INVESTIGATION OF SYSTEM-LEVEL ESD INDUCED SOFT FAILURES

BY

NICHOLAS ADAM THOMSON

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2017

Urbana, Illinois

Doctoral Committee:

Professor Elyse Rosenbaum, Chair
Professor Deming Chen
Assistant Professor Songbin Gong
Professor José Schutt-Ainé

ABSTRACT

Electrostatic discharge (ESD) is a common phenomenon that can have negative implications for the performance of systems during operation. ESD is a brief, high-current stress that when applied to a functional system, can cause a variety of problems ranging from temporary system malfunction to permanent failure of components within the system and/or the system itself. Permanent failure of components, also known as hard failure, is a well-studied area among the ESD community with established practices in place for mitigating the stress and improving the reliability of the system. While further improvements with regard to hard-failure protection can always be made, this work will focus on the aspects of soft failures. Soft failures encompass any type of disruption to the system that, in general, can be recovered from, for example by power-cycling the system. An example of a soft failure could be the erasure of data within some memory elements of an embedded computer that cause programmed routines to fail. By restarting the system, the data will be regenerated and written into the memory elements and proper functionality is returned. While seemingly innocuous, the temporary failure of crucial systems, for example medical or automotive systems, could lead to severe consequences. This dissertation surveys the type of sensitive logic that could be susceptible to soft failures. Custom hardware which utilizes these elements has been designed and fabricated. Experiments with this hardware has yielded evidence of soft failures and lead to a preliminary understanding of the mechanisms responsible for those soft failures. While strong support for these hypothesized mechanisms has been obtained, future work must be completed to ascertain the validity of the conclusions drawn from the preliminary results.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION 1

CHAPTER 2: CHIP-BOARD ESD PROTECTION 5

2.1 IO Protection 5

2.2 Supply-Rail Protection 6

2.3 Board Protection 8

CHAPTER 3: SYSTEM-LEVEL TESTING 10

3.1 IEC 61000-4-2 Test Setup 10

3.2 Discharges to Mobile vs. Tethered Systems 16

3.3 Measurements during IEC 61000-4-2 Testing 17

CHAPTER 4: SOFT FAILURES OBSERVED WITHIN AN EUT 22

4.1 Test Chip 1 Overview 22

4.2 Test System 1 Overview 27

4.3 IEC 61000-4-2 Testing of Functional System 28

4.3.1 Glitches on Signal Lines 29

4.3.2 Static Latch Upsets 36

4.3.3 Upsets in Dynamic Logic 39

4.3.4 Bit Flips in the Shift-Register 42

4.3.5 Upsets to USB Functionality 48

4.3.5 Hard Failures 49

CHAPTER 5: ADVANCED NOISE MONITORS 51

5.1 New Test Chip Overview 52

5.1.1 Supply-Voltage Monitors 54

5.1.2 Static Latches 70

5.1.3 Dynamic OR Gates 72

5.1.4 Shift-Registers 74

5.2 Circuit Board and Test Set-Up Designs 77

5.3 IEC 61000-4-2 Testing of New Test System 81

5.3.1 Power Supply Noise Detected by Voltage Monitors 82

5.3.2 Static Latch Upsets 88

5.3.3 Bit Flips Due to Input Glitches 92

5.3.4 Effects of Reduced Supply Voltage 102

5.3.5 Zaps to a Protected System 104

CHAPTER 6: FUTURE WORK 108

REFERENCES 110

CHAPTER 1:

INTRODUCTION

Semiconductor devices may be subjected to electrostatic discharge (ESD) many different times during their fabrication, packaging, handling, installation, and use [1]. ESD is a brief stress on the device that can last hundreds of nanoseconds with current amplitudes reaching tens of amps. So much current being shunted through the semiconductor device causes heating that can lead to thermal failure. This current may also create potential differences on the device that can damage gate dielectrics of metal-oxide-semiconductor (MOS) structures. Because of this, devices must be designed to withstand ESD.

Typically, dedicated ESD protection devices are added to a system. These protection devices are placed within a semiconductor device, on a printed circuit-board within the system, or at both locations. Deciding where to place ESD protection is not straightforward. This ESD device needs to be sized large enough to handle the current levels that result at the targeted ESD level. Placing large devices on-chip results in additional area requirements that increase the costs of the integrate-circuit. Placing stand-alone protection devices on the printed-circuit board (PCB) requires additional parts to be ordered and large PCB area to accommodate the product. Minimizing the total system cost may lead to discontent between integrated-circuit (IC) manufacturers and system designers. Thus it is important to understand how the ESD stress propagates through the system and to identify potential vulnerabilities. This will allow the smallest amount of protection that is sufficient for adequate device operation during an ESD event.

Failure to adequately protect against the ESD stress can result in two types of failures: hard failure or soft failure. Hard failures are permanent catastrophic failure of the device. Hard failures have long been studied and mature methodologies for protection have been developed.

Soft failures represent a relatively new challenge to ESD designers. A soft failure is a type of disturbance to a system that is recoverable [2]. These could include erasure of data stored in a memory element that results in an inability of the system to continue operating, or the failure of an input/output (IO) cell to properly read in logical values. Due to the proliferation of embedded computers and sequential-logic, soft failures have quickly become a topic of interest in the ESD community. Some industries, e.g. automotive or medical, have a strict intolerance to soft failure [3]. In these industries, a soft failure could result in a temporary system failure that leads to a safety hazard.

Since large amplitude current or voltage is typically required to cause hard failures, direct current injection usually causes hard failures to occur along the main discharge current path. Thus, hard failures have traditionally been studied on the component level. This could mean that unpackaged wafers are probed and ESD current is injected, or the leads of package parts could be probed and stressed. Standards such as Human-Body Model (HBM) [4], and Charged-Device Model (CDM) [5] exist for characterizing a component's hard-failure level resulting from different ESD stresses. However, soft failures can occur from logic-level switching and thus could appear due to noise that couples to logic from distant zaps. For this reason, studying soft failures within a component under direct ESD injection is not sufficient. The interactions of several components within a functioning system must be considered. A system, in this context, is a functioning device that may consists of several packaged components on multiple PCB. Interconnects between components and ESD could be the means of coupling noise. Thus, system-

level ESD-generated soft failures are not necessarily restricted to devices along the main discharge path. Soft failures could result at any location within a semiconductor due to power noise generated at one IO.

In [6] a system is characterized for susceptibility to soft failures by scanning a probe across a test system while applying a stress in the form of radiated fields. This method allows for sensitive signal traces to be identified, but lacks the sensitivity required to localize the source of soft failures within an IC. Further, this method does not apply a stress to the system that is compatible with the widely used standard for system-level ESD testing (IEC 61000-4-2 [7]) and thus may not reflect the soft failures that would occur during such a test. In [8] and [9], a test system is stressed with injected ESD current while various traces are directly probed with a voltage probe. While this technique could be adapted to work with the IEC 61000-4-2 standard, the scope is once again limited to board-level analysis as it is impractical to probe traces within an IC. Also, the use of conductive probes will tether the system to earth ground which in turn can greatly affect the current path and lead to unrealistic upsets. In [10] and [11] direct probing was used in conjunction with accurate IC modeling to glean insight into what is occurring on the IC. However, this method still suffers from system alterations required for direct probing. In [12] a custom IC was fabricated with a high-speed oscilloscope incorporated on die. This method does allow on-chip noise waveforms to be measured, as opposed to the previous measurement practices. However, the test chip and test system used are highly specialized and thus not suitable for use outside of dedicated experimental work.

This work attempts to observe soft failures within a test system and uncover the mechanisms that cause these failures at the board and IC levels. Methods used within this work will strive to preserve the integrity of the testing standards and be applicable to all types of test

systems by avoiding the addition of conductive tethers to the test system and by avoiding the use of external measurement equipment during ESD testing.

CHAPTER 2:

CHIP-BOARD ESD PROTECTION

An important topic when evaluating any ESD failure is the method that was used to protect the various devices within the system from the ESD stress. Protection elements could be circuits within components, stand-alone circuits on a PCB, or discrete devices placed at key locations. The protection devices control the flow of current through the system and thus impact what failure mechanisms could arise. This chapter will introduce the various protection elements used within this work.

2.1 IO Protection

If ESD current reaches an integrated circuit (IC, chip) within a system, that current may enter into that component. If it enters through an IO, ESD protection devices are included to shunt that current away from the sensitive IO circuitry and clamp the ESD voltage to safe levels. In this work, diodes and silicon-controlled rectifiers (SCR) are used. Figure 1 illustrates the use of dual-diode (DD) protection on the left and SCR protection on the right.

