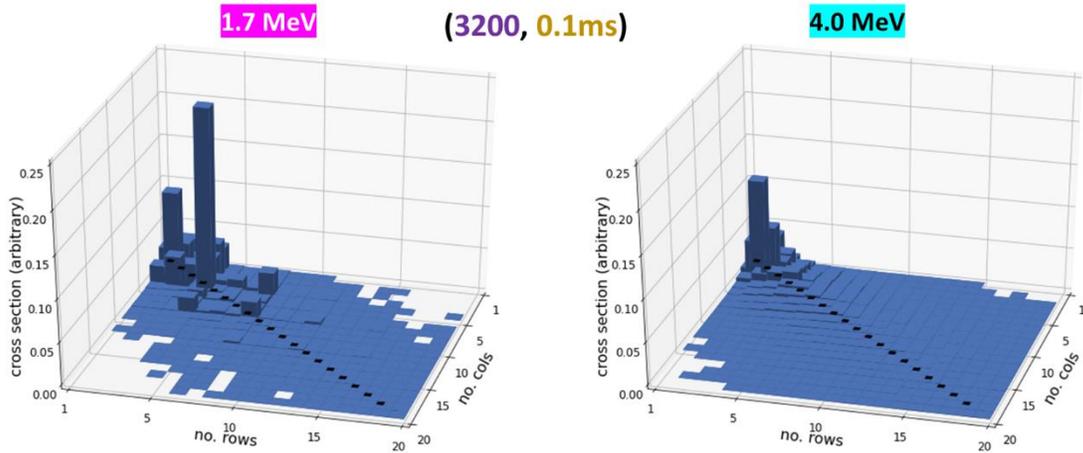
Fig. 26. Cross Section vs Row/Col Dimensions comparison for (3200, 0.1 ms).
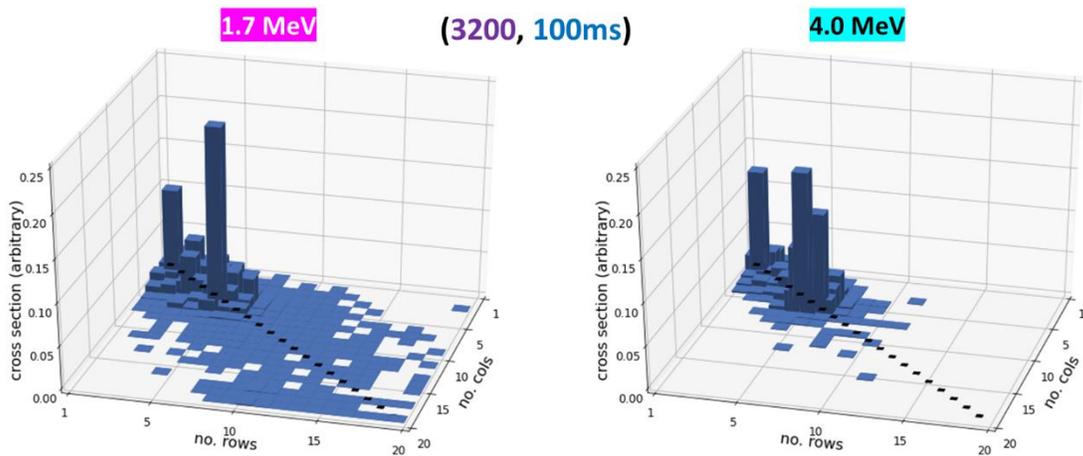
Fig. 27. Cross Section vs Row/Col Dimensions comparison for (3200, 100 ms).

CHAPTER VI

DISCUSSION

VI.1   The Pixel Layer's Apparent Dominance

As noted in the previous section, event cross-section distributions appeared somewhat independent of alpha energy for all settings (as seen in Fig. 24, Fig. 25, and Fig. 27), except the (3200, 0.1 ms) setting (explored in the next sub-section). This suggests a common source of apparent event shape. The 1.7 MeV data set was captured in such a way that the alphas could not reach the DRAM layer. However, the aggregate device's response was very similar across energies. Therefore, it is reasonable to assume that the pixel layer's response dominates the device's response for both alpha energies. This result was expected; at a high level, the layer designed to collect charge from external stimuli was found to be more sensitive to external stimuli than a layer that was not designed for this purpose. Although each image is stored in the DRAM at least once, the DRAM operates at 200 MHz, a speed significantly faster than capture rate of the pixel layer. It is possible that the DRAM does in fact produce a similar number of errors, but because of the nature of DRAM, it is possible that transient effects are cleared out before they can impact the imager's final video output.