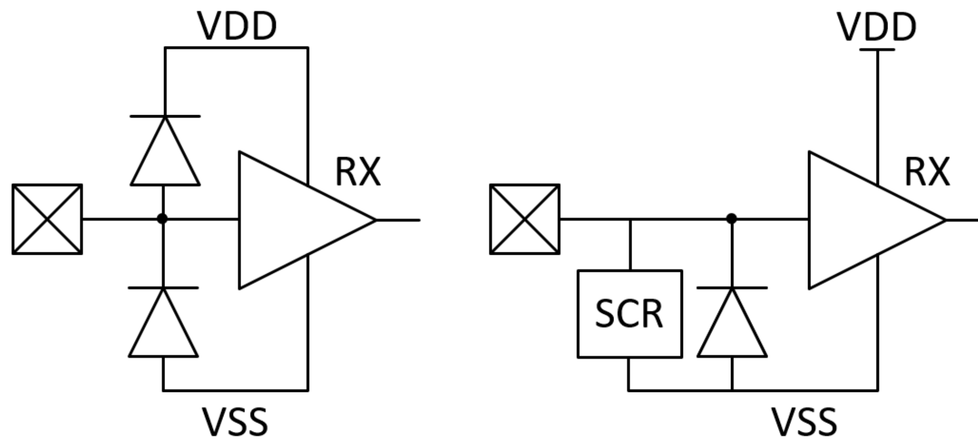

Figure 1: Schematic representation of dual-diode (DD) protection (left) and SCR protection (right).

Dual-diode (DD) protection consists of a “top diode” placed between the positive supply rail (VDD) and the IO and a “bottom diode” placed between the IO and the ground rail (VSS). With this configuration positive stress with respect to VDD will turn on the top diode and current will be shunted to the VDD rail away from the sensitive receiver (RX) circuitry. Similarly, for negative stress with respect to VSS will turn on the bottom diode and shunt current to VSS. SCR protection utilizes an SCR with some triggering circuit which will turn on the main protection device once the IO-to-VSS voltage achieves the designed threshold voltage. Thus, the SCR protects against positive ESD stress with respect to VSS. For negative stress with respect to VSS the “reverse-diode” will turn on and shunt the current. The SCR has a parasitic diode associated with its layout that will act as a reverse-diode, however, additional PN junctions can be added to augment the parasitic diode.

2.2 Supply-Rail Protection

Dual-diode and SCR protection elements shunt current to the supply rails, but an additional protection device must be included to maintain the chip’s power-supply voltage or chip-wide damage could result. Further, for component-level ESD, ESD stress can be applied between any IO, VDD, or VSS pin. Thus, for a positive stress applied between an IO and VSS, for example, DD protection will shunt the stress up to VDD, but cannot supply a return path to VSS. To satisfy these needs a rail clamp is used. Figure 2 depicts DD protection with the rail clamp between the power supply rails. When the rail clamp turns on, it shunts current from the VDD rail to the VSS rail. The parasitic diodes associated with all N-well and P-substrate contacts act together with any ESD diode placed between VDD and VSS to shunt current from VSS to VDD when required. The rail clamp also acts as the main protection device for stress injected between the chip’s power-supply pins.

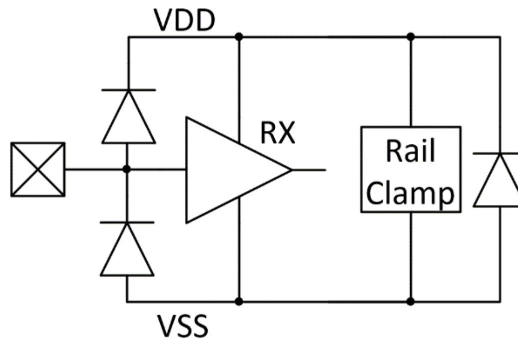

Figure 2: DD ESD protection with rail clamp.

In many ICs, several power domains are used to power different types of circuit blocks. For example, the IO ring of a chip may be powered with a higher voltage than the core circuitry. Further, the core logic may also have several domains to isolate the noise generated from fast-switching digital logic from the power supply used for analog circuitry. To ensure noise isolation, the references for these different domains must also be isolated. Having separate references can cause vulnerabilities for ESD stress. For example, if the power-supply pin for the IO domain is stressed with respect to the reference for core-logic domain, there is no dedicated path for that ESD current to pass from one domain to the other. Thus the current would likely cause a failure as it creates a path through the substrate. Anti-parallel diodes (APD) are used in this situation to allow a path between the reference nets of various power domains. Figure 3 illustrates a complete on-chip ESD protection strategy with APD. The APD allows current to flow in either direction once a potential difference develops between the two supply-domain reference-nets that is large enough to forward bias one of the diodes. Small magnitude noise will be insufficient to turn on a diode, allowing it to be isolated from the other net. During an ESD event, any potential difference between the two nets will turn on a diode before it becomes large enough to cause breakdown within the chip. The current flow will allow the potentials to equalize within safe levels.

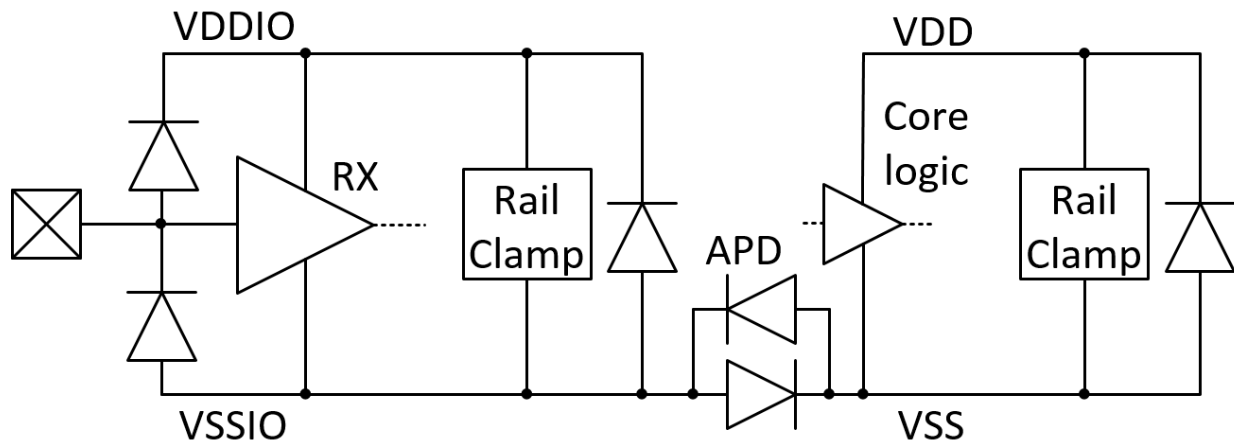

Figure 3: Cross-domain ESD protection. The APD circuit allows current to flow, in either direction, once a large enough potential difference develops between two domain reference nets.

2.3 Board Protection

So far only on-chip protection has been discussed. Stand-alone protection devices can be placed on the same printed circuit-board (PCB) as the main IC. These protection devices work to shunt current away from other sensitive devices on the PCB and reduce the stress seen by those devices. By using on-board protection, chip-level protection can be reduced (not removed) and free up the costly silicon area for other uses. A common dedicated on-board protection element is the transient voltage suppressor (TVS). For ESD applications, the TVS is usually made up of diodes. Shunt paths can be placed between the IO to be protected and board ground, board power, or both. Further, TVS devices can have unidirectional or bidirectional shunt paths. One example of a bidirectional shunt from the IO signal line to board ground is a Zener diode. During a positive ESD stress injected into the board IO trace with respect to ground, the Zener diode will turn on and start shunting current to ground once the breakdown voltage of the Zener diode is reached. For a negative zap, the diode will conduct once it is forward biased.

TVS diodes are selected based on their topology, ESD protection rating, and their input capacitance. The TVS topology dictates what stresses can be protected against. For example, a TVS between an IO and board ground cannot protect against a stress between the IO and the board power-supply unless other shunt paths exist between the board power-supply and the board ground. The ESD protection rating is the maximum stress that the TVS is designed to handle without failure. Finally, the input capacitance limits the types of signal lines the TVS can be placed at. Small form-factor TVS with very low input capacitance (< 1 pF) are useful for high-speed signal lines where larger devices with higher capacitance may only be suitable for low-speed lines.

CHAPTER 3:

SYSTEM-LEVEL TESTING

As mentioned in Chapter 1, system-level testing is required for the study of soft failures. Soft failures are, by definition, interruptions to a system that is functional. Several standards have been developed for system-level ESD testing. This work will focus on the International Electrotechnical Commission (IEC) 61000-4-2 standard [7]. This standard is widely used for the study of ESD to portable devices (such as cellphones, tablets, or laptops) as well as larger equipment (e.g. desktop computers).

3.1 IEC 61000-4-2 Test Setup

The IEC 61000-4-2 test setup is shown in Figure 4. The test setup consists a horizontal coupling plane (HCP) elevated above a ground plane on top of a non-conducting table. The HCP is electrically isolated from the ground plane except for the inclusion of two 470 kΩ resistors. These resistors slowly bleed charge from the HCP to the ground plane after zaps, but prevent the majority of current from flowing between the HCP and ground plane during the ESD event. Similarly, when the equipment under test (EUT) is tested in a mobile configuration, two 470 kΩ resistors are included to bleed any stored charge from the EUT to the HCP. The ground plane is connected to earth ground via a connection to the building's power outlet as well as tied to the ESD generator's ground via the ground strap from the ESD gun.

The EUT is placed on the HCP atop an insulating sheet. The insulator prevents any electrical connection between exposed metal on the EUT and the HCP. It also provides a more robust dielectric between the EUT and the HCP. The EUT can, and should, be tested in all of its common use scenarios. Most importantly certain mobile devices should be tested in a truly mobile

fashion, i.e. with no electrical connection to any other equipment or ground, or in a tethered fashion, e.g. plugged into a charging apparatus that may supply a connection to ground. The differences between mobile and tethered setup and the effect it has on the discharge waveform will be presented in Section 3.2.

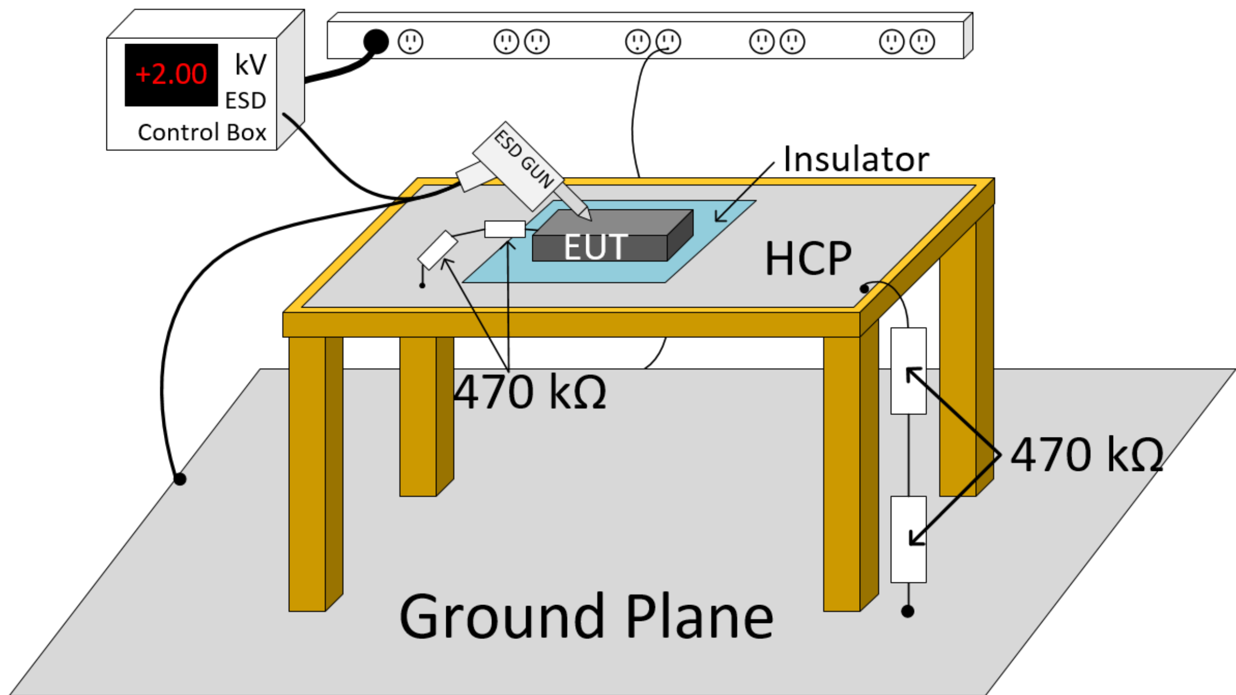

Figure 4: IEC 61000-4-2 test setup.

The ESD stress is applied via the ESD gun. The gun is charged from the HV ESD generator box and discharge via an internal relays when the gun is triggered. A simplified schematic of the gun's internal circuitry is shown in Figure 5. The HV DC supply represents the high-voltage supply within the ESD control box (see Figure 4). The storage capacitance, C_s, is charged through the charge relay and the charge resistance, R_c. Capacitance is also formed between the gun and objects in its surroundings. This capacitance, represented by C_d, will also charge along with C_s. When the gun is triggered, the discharge relay closes discharging C_s and C_d

through the discrete discharge resistance, R_d, and out of the gun tip. An inductance associated with the gun tip is included. It is used to meet the risetime specifications.

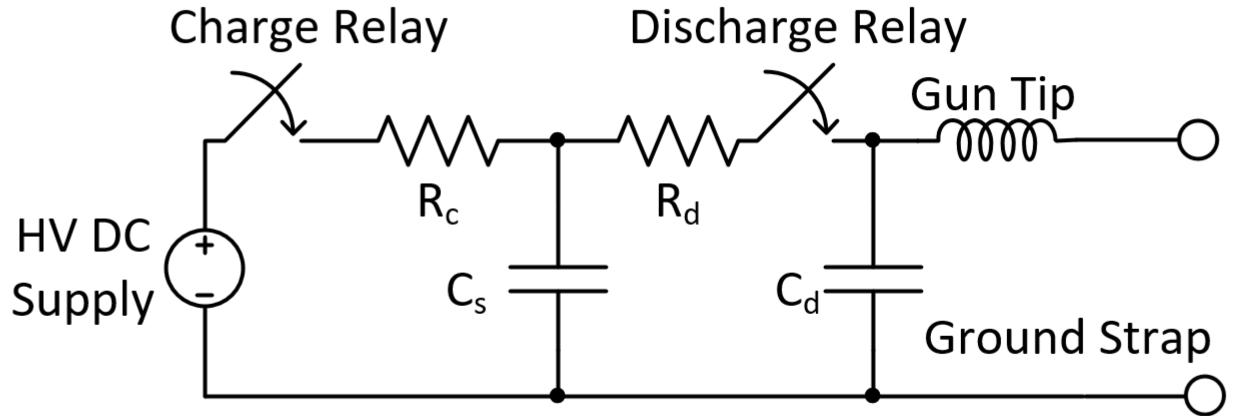

Figure 5: Schematic of circuitry within the ESD gun. R_c is the charging resistance, R_d is the discharge resistance. C_s is the discrete storage capacitance within the gun. C_d represents the capacitance that exists between the gun and its surroundings.

Several types of testing are specified in [7]. Contact discharge is carried out by bringing the gun tip into contact with the EUT prior to triggering the gun. This is the most severe zap, with respect to hard failures, as the current from the gun is directed into the EUT. A second type of discharge is an air discharge. During and air discharge, the gun is triggered before making contact with the EUT. Subsequently the gun is brought into close proximity with the EUT until a discharge is detected. A third type of discharge is an indirect discharge. These stresses are performed by contact discharge to the vertical coupling plane (VCP; not shown in Figure 4), or HCP. Similar to the HCP, the VCP is an isolated conductor that is vertically oriented and placed near the EUT. Since contact discharge represent the most severe stress, they are the only discharges that have been studied up to this point. The number of tests, stress level, and location of testing depend on the discharge type being used. For this study, contact discharges have been applied up to the failure level of the EUT (determined through measurement) to system ground and to signal lines. Further the EUT should be tested in all realistic application scenarios. For this

reason the EUT in this study is subjected to stress in a mobile (battery-powered) and tethered (powered by a DC supply connected to the building's power supply) configurations.

According to the specification [7], the waveform produced by the ESD apparatus is defined for zaps to a 2 Ω target. The target used in the standard is constructed of two concentric circular conductors separated by a resistive gap that provides the 2 Ω impedance. A 50 Ω conical adapter line is used to taper the form factor of the target down to a BNC connector, maintaining a 50 Ω characteristic impedance along its length. The voltage across the target is measured by an oscilloscope connected to the BNC connector. Since the ESD gun was professionally calibrated by the manufacturer a high-accuracy target is not needed for periodic verification of the ESD gun's calibration. Thus to avoid the high cost of a commercial calibration target, a simplified target was constructed for calibration verification. A cross section of the 2 Ω target manufactured for this study is illustrated in Figure 6.