Monte Carlo Radiative Energy Deposition (MRED) is a Monte Carlo-based application that simulates radiation transport [23] [24]. A MRED simulation was performed to verify the energies and fluences of the incoming alpha particles at the experimental distances. The packaged thin-film Am-241 button source was modeled as a planar region that fired alpha particles at random angles, where alpha energy was calibrated to match previous dosimetry data. The imager was modeled as a uniform 30-µm thick block of Silicon with a cross section of 1x1 µm^2, representative of the area directly beneath the pinhole. These simulations revealed that the use of this button source results in incoming alpha particles that are non-monoenergetic at the surface of the DUT, with the most common energy alphas attenuated to ~1.7 MeV and ~4.0 MeV. As a result, the 4.0 MeV spectrum also contains alpha particles at 1.7 MeV, albeit at significantly fewer number. The results from the MRED simulations are plotted in Fig. 28 as counts versus energy bin. For reference, 2 million alpha particles were fired in each simulation.
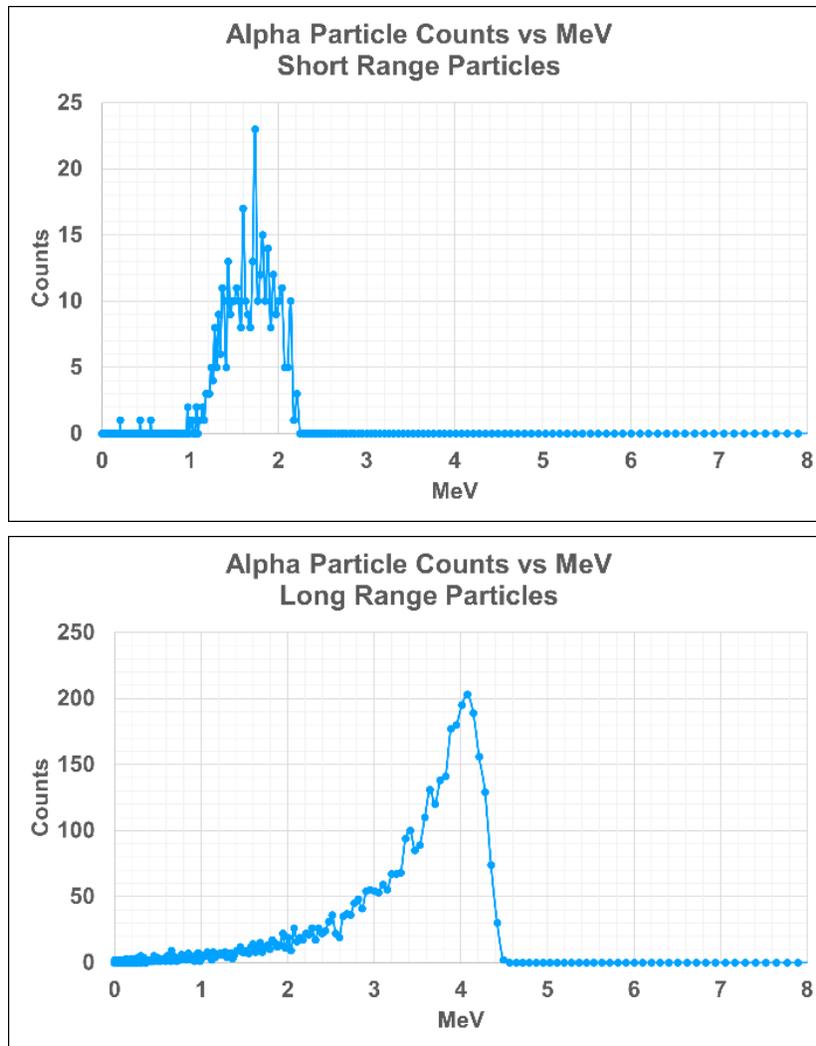
Fig. 28. MRED simulations of count vs MeV, with ~18.5 mm spacing on top and ~0 mm spacing on bottom. Note the Y scales are different.