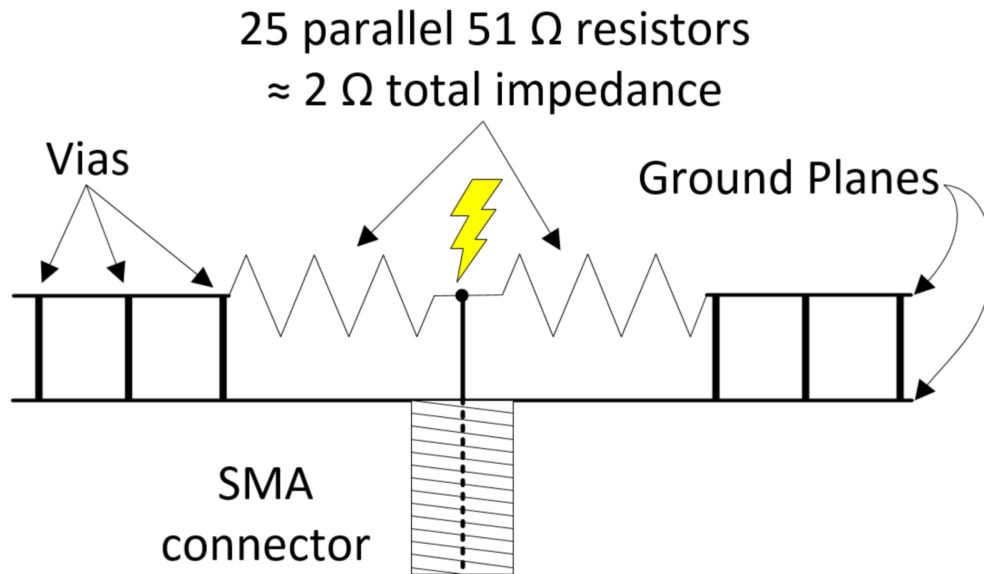

Figure 6: Cross section of the 2 Ω target. The resistors represent a disk resistor with 2 Ω total impedance. The center of the disk resistor is the zap location and is connected to the center conductor of the SMA connector. The top and bottom ground planes are connected together with vias and are connected to the SMA connector's outer conductor.

The target was fabricated from a two-sided copper clad board. A circle was etched into the top copper and 25 51-Ω resistors were soldered around the circle to form approximately 2 Ω total resistance with low series inductance. An SMA connector was soldered on the backside with the center conductor passed through the board to the center of the etched circle and the outer conductor soldered to the bottom ground plane. An array of holes were drilled through the ground planes and wires were inserted to act as vias connecting the two planes.

By connecting this SMA connector to an oscilloscope while zaps are applied to the center terminal, the discharge waveform from the gun can be captured. The standard recommends placing the oscilloscope inside a shielded enclosure, but this is not required unless measurements indicate that indirect coupling paths influence the calibration results. Figure 7 shows the waveform as defined by the specification and Figure 8 shows the waveform measured using the 2 Ω target. The parameters marked on the graph in Figure 7: I_p, I_{30}, I_{60}, and t_r are specified (within tolerances) by [7]. Table 1 indicates the values of these parameters for the four levels of contact discharge. Note that the zero reference for the time axis is defined as the point where the current waveform reaches 10 % of its peak (labeled as 10 % I_p in Figure 7). As a rule of thumb, the peak current, I_p, is 3.75 A/kV, the current at 30 ns, I_{30}, is 2 A/kV, and the current at 60 ns, I_{60}, is 1 A/kV. Comparing the waveforms of Figure 7 (illustrated) and Figure 8 indicate that the waveforms measured across the simplified 2 Ω target exhibit a similar shape, but have different time characteristics. These differences can be attributed to the simplifications made to the calibration setup.

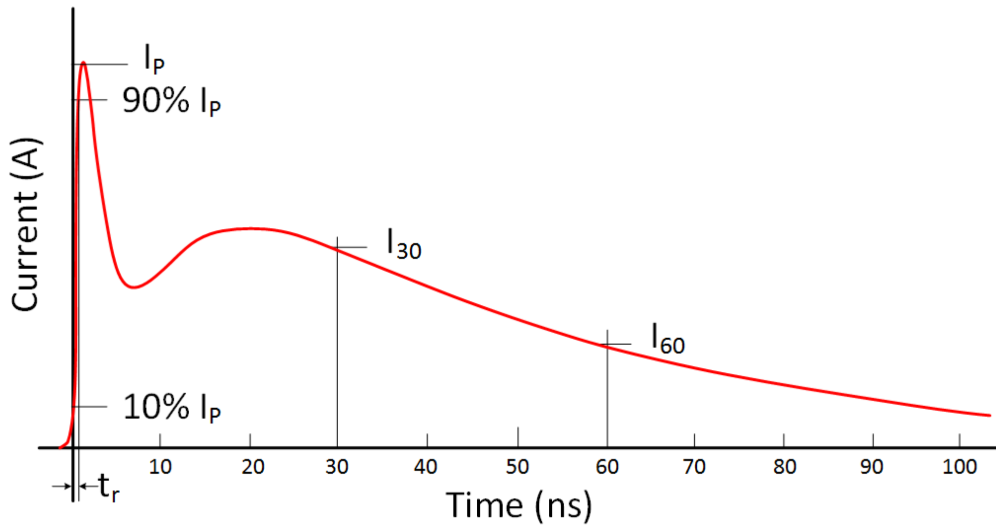

Figure 7: Discharge waveform from the ESD gun into a 2 Ω target. Important values that are specified in the IEC 61000-4-2 standard have been marked.

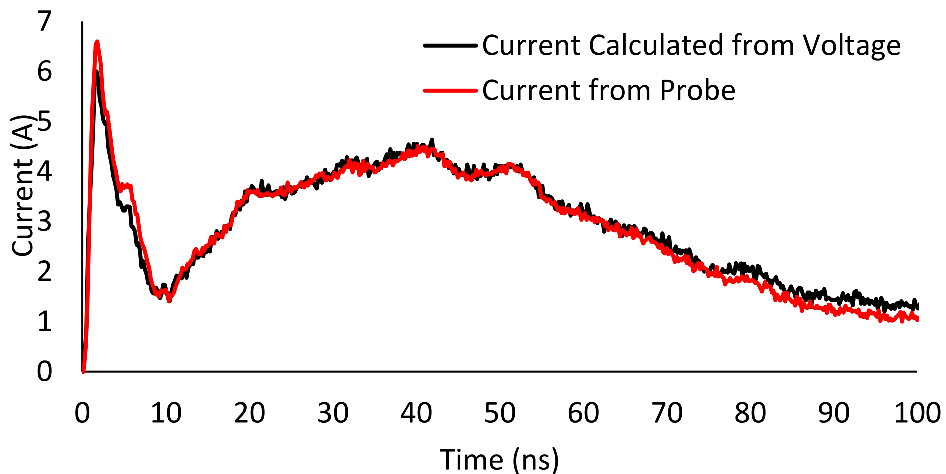

Figure 8: Current calculated from the voltage measured across the 2 Ω target compared to the current measured with a current probe placed around the gun tip during a 2 kV discharge.

Table 1: Waveform parameters specified for contact discharge into a 2 Ω target.

Level	Precharge Voltage V_{pre} (kV)	First Peak Current Ip (A), ±15 %	Rise Time t_r (ns), ±25 %	Current at 30 ns I_{30} (A), ±30 %	Current at 60 ns I_{60} (A), ±30 %
1	2	7.5	0.8	4	2
2	4	15	0.8	8	4
3	6	22.5	0.8	12	6
4	8	30	0.8	16	8

3.2 Discharges to Mobile vs. Tethered Systems

Since many systems today are designed to be completely mobile, testing needs to be carried out on systems in this configuration. On the flip side, mobile devices spend a good amount of time tethered to other devices, primarily for charging purposes. The design of a robust product thus requires that these systems be tested in a tethered configuration as well. While being produced from the same ESD gun, the waveforms delivered to a mobile system vary greatly from those delivered to a tethered system. Figure 9 shows the waveforms of discharges into a tethered and mobile EUT. The waveforms were measured using a current probe placed around the tip of the ESD gun during a 2 kV contact discharge. Measurement techniques will be discussed in more detail in Section 3.3. The two waveforms show a similar first peak, but vastly different second peak. The discharge path of the mobile device is simply made up of displacement currents between the EUT and the HCP or other surrounding objects. Since the capacitance associated with these objects is small, this discharge path has a high-pass filter characteristic. Thus only high-frequency components can travel to a mobile EUT. This can be seen in the mobile waveform of Figure 9 where the fast first peak and subsequent oscillations are observed. The first peak of the discharge waveform is similar to both the tethered and mobile EUT because the first peak is attributed to the parasitic capacitance between the gun and its surroundings (C_d in Figure 5), and the discharge path through displacement current which is common to both the tethered and mobile configurations.