LET curves (normalized to Silicon) corresponding to 4.0 MeV and 1.7 MeV alpha particles are reproduced in Fig. 29. The 1.7 MeV particles have relatively high LETs within the pixel layer, while the 4.0 MeV particles have relatively low LETs within the pixel layer. The 4.0 MeV particles originated closer to the DUT and have a larger range of incident angles (and therefore a more variance in particle range) than the 1.7 MeV particles. Despite these similarities and differences, particles at both energies largely caused similar responses from the DUT. This further suggests that the pixel layer dominates the DUT's response; despite seeing different LETs and different effective ranges, the DUT's response is largely similar, as seen in Fig. 24, Fig. 25, and Fig. 27.
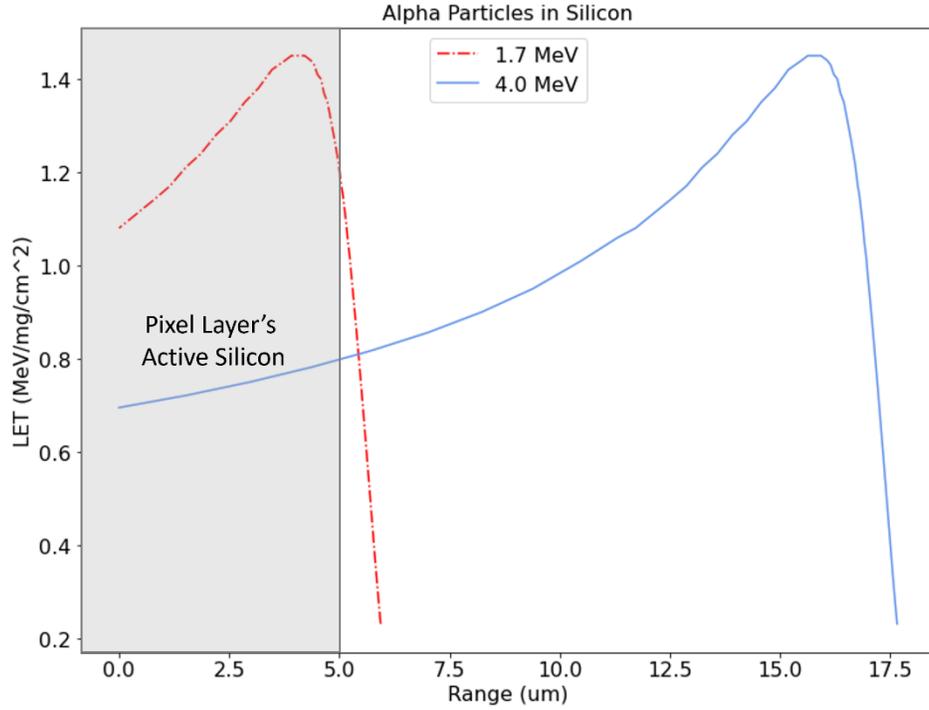
Fig. 29. LETs of Alpha particles in silicon. The gray overlay indicates the active silicon in the pixel layer.

SRIM [25] simulations were performed to verify stoppage of the alpha particles into the DUT. The SRIM simulations revealed that the alpha particles had acceptable straggle in the relatively thin layers of the DUT (~10 µm per layer including metals). The SRIM simulations are reproduced in Fig. 30 and are plotted such that the alphas enter the volume from the left side of the figure.
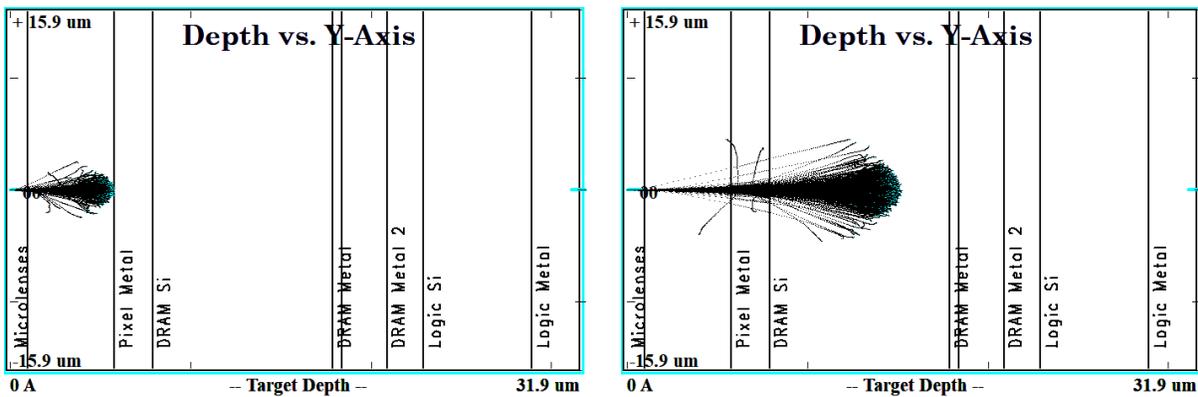
Fig. 30. SRIM simulations for alphas of 1.7 MeV (left) and 4.0 MeV (right) into the DUT.

VI.2   Observation of Single-Event Latch-Up with 4.0 MeV particles at (3200, 0.1 ms)

During experiments, only one combination of ISO, SS, and particle energy resulted in repeated fatal errors in the DUT: 4.0 MeV particles at the (3200, 0.1 ms) capture condition. This unique combination showcases the most prominent example of depth profiling in this pixel-dominated device. At 4.0 MeV, the alphas possess sufficient energy to pass fully through the pixel layer and continue into the DRAM layer. A micrograph of the DRAM layer is reproduced in Fig. 31, where control circuitry is visible in the center of the DRAM layer.

All data captured at 4.0 MeV was captured at the location circled in black in Fig. 31. However, only the operation settings of (3200, 0.1 ms), the highest available gain and the shortest shutter speed, resulted in fatal errors that shut down device operation. During these errors, the imager was noticeably warmer to the touch and was able to continue operation after a power cycle. This behavior is consistent with single event latch-up: increased current draw increases power consumption and temperature while returning to normal operation upon power cycle. This suggests that the device is more susceptible to latch-ups in this peripheral/control region when operating at the highest gain and shortest shutter speed. It should be noted the suspected latch-ups may occur in all data sets, however, they only caused fatal errors in this region under this capture condition.
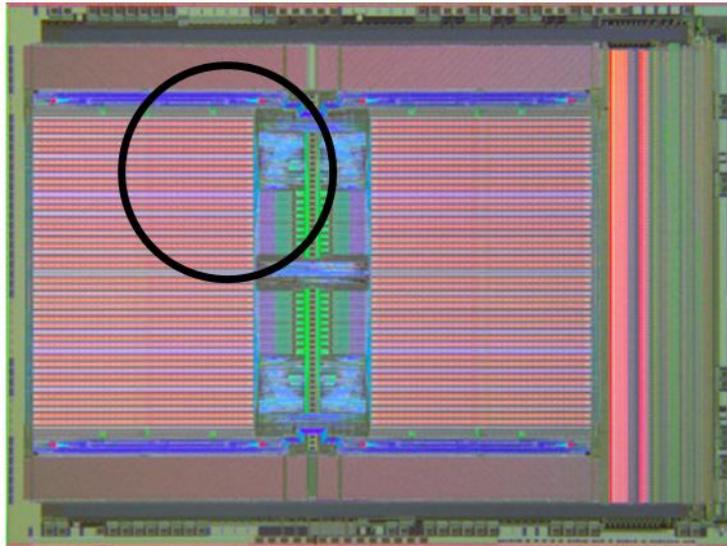
Fig. 31. A micrograph of the DRAM layer, with an overlay of the pinhole. The central region of this layer is peripheral control circuitry.

VI.3   The Impact of ISO, SS on Threshold

The analysis of the data used a per-pixel intensity threshold of $\mu+8\sigma$, where μ and σ were determined from the dark field measurements, to distinguish radiation-induced events. Tables of μ and σ are provided below in Table 3. If a group of adjacent pixels exceeded their respective thresholds for a single frame, the group was considered a transient event. This is typically a sufficient method of defining a threshold for event classification. However, it has limitations when studying this imager's response to radiation because both μ and σ for any given pixel change with the gain and shutter speed used during data capture. In general, an increase in ISO yielded an increase to both μ and σ, with the same trend for shutter speed increases. Moreover, changes in ISO and SS affect the scale of change in μ and σ differently. For example, when going from (50, 0.1 ms) to (3200, 0.1 ms), both μ and σ increase. However, the change in μ tends to be significantly smaller than the change in σ, and thus the threshold becomes σ-dominated.

This relation is important to note because of dose effects. Even in modern CMOS devices, leakage current can significantly increase with total ionizing dose (TID) [26], which may skew the mean μ used in threshold calculation. However, TID to the DUT in these experiments was small.

Table 3. Table of μ and σ.

		50, 0.1ms	50, 100ms	3196, 0.1ms	3196, 100ms
1.7 MeV	Average μ	1.02	1.05	1.03	1.35
	Average σ	0.13	0.21	0.18	0.63
4.0 MeV	Average μ	1.02	0.95	1.04	1.40
	Average σ	0.13	0.22	0.22	0.67

VI.4   The Effect of Threshold on Distribution Characteristics

Because of the changes in both factors of how thresholds are calculated, the distributions themselves of apparent event shapes can be influenced by the ISO/SS settings. Fig. 32 illustrates how threshold may affect the appearance of events. If the event shown on the left is how the event truly manifests, then a difference in threshold changes how it is perceived. With a threshold of 2,

the event appears as a 6x5, but with a threshold of 3, it appears as 2 separate events of 4x4 and 1x1. In the distribution plot, this changes how tall the peaks appear and how sparse the distributions appear. This effect also leads to some of the unintuitively larger events seen in Fig. 22 through Fig. 27. However, much of those large events are valid, as the next section explains.
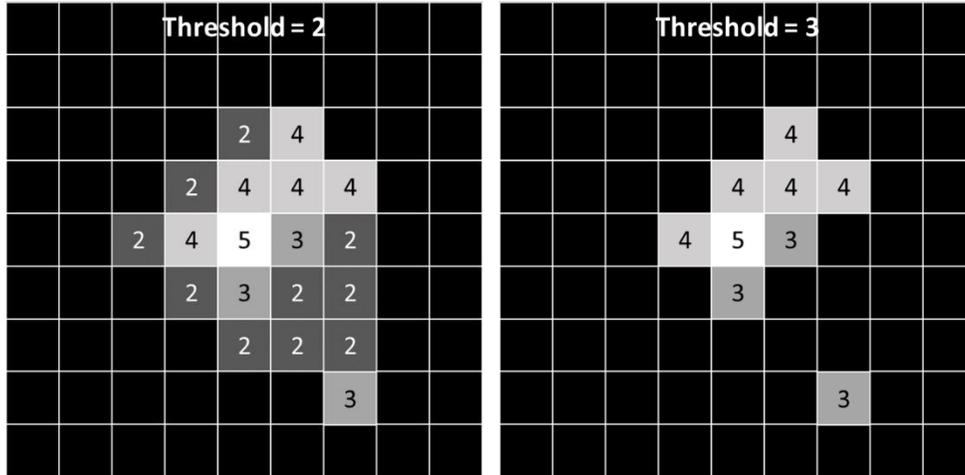
Fig. 32. An example of threshold affects the appearance of events.