The tether used to connect the EUT to the power supply is typically two wires, similar to a twisted pair, of considerable length (0.5 m). Due to its length and construction, the tether appears to the system as a highly inductive tie to the power supply. This inductance acts as low-pass filter, thus allowing lower-frequency components of the discharge to travel through the EUT. The main

charge transfer from the gun to the EUT is seen during the second peak which is created by the storage capacitor, C_s (see Figure 5), which is much larger than C_d and thus has a slower discharge. Further, the second peak of the waveform in Figure 9 looks much different than that of the specification shown in Figure 7. This is again due the inductance of the power-supply cable that connects the EUT to ground.

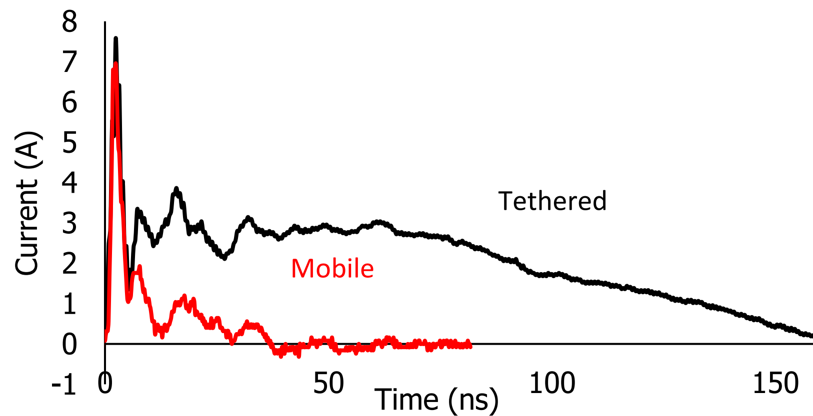

Figure 9: Comparison of discharge waveform into a mobile system and tethered system.

3.3 Measurements during IEC 61000-4-2 Testing

In order to study the cause of soft failures, it is useful to study the waveforms that occur during system-level ESD testing. The specifications of [7], do not provide methods for waveform measurement outside of the required measurements through the 2 Ω target. Two methods of making these measurements are considered: the use of a current probe around the gun tip, and direct voltage measurement within the system.

A current probe placed around a conductor will generate a voltage proportional to the current traveling through the conductor. By placing a current probe around the tip of the ESD discharge gun, the current that is discharged into the EUT can be measured. Current probes will add a small series impedance into the conductor which is being measured, but overall has little effect on the system. However, since a current probe can only detect changes in magnetic field,

the probe will have a high-pass filter characteristic. Within this work a Fischer F-65A current probe is used. This probe has a -20 dB gain with a bandwidth from 100 kHz to 1 GHz. To test the accuracy of the current probe during system-level ESD zaps, the current probe was used during a discharge to the 2 Ω target. The current through the 2 Ω target can be calculated from the voltage waveform obtained through the SMA connector (see Figure 6). This is compared to the current measured directly from the current probe in Figure 8. Overall the two waveforms show very good agreement. Differences can be seen in the first peak; these can be attributed to the assumption that the disk resistor is a perfect 2 Ω load. In reality the 2 Ω target is constructed with large ground planes which will form a parasitic capacitance, C_p, between the zap point and ground. This capacitance is in parallel with the 2 Ω target resistance and will result in a low-pass filter. This will affect the measured first peak especially, since the higher-frequency components will see a smaller impedance.

While the above method is simple it can only supply information about what current is entering the EUT. It is highly desirable to capture waveforms appearing within the system, for example at an IC pin. The current waveform at the IC pin may be different from that at the gun tip due to the filtering effect of the circuit board. Direct voltage measurements can be taken within an EUT by soldering a rigid coaxial cable to the desired test-point within the system. In order to ensure that the highest frequency content of the ESD event (0.8 ns rise-time) is captured, a 3 GHz oscilloscope with 50 Ω input impedance is used. Inserting a low-impedance (50 Ω) shunt to ground into a functional EUT would have a large impact on the system's performance, thus in order to reduce the current flow through the oscilloscope a high-impedance pick-off is used. It is important to note that added resistance will degrade the bandwidth of the measurement, thus the series resistance must be carefully chosen to prevent the measurement bandwidth from becoming

too small. Unlike current-probe measurements, direct voltage probing is highly invasive; the center conductor is isolated with a large resistance, but the ground of the cable is directly connected to the ground of the EUT.

In one experiment, an ESD gun was discharged into the ground plane on a special test board. In the center of this test board is a mock-IC whose IOs consist of SMD resistors connected between the IO pad and board ground. The mock-IC also has power supply pads and ground pads with ample decoupling capacitance placed nearby. This test board was placed on the IEC61000-4-2 test bed in a mobile configuration, so the current return path is through the various board-to-ground capacitances, e.g. board to table and table to ground plane. A high-impedance pick-off was used so that the induced voltage waveform at a mock-IC power pad could be measured using an oscilloscope; however, this introduces a second ground return path. ESD current was injected into one of the board traces that terminates at a mock-IC IO. During ESD, the incoming current wave initially "sees" only the near-end of the pick-off, which is a 450 Ω resistor in series with the 50 Ω coaxial line going to the scope. Depending on the length of the cable, the effects of the oscilloscope (or other measurement device) will appear after a delay. The waveform at the oscilloscope was recorded for two cases: (1) the oscilloscope chassis is connected to earth-ground via the power cable and (2) a large resistor is inserted in the oscilloscope ground return. In the two cases, the impedance of the ESD current return path to the gun is different. Both measurement results are shown in Figure 10. The two waveforms appear nearly identical during the first 50 ns of the discharge, but an additional current peak is observed when the oscilloscope is connected to earth-ground in the usual manner.

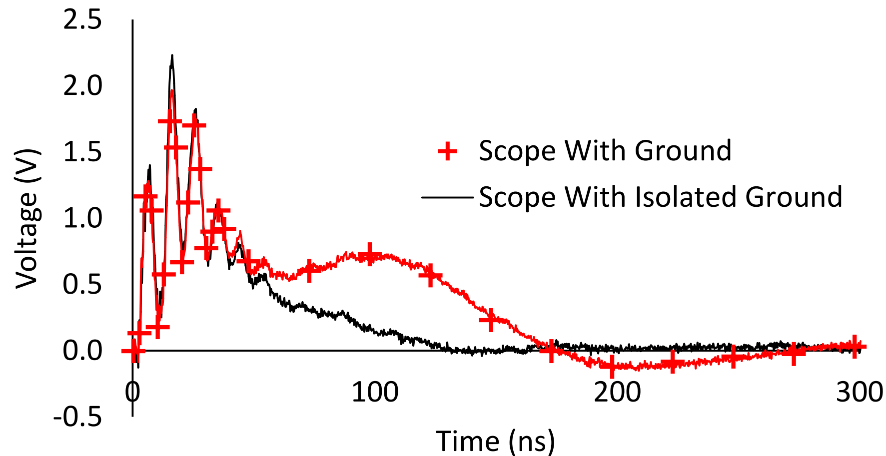

Figure 10: Direct probing of signals on a test board; traces are terminated to emulate an IC load. The current flowing to the grounded oscilloscope (proportional to the voltage waveforms shown) is much different than with a ground-isolated oscilloscope.

This additional peak is caused by excess current flowing through the oscilloscope's grounded chassis. The ESD current charges the board with respect to earth ground. When the oscilloscope's chassis is isolated, the current will leak through the oscilloscope and through the high impedance. However, when the oscilloscope's chassis is shorted to earth ground, a much larger current can flow through this short. To confirm this, the 2 Ω target was again zapped, this time the current probe was placed around the coaxial cable that goes from the target to the scope. This measurement yields the current waveform traveling along the shield of the coaxial cable, but not opposed by current traveling through the center conductor. Figure 11 shows the current injected into the 2 Ω target and current measured by the current probe. Almost all of the current associated with the second peak of the ESD discharge can be attributed to this current flowing through the oscilloscope cable, even with the oscilloscope isolated from earth ground. Once more, an additional current peak (around 140 ns) is seen in the case where the oscilloscope is not isolated. Thus it is clear that the addition of direct voltage measurements effectively tethers the system precluding it from use on mobile systems. Even if direct voltage measurements are applied

to tethered systems, the additional ground return path will affect the discharge path and could change the system's response.

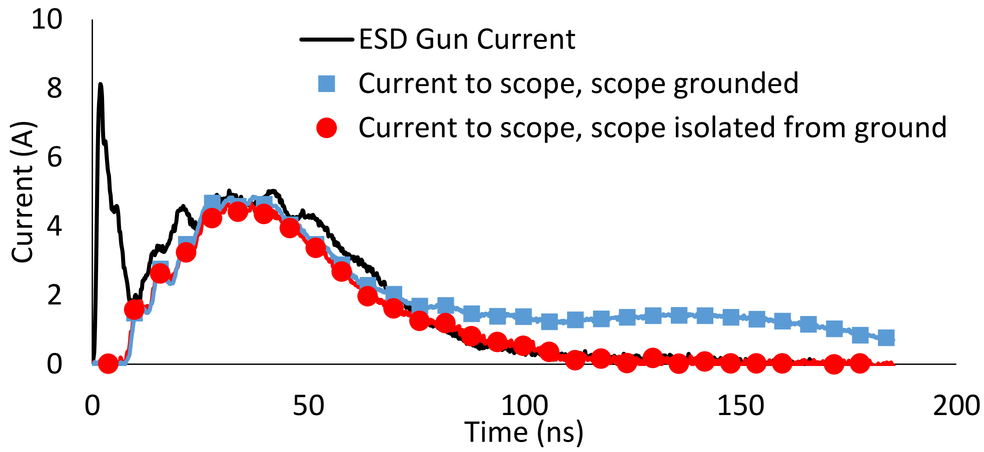

Figure 11: Current measured by a current probe placed around the coaxial cable from the 2 Ω target to the oscilloscope. Ideally the current to scope would be zero, but significant current is allowed to flow. Further, when the oscilloscope is grounded to earth ground, an additional peak is detected around 140 ns.

CHAPTER 4:

SOFT FAILURES OBSERVED WITHIN AN EUT

In order to study the mechanisms responsible for soft failure within a system during ESD stress, a suitable system must be acquired. While it is possible to purchase any consumer device and subject it to ESD stress, such experiments would be unwieldy for investigative work. Without detailed knowledge of the inner workings of the system, identifying the source of soft failures may prove difficult or impossible. Further, intricate systems such as cellphones, may have many complex circuits that could experience soft failure concurrently and obfuscate the results. Instead a custom test chip was fabricated in a 130-nm CMOS technology [13]. Further, to test this test chip in a functional state a corresponding test board was also designed and fabricated. Together this system becomes the EUT for the following investigations. By utilizing custom test equipment, detailed system design is on-hand facilitating a deeper knowledge of the mechanisms behind observed soft failures.

4.1 Test Chip 1 Overview

The test chip layout is shown in Figure 12. The test chip has two main power domains. The core circuitry and the SSTL IO test circuits lie within the 1.5 V domain, V_{DD}. Each V_{DD} cell contains a 2 mm active clamp. The 3.3 V supply, V_{DDIO}, provides power to all the CMOS IO circuits; each V_{DDIO} cell contains a 5.5mm active clamp. The rail clamp trigger circuits are optimized for power-on ESD [14], [15]. The V_{DDIO} bus was subjected to 100 ns power-on TLP (w.r.t. V_{SSIO}); leakage current measurements indicate that at least one of the several V_{DDIO} clamps distributed around the pad ring suffers hard failure when 18 A is injected onto the pad ring. In an

analogous measurement, one of the several V_{DD} clamps fails when 8 A is injected between the V_{DD} and V_{SS} buses. Each V_{SSIO} cell contains 500 μm perimeter anti-parallel diodes to V_{SS}.

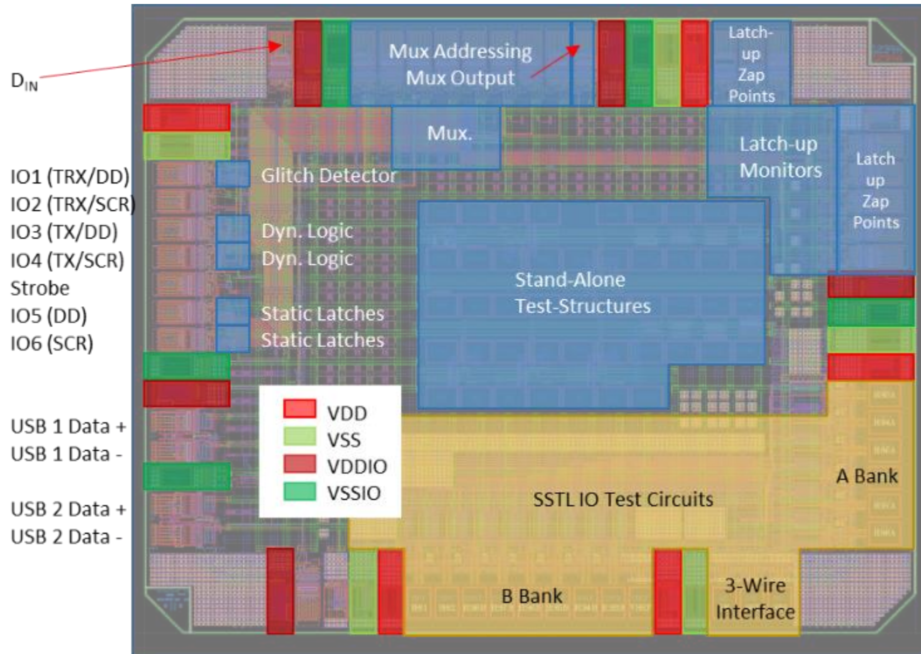

Figure 12: Chip layout view.

In Figure 12, the pad cells labeled IO1 through IO6 are CMOS IOs that connect to external pins which will undergo ESD zapping; these pads have either dual-diode or DTSCR protection, as indicated by the labels "DD" and "SCR." DTSCR protection devices are augmented by a reverse diode to achieve good protection against stress of both polarities. All six of these pad cells contain a bi-directional transceiver ("TRX"); however, by selective exclusion of key interconnects, IO3 and IO4 were configured as transmitters ("TX"), and IO5 and IO6 were configured as ESD-protected dummy cells. These dedicated zap IOs, as well as all of the 3.3 V control signal receivers, include a Schmitt trigger for noise resilience.

Two USB transmitters are included on-chip; each outputs a differential square-wave near the 480 Mbps operating speed of USB 2.0. Each USB transmitter is driven by an independent on-

chip ring oscillator. The output pins of USB transmitter #1 have dual-diode ESD protection, while the output pins of transmitter #2 have DTSCR protection.

Fourteen of the pad cells contain bidirectional SSTL IO test circuits [16]. As indicated on the bottom-right side of Figure 12, the SSTL IOs lie in two groups, labeled bank A and bank B. Each side contains six working SSTL transceivers; additionally, two non-functional SSTL transceivers were included on bank B for a component-level ESD study. The operating state of the SSTL block is programmable, allowing for selection between TX and RX modes and for adjustment of the on-die termination resistance. The SSTL block is programmed by an on-chip shift-register via a three-wire interface (serial data, serial clock and latch enable). The three-wire interface receivers are powered from the low-voltage domain and consist of a simple CMOS buffer with DD protection. Each time this block is reconfigured, the chip power consumption is changed. Banks of logic gates are included on the test chip. The stored data may be read-out before and after an ESD zap, in order to detect ESD-induced logic errors. The output of the selected logic gate is connected to an output pin through the on-chip multiplexer ("MUX").

In Figure 12, each of the two blocks marked "Dyn. Logic" contains eight Domino Logic OR gates [17]; the schematic is shown in Figure 13(a). These logic gates are labeled 0-15, corresponding to their addresses. Table 2 describes these gates. In Figure 12, the two blocks labeled "Static Latches" each contain 16 latches, labeled 16-47, corresponding to their addresses. Each static latch consists of a bistable element (cross-coupled inverters) accessed via a transmission gate; the schematic is shown in Figure 13(b). The latches vary in terms of the inverter sizing, layout orientation, and the distance from the IO ring [17]. Table 3 describes the various latches.

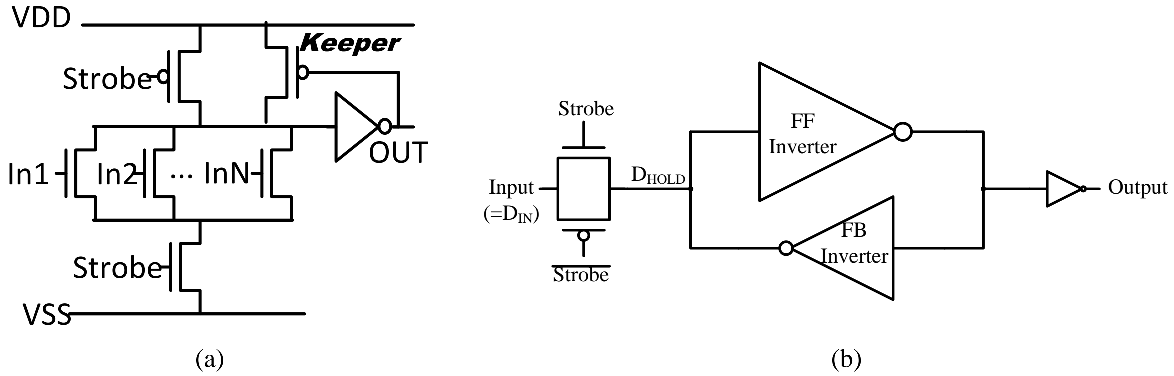

Figure 13: Dynamic logic (a) and static latch (b) topology.