The effect of threshold on the distribution's appearance can be used to our advantage. For example, by simply setting a high threshold, we can focus on events where the largest pixel saturation – and therefore the most deposited charge – are recorded. Similarly, by choosing a lower threshold, we can observe how far charge may spread within the device.

VI.5   Mechanisms for Disproportionately Large Events

In all plots within Chapter V's experimental results section, results were presented that suggest a single alpha particle contributed to events that seemed unreasonably large. Intuitively, an event caused by a single alpha particle seems unlikely to manifest as large as 400 affected pixels (shapes of "20x20", in Fig. 22-Fig. 27), especially when each pixel is ~1.5 μm^2. This section will demonstrate that these events are, in fact, real manifestations of alpha-induced transient events and not the result of improper data analysis.

A common interpolation technique in the field of image processing is bilinear interpolation. Bilinear interpolation is simply linear interpolation in 2 dimensions: Z as a function of X, and Z as

a function of Y. Bilinear interpolation is a common method of performing demosaicing, or the assignment of color information to neighboring detectors of a different color.

An example scenario illustrating the demosaicing process is presented below in Fig. 33. In this scenario, an array of physical pixels is freshly zeroed out (Fig. 33.a). An alpha particle strikes a zero-valued blue detector head-on, with no charge shared (Fig. 33.b). The blue pixel is then assigned an intensity of 100/255 (Fig. 33.c). Bilinear interpolation results in the struck pixel's nearest neighbors being assigned the average value of the struck pixel and the struck pixel's 1^{st}-degree separated neighbors. That is, the struck pixel looks at the 2^{nd} nearest neighbor to the left, containing a value of 0, and assigns the average value of 100 and 0 (or 50) to the 1st left neighbor. This stage is illustrated in Fig. 33.d. In Fig. 33.e, the non-blue neighbors of the struck pixel are assigned blue color values in a similar fashion, resulting in the appearance of 13 pixels experiencing higher-than-threshold values spanning an area of 5x5 (Fig. 33.f).
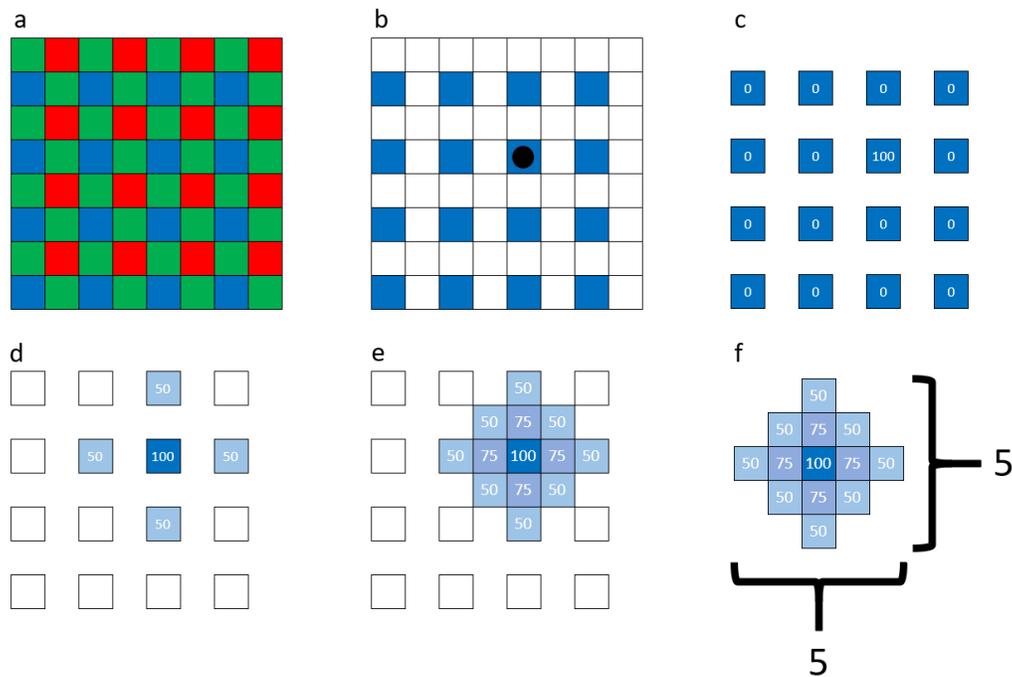
Fig. 33. An example of a direct hit of a blue detector. Subfig. (a) represents a freshly cleared pixel array in which all values are 0. In (b), a blue detector is hit head on, resulting in the assigned intensity level of 100 in subfig. (c). In subfig. (d), the struck pixel's nearest blue neighbors are assigned the average value of the struck pixel and its neighbors. In Subfig. (e), the non-blue neighbors are assigned blue values, again an average of neighbors. Subfig. (f) illustrates that the resulting event appears to have 13 lit pixels within a span of 5x5.