Table 2: Dynamic logic gates are located near IO3 and IO4. Each of these gates is a Domino OR circuit. The MUX address selects the indicated gate's output for read-out. In the column marked "Input," an entry of "Tied low" indicates that all the gate's inputs are tied low. An entry of "D_{IN} pin" indicates that N-1 inputs are tied low, and one input is driven through a buffer chain from the D_{IN} pin.

MUX Address		*N*: # of inputs	Input
IO3 (DD)	IO4 (SCR)		
0	8	4	Tied low
1	9	16	Tied low
2	10	4	Tied low
3	11	16	Tied low
4	12	4	Tied low
5	13	4	Tied low
6	14	4	D_{IN} pin
7	15	16	D_{IN} pin

Table 3: Static latches are located near IO5 and IO6. The MUX address selects the indicated gate's output for read-out. The distance indicated is given in arbitrary units, with larger numbers indicating a greater distance.

MUX Address		Distance from IO	Relative size of FF inverter	Relative size of FB inverter
IO5 (DD)	IO6 (SCR)			
16	32	1	2	1
17	33	1	4	2
18	34	1	6	3
19	35	1	6	2
20	36	1	6	1
21	37	1	6	3
22	38	1	6	2
23	39	1	6	1
24	40	2	2	1
25	41	3	2	1
26	42	4	2	1
27	43	5	2	1
28	44	6	2	1
29	45	7	2	1
30	46	8	2	1
31	47	9	2	1

Data are written into the logic circuits described above only when an external strobe signal is set to the appropriate value; the strobe is applied at the pad labeled "Strobe." When strobe is low, the static latches hold their previous data and the dynamic logic gates are in evaluate mode.

A "glitch detector" was included because coupled noise at an input pin can produce logic errors [13]. The glitch detector circuit is shown in Figure 14. This circuit was placed inside the IO1 pad cell and will detect a logic level change at the IO1 pin due to ESD discharges elsewhere in the system. The glitch detector is active when the OE control signal is low. When the logic level at the input changes, GD switches from low to high, maintaining this value until the latches are reset. The rightmost multiplexer allows the IO to output either GD or D_{OUT}. The leftmost multiplexer keeps the signal at the input of the glitch detector constant when OE is high (activating the output driver); without it, the signal output by the glitch detector would be fed back into the glitch detector.

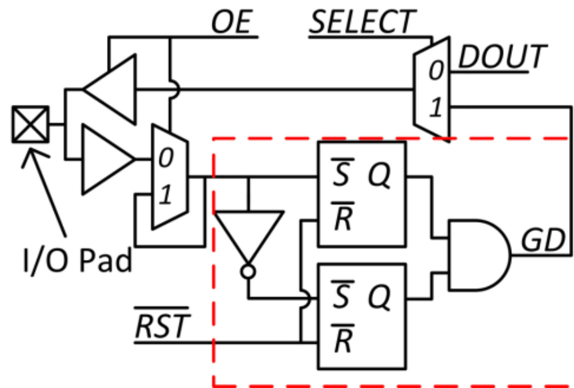

Figure 14: Glitch detector schematic.

ESD-induced logic upset (data change) can be detected by reading out the contents of each static latch and dynamic OR gate after an ESD zap. The output of a USB transmitter can be observed before, during and after an ESD zap to observe the effect of ESD on a high-speed, free-running circuit.

4.2 Test System 1 Overview

The system-under-test is a four-layer FR4 circuit board, shown in Figure 15. The test system can be powered by a DC supply or by a battery pack (pictured). Independent linear low-dropout (LDO) regulators supply each of the power domains: V_{DDIO} (3.3V), V_{DD} (1.5V), and V_{DDLED} (3.3V), the last of which provides power for the LEDs described below. Adequate decoupling capacitance is included on each of the power nets, including large tantalum capacitors and smaller ceramic capacitors. SMD decoupling capacitors were placed near the chip, following best practices.

An on-board LED provides a visual readout of logic high signals from the multiplexer output and a second LED enables data readout from IO1, the glitch detector circuit's output. The LEDs are driven by a separate buffer IC that is powered from the V_{DDLED} domain.

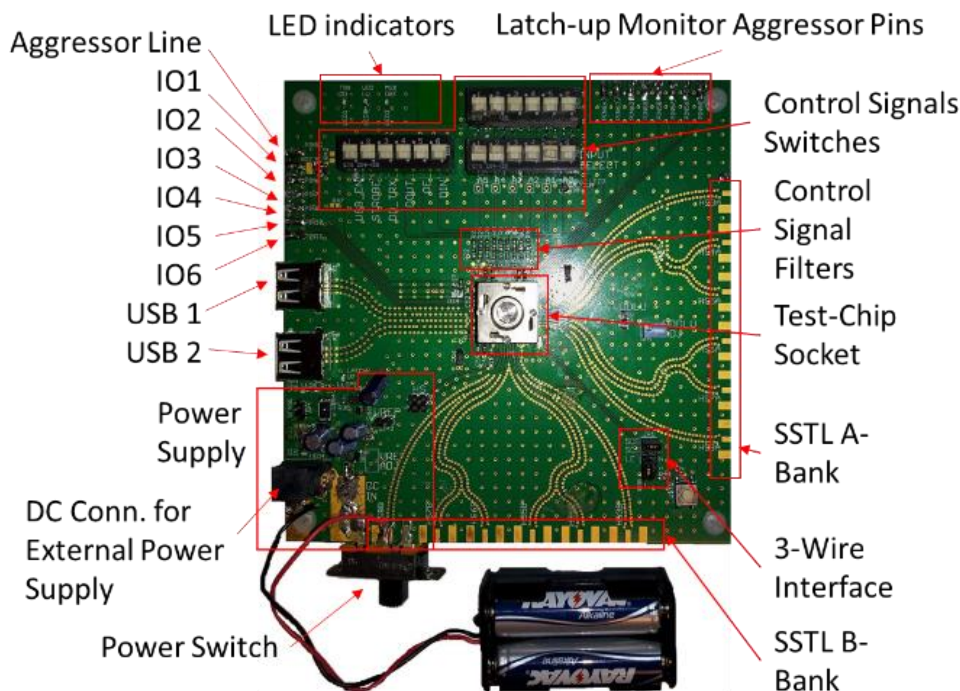

Figure 15: Board photograph.

Some of the chip IO pins are intended to undergo ESD zapping; test points are placed at the board-edge ends of traces that terminate at these pins. A zap is initiated by a contact discharge

to a test point. Test points provide direct access to IO1-IO6. An additional test point is connected to a trace that is adjacent to the signal trace that goes to IO1. This neighbor line is referred to as the “aggressor line;” the aggressor line is terminated near the test chip by a short circuit to ground. Zaps are applied to the aggressor line in experiments that utilize the glitch detector inside IO1; in these experiments, IO1 is set to receive mode and its input is set to either logic low or logic high. The logic input for IO1 is supplied by an on-board buffer IC. IO1 can also be directly connected to a test-point on the board’s edge. The desired input source for IO1 is selected by soldering a 0 Ω resistor to one of two pads on the board.

The control signal switches drive the multiplexer address pins and the on-chip control lines, and are used to input data to the latches and dynamic logic. Address and control pins are not intended to be zapped; robust signal filters are placed on board near these pins of the test chip to minimize ESD-induced disturbances.

A 0.1 Ω precision resistor is inserted in series with the V_{DD} on-board voltage regulator. A multimeter is used to measure the voltage across this resistor. This allows the quiescent current draw (I_{DDQ}) to be measured before and after an ESD gun discharge. I_{DDQ} changes with the operating state of the SSTL block. An unintentional change in the circuit operating state constitutes a soft failure.

4.3 IEC 61000-4-2 Testing of Functional System

ESD discharges were applied to the powered-up system according to the IEC61000-4-2 standard. To observe the effects of an earth-ground connection, the system was tested in a mobile configuration (i.e. powered from a battery pack) and in a tethered configuration (i.e. powered from a grounded DC supply). Gun discharges were applied to several key locations on the board, which varied in accordance with experiment being carried out.