Fig. 34 illustrates an example where the corner between 4 detectors is hit, resulting in the appearance of 24 super-threshold pixels spanning an area of 6x6. Notice that these situations both result in seemingly disproportionately large events without any charge sharing between pixels. Research presented in [20] and [21] find that charge sharing between pixels is indeed possible in a backside-illuminated CMOS imaging sensor with shallow isolation oxides. In this context, the shallow isolation oxides do not span the full thickness of the detector elements. The findings of these works along with the presence of shallow isolation oxides in the SEM images suggest that it is indeed possible for charge sharing to occur between pixels. It is therefore reasonable to assume that charge sharing from any of the struck pixels would result in events appearing larger than these presented examples.
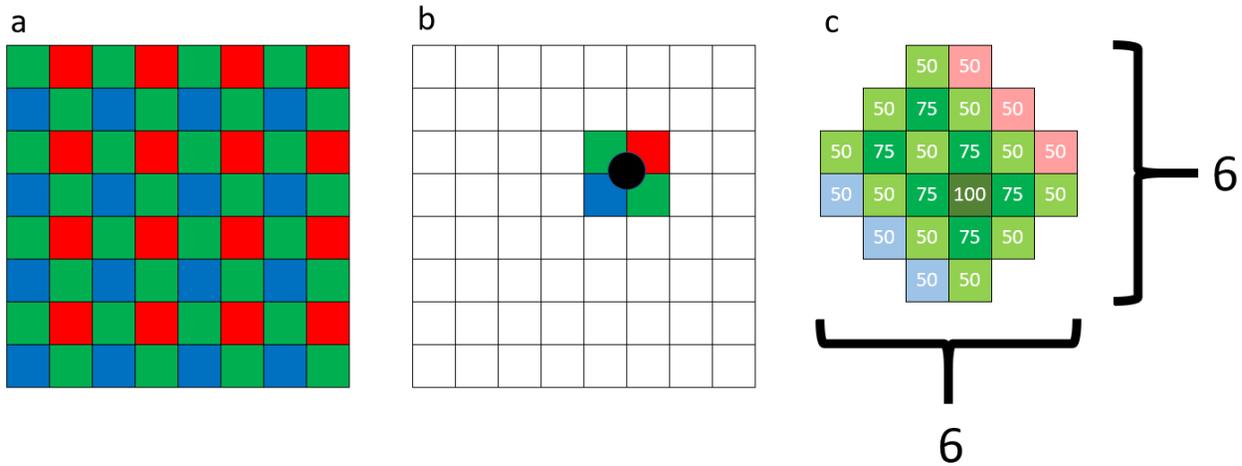
Fig. 34. An example of a corner between detectors being hit. In this scenario, the apparent event has 24 pixels lit within a span of 6x6.

Events that seemed even more unreasonably sized were recorded during (50, 100 ms) operation. At both energies, events were recorded that spanned apparent areas of up to 780x4, or the length of the entire analyzed area. It is reasonable to assume that these are caused by row/column driver upsets in the pixel layer, because they are recorded at both alpha energies and the 1.7 MeV alphas do not have sufficient energy to enter the DRAM layer. Examples of such events are reproduced below in Fig. 35. The selected examples are both row-major (for printing convenience), though column-major events of such scale are also observed.

Although events of such scale were only recorded during (50, 100 ms), it is believed that these row/column upsets are happening in all data sets but are undetectable. The (50, 100 ms) operation setting has the unique characteristics of having the lowest μ and σ while simultaneously having the longest time between pixel resets. This situation provides sufficient time to capture short-lived row/column transients, while the low threshold allows the entirety of the row or column to be considered within 1 event by the analysis technique. These events are masked by the short capture window in 0.1 ms operation. The heightened dark field values in ISO 3200 operation also mask these events, which is why they are not recorded in the (3200, 100 ms) data set.

Fig. 35. Examples of "extremely large events"; the top extends to the left for ~100 pixels, while both extend off to the right for hundreds of pixels. The top example occurred in the 1.7 MeV data set, while the bottom example occurred in the 4.0 MeV data set. Both examples were captured at (50, 100 ms), and the reproduced cutouts are 154×33.

CHAPTER VII

CONCLUSIONS

3D-ICs are becoming increasingly common due to their form factor and decreased power, including in cutting edge imaging devices. This work presents a standalone method of data analysis and the "depth profiling" data acquisition approach. When the two are combined, it is possible to obtain informative insight into the radiation response of 3D-IC devices, even when minimal information is known. Depth profiling is performed by choosing ion species and energies that will cause excitation events in specific layers. Data captured from multiple desired combinations of layers and orientations grant a rich data set for the analysis technique. The analysis technique takes advantage of a simple piece of image processing and storage: the intensity of color bands. From intensity, determining the event size and shape is possible. With this information, it is possible to reasonably determine the relative contribution of layers in the 3D-IC to the total output.

A case study was designed such that alpha particles of predominately 1.7 MeV and 4.0 MeV penetrate different layer combinations of a 3D-IC imager. It was found that despite differences in energy spectrums, the aggregate output of the device was largely the same for both data sets. This suggests that the sensitivity to external stimuli and response of the pixel layer dominate the holistic response of the device. The device was found to exhibit alpha particle induced row/column transients when operated at (ISO, SS) = (50, 100 ms). Furthermore, the device was found to exhibit alpha particle induced SEL when particles of sufficient range can deposit charge in the 2^{nd} layer during operation at (3200, 0.1 ms).

It is apparent from these results that this analysis technique provides a convenient way to analyze results from experiments performed while using the depth profiling approach. Use of this technique coupled with experiments at different energies allows quick determination of trends due to radiation events and their manifestations in the output. Furthermore, even in 3D integrated COTS devices with outputs as heavily post-processed as video files, it is possible to distinguish differences in each layer's response trends.

REFERENCES

This work references figures published under the Creative Commons Attribution-Share Alike 3.0 Unported License. https://creativecommons.org/licenses/by-sa/3.0/deed.en

[1]   Wikimedia Foundation, "Three-dimensional integrated circuit," 22 Jan. 2008. [Online]. Available: https://en.wikipedia.org/wiki/Three-dimensional_integrated_circuit. [Accessed 2020].

[2]   J. Knechtel et al., "Large-Scale 3D Chips: Challenges and Solutions for Design Automation, Testing, and Trustworthy Integration," *IPSJ Trans. on System LSI Design Methodology,* vol. 10, pp. 45-62, Aug. 2017.

[3]   R. Reif et al., "3-D Interconnects Using Cu Wafer Bonding: Technology and Applications," [Online]. Available: https://web.archive.org/web/20190715235342/https://pdfs.semanticscholar.org/e996/825ace86fa01a3492950f2d61fcbf780f515.pdf. [Accessed 2021].

[4]   Wikimedia Foundation, "3DS die stacking concept model," Wikipedia, 11 Nov. 2011. [Online]. Available: en.wikipedia.org/wiki/File:3DS_die_stacking_concept_model.PNG. [Accessed 2020].

[5]   N. Weiner, "Cybernetics: or the Control and Communication in the Animal and the Machine," *MIT Press,* no. ISBN 0-262-73009-X, p. xi, 1961.

[6]   P. E. Dodd and L. W. Massengill, "Basic mechanisms and modeling of single-event upset in digital microelectronics," *IEEE Transactions on Nuclear Science,* vol. 50, no. 3, pp. 583-602, June 2003.

[7]   L. W. Massengill, B. L. Bhuva, W. T. Holman, M. L. Alles and T. D. Loveless, "Technology scaling and soft error reliability," in *IEEE International Reliability Physics Symposium (IRPS)*, Anaheim, CA, 2012.

[8]   H. Barnaby and M. Marinella, "Total Ionizing Dose and Displacement Damage Effects in embedded Memory Technologies," in *Radiation Effects on Components and Systems Conference,*, 2013.

[9]   P. V. Dressendorfer, "Basic mechanisms for the new millennium," in *IEEE NSREC Short Course*, 1998.

[10]   R. A. Reed, "Fundamental Mechanisms for Single Particle-Induced Soft-Errors," in *Nuclear and Space Radiation Effects Short Course*, 2008 .

[11]   Wikimedia Foundation, "Bragg peak," 20 Dec. 2005. [Online]. Available: en.wikipedia.org/wiki/Bragg_peak. [Accessed 2020].

[12]   Wikimedia Foundation, "Bayer filter," 2 Apr. 2004. [Online]. Available: en.wikipedia.org/wiki/Bayer_filter. [Accessed 2020].

[13]   Wikimedia Foundation, "Demosaicing," 4 Nov. 2005. [Online]. Available: https://en.wikipedia.org/wiki/Demosaicing. [Accessed 2020].

[14]   O. Yadid-Pecht, "Active Pixel Sensor (APS) Design – From Pixels to Systems," Lecture.

[15]   A. Peters, "Image Histograms," 2019 Lecture. [Online]. Available: archive.org/details/Lectures_on_Image_Processing. [Accessed 2020].

[16]   T. Haruta et al., "A 1/2.3inch 20Mpixel 3-Layer Stacked CMOS Image Sensor with DRAM," *ISSCC,* 2017.

[17]   MIPI Alliance, Inc., "MIPI D-PHY," [Online]. Available: https://www.mipi.org/specifications/d-phy. [Accessed 2020].

[18]   H. Tsugawa et al., "Pixel/DRAM/logic 3-layer stacked CMOS image sensor technology," *IEDM,* 2017.

[19]   S. Buchner et al., "Variable Depth Bragg Peak Method for Single Event Effects Testing," *IEEE Transactions on Nuclear Science,* vol. 58, no. 6, 2011.

[20]   V. Lalucaa et al., "Single-Event Effects in CMOS Image Sensors," *IEEE Transactions on Nuclear Science,* vol. 60, no. 4, pp. 2494-2502, Aug. 2013.

[21]   V. Lalucaa et al., "Single Event Effects in 4T Pinned Photodiode Image Sensors," *IEEE Transactions on Nuclear Science,* vol. 60, no. 6, pp. 4314-4322, Dec. 2013.

[22] Wikimedia Foundation, "Lossless compression," 30 Oct. 2001. [Online]. Available: https://en.wikipedia.org/wiki/Lossless_compression. [Accessed 2020].

[23] Robert A. Weller, Marcus H. Mendenhall, Robert A. Reed, Ronald D. Schrimpf, Kevin M. Warren, Brian D. Sierawski, Lloyd W. Massengill, "Monte Carlo Simulation of Single Event Effects," *IEEE Trans. Nucl. Sci.,* vol. 57, no. 4, pp. 1726-1746, Aug. 2010.

[24] Robert A. Reed, Robert A. Weller, Marcus H. Mendenhall, Daniel M. Fleetwood, Kevin M. Warren, Brian D. Sierawski, Michael P. King, Ronald D. Schrimpf, Elizabeth C. Auden, "Physical Processes and Applications of the Monte Carlo Radiative Energy Deposition (MRED) Code," *IEEE Trans. Nucl. Sci.,* vol. 62, no. 4, pp. 1441-1461, Aug. 2015.

[25] J. F. Ziegler, J. P. Biersack, and M. D.Ziegler, "SRIM - The Stopping and Range of Ions in Matter," SRIM Co, 2008 ISBN 978-0-9654207-1-6.. [Online]. Available: www.srim.org.

[26] D. M. Fleetwood, "Evolution of Total Ionizing Dose Effects in MOS Devices with Moore's Law Scaling," *IEEE Trans. Nucl. Sci.,* vol. 65, no. 8, pp. 1465-1481, 2018.

[27] W. Cauer, "Theorie der linearen Wechselstromschaltungen," *Akademische Verlags-Gesellschaft Becker und Erler, Leipzig,* vol. I, 1941.

[28] Lawrence Berkeley National Lab, "https://cyclotron.lbl.gov/base-rad-effects/bragg-curves," LBNL, 17 Sept 2020. [Online]. Available: https://lh6.googleusercontent.com/QvKC9gJ_wtsmIkot4GvY4BgQfa6MmMa9gEzkxNp8BPdVml30UL66 k4Jrczav1ZDKMudFfq7BIhee9anXXBgaZm0ZBMU1Mju-MDzsW_9ZgxLMAsZ7V2fn=w1280. [Accessed 2020].

[29] Wikimedia Foundation, "Black box," 12 Oct. 2001. [Online]. Available: en.wikipedia.org/wiki/Black_box#History. [Accessed 2020].

[30] A. Peters, "RGB and HSV Color, part a: representations," 2019 Lecture. [Online]. Available: archive.org/details/Lectures_on_Image_Processing.

[31] Wikimedia Foundation, "Grayscale, subsection 'Converting color to grayscale'," 10 Feb. 2004. [Online]. Available: en.wikipedia.org/wiki/Grayscale#Converting_color_to_grayscale. [Accessed 2020].

[32] Wikimedia Foundation, "Cone cell," 18 Mar. 2004. [Online]. Available: https://en.wikipedia.org/wiki/Cone_cell. [Accessed 2020].

[33] C. Virmontois et al., "Radiation-Induced Dose and Single Event Effects in Digital CMOS Image Sensors," *IEEE Transactions on Nuclear Science,* vol. 61, no. 6, pp. 3331-3340, Dec. 2014.

[34] C. Virmontois et al., "Displacement Damage Effects in Pinned Photodiode CMOS Image Sensors," *IEEE Transactions on Nuclear Science,* vol. 59, no. 6, pp. 2872-2877, Dec. 2012.

[35] G. R. Hopkinson, "Radiation Effects in a CMOS Active Pixel Sensor," *IEEE Transactions on Nuclear Science,* vol. 47, no. 6, pp. 2480-2484, Dec. 2000.

[36] C. Virmontois et al., "Similarities Between Proton and Neutron Induced Dark Current Distribution in CMOS Image Sensors," *IEEE Transactions on Nuclear Science,* vol. 59, no. 4, pp. 927-936, Aug. 2012.

[37] Nikon, "Understanding ISO Sensitivity," [Online]. Available: https://www.nikonusa.com/en/learn-and-explore/a/tips-and-techniques/understanding-iso-sensitivity.html.