4.3.1 Glitches on Signal Lines

The first experiment is designed to identify the conditions under which ESD-induced coupled noise at an IO pin is large enough to produce a logic error. The glitch detector, located inside IO1, was activated, and the IO was set to receive mode. The input to IO1 was driven through a buffer whose input is shorted to either supply rail by one of the control switches. Gun zaps were applied to the aggressor line adjacent to the IO1 signal trace. A minimum of five zaps were applied for each test and the test was repeated on a second test chip. Finally, the entire experiment was repeated using a tethered configuration by powering the test system from a precision DC supply. Table 4 summarizes the occurrence of detected glitches. The following syntax is introduced to ease this discussion. For a signal line that is initially driven low, a glitch

Table 4: Occurrence of glitches at IO1 when zaps are applied to the aggressor line. The percentage of zaps resulting in a detected glitch is indicated. A dash indicates that no tests were performed at that particular voltage level. The EUT is tested in both mobile and tethered configurations.

Aggressor Line Zaps					
		Mobile		Tethered	
$\lvert V_{pre}\rvert$ (kV)	Glitch Type	Positive Zap Upsets	Negative Zap Upsets	Positive Zap Upsets	Negative Zap Upsets
0.5	(0→1)	0%	0%	0%	0%
0.55	(0→1)	0%	20%	-	20%
0.57	(0→1)	0%	80%	-	60%
0.6	(0→1)	0%	100%	0%	80%
0.7	(0→1)	0%	100%	0%	100%
0.8	(0→1)	0%	-	0%	-
0.85	(0→1)	0%	-	60%	-
0.9	(0→1)	40%	-	100%	-
0.95	(0→1)	100%	-	100%	-
≥ 1	(0→1)	100%	100%	100%	100%
0.5	(1→0)	0%	0%	0%	0%
0.55	(1→0)	0%	-	-	-
0.6	(1→0)	40%	-	60%	0%
0.65	(1→0)	100%	-	-	-
0.7	(1→0)	100%	0%	100%	0%
0.8	(1→0)	-	0%	100%	0%
0.9	(1→0)	-	0%	100%	80%
0.95	(1→0)	-	100%	-	-
≥ 1	(1→0)	100%	100%	100%	100%

occurs when that line is subsequently driven high (perhaps briefly). This will be denoted as (0→1). The polarity of the zap that caused the glitch will be prefixed. Thus, +(1→0) indicates that a positive zap caused a signal level to go from high to low. From Table 4, it is observed that the minimum pre-charge voltage at which -(0→1) glitches occur is lower than that of +(0→1). Conversely, +(1→0) has a lower threshold than -(1→0). This indicates that the trace to IO1 is inductively coupled to the aggressor line. The initial di/dt of the rising edge of positive zaps (or falling edge of negative zaps) induces a current in the IO1 signal line that flows out of the chip (or into the chip for falling edges). An example of this coupled current can be seen in Figure 16.

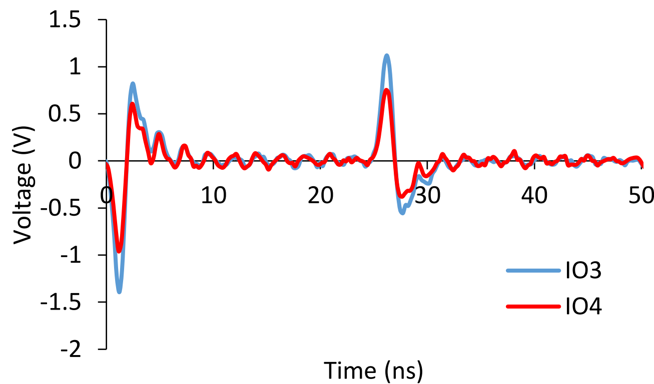

Figure 16: Coupled waveform observed on an unpopulated board at the IO1 pad location. 100 V TLP discharges were applied to IO3 and IO4 as indicated. The TLP pulse had 25 ns width and 1ns rise/fall times.

Positive TLP current pulses were injected into IO3 and IO4 on an unpopulated test board. The voltage induced at the IO1 pad was monitored. Large negative peaks are observed for the leading (rising) edges of the TLP pulses; conversely, large positive peaks are observed for the trailing (falling) edges. Despite the susceptibility for +(1→0) and -(0→1) glitches, +(0→1) and -(1→0) glitches do still occur; just at higher stress levels. This suggests that the second edge of the fast transient, while able to cause glitches, is less severe than the first edge. The presence or absence of a tether has no apparent effect on the glitch occurrence thresholds. Thus the amount of delivered charge has little impact on the occurrence of glitches, but rather the first peak of the IEC pulse, which is common to discharge waveforms into both tethered and mobile configurations

(see Figure 9), is the source of the noise. This too is expected as the first peak has much larger di/dt than the second peak and it is the di/dt that is responsible for inductive coupling.

In addition to inductive coupling, there is another possible source for the observed glitches. ESD noise may have induced a logic error at the buffer IC that drives this trace and the pin. To ascertain whether this contributes to the cause of the glitches, an additional experiment was conducted. IO1 was driven low either by the buffer IC (as above) or via a small-valued pull-down resistor. However, since the glitch detector's input and output share a pin, a suitable pull-down resistor must not overly load the glitch detector's output driver. A 47Ω resistor was used, which allows for a 0.7V output-high voltage and a 0V output-low voltage. Three zaps were applied to the aggressor line and again to the USB shield for each configuration and the glitch detector was monitored after each. Table 5 summarizes the results of this experiment. The occurrence of glitches is not reduced by eliminating the buffer IC, indicating that noise is coupled directly to the test chip rather than indirectly to the test chip via the buffer IC.

Table 5: Glitch detection results with different signal trace drivers. IO1 is being driven low by either a buffer or a resistive pull-down. Each entry represents the percentage of +(0→1) glitches detected out of three zaps.

V_{PRE} (kV)	Zaps to Aggressor Line	
	Buffer IC	Resistive Pull-Down
0.2	0 %	0 %
0.5	0 %	0 %
1	0 %	33 %
1.5	66 %	100 %
2	100 %	100 %
V_{PRE} (kV)	Zaps to USB Shield	
	Buffer IC	Resistive Pull-Down
4	0 %	0 %
5	0 %	0 %
6	0 %	0 %
7	33 %	0 %
8	33 %	0 %
9	0 %	0 %
10	0 %	33 %
11	0 %	66 %
12	66 %	100 %

ESD zaps were also applied to the signal traces for IO2 through IO6. In this case, ESD current will enter the test chip, in contrast to the previous experiment. The results are shown in Table 6; entries are grouped by their ESD protection methods. There is a proximity effect; when zaps are applied to signal traces that are progressively farther from IO1 (while comparing similar protection devices) a higher pre-charge voltage is needed to cause an upset (glitch) at IO1. This result was to be expected since inductive coupling decreases with increased distance between aggressor and victim. For discharges into IOs, upsets occur at even lower pre-charge voltages than in the previous experiment in which the "aggressor line" was zapped. This is attributed to the presence of additional noise mechanisms; when an IO pin is zapped, ESD current enters the chip produces supply noise and/or couples strongly to neighboring bondwires. Considering zaps to IOs regardless of the protection devices used as the distance from IO1 to the zap pin increases, the probability that a glitch is induced does not decrease in a perfectly monotonic fashion. This is attributed to the differences in the ESD protection used at the various IOs. IO3 and IO5 have DD protection while IO4 and IO6 have SCR protection. During an ESD event, the input impedance at

Table 6: Occurrence of +(1→0) glitches at IO1. Zaps were applied to the indicated IO. The EUT was in its mobile configuration. As zaps are applied farther from the IO1 the threshold for upsets increases. Only IOs with the same protection (SCR) are compared.

V_{pre} (kV)	Zapped IO#				
	SCR Protection			DD Protection	
	2	4	6	3	5
0.2	20%	0%	0%	0%	0%
0.3	100%	40%	0%	80%	0%
0.4	100%	100%	0%	100%	0%
0.5	100%	100%	0%	100%	0%
0.6	-	-	0%	-	0%
0.7	100%	100%	80%	100%	0%
0.8	-	-	100%	-	100%
0.9	-	-	100%	-	100%
1	100%	100%	100%	100%	100%
2	100%	100%	100%	100%	100%

the pin (impedance from the pin to the board ground) is different for the cases of DD and SCR protection. The input impedance affects the ESD current waveform at the pin and therefore the magnitude of the noise coupled to IO1 and the glitch detector.

The schematic of the chip-level IO being zapped and the victim IO (e.g. Strobe) where glitches occur is shown in Figure 17(a). When ESD current is injected into the zap IO, a glitch that results from the induced noise may be detected at the victim IO. Such noise may drive the victim IO temporarily to logic high or logic low. Figure 17 depicts the rising-edge of an ESD discharge into an IO. Two on-chip mechanisms work in tandem to cause a (1→0) glitch. The first mechanism is induced noise on the on-chip supply rails. For this example, the stress current must pass through the ESD protection to the VDDIO supply rail. The dV/dt caused by this current will initially trigger the rail clamp to turn on. The rail clamp and on-chip decoupling capacitance will both provide paths to the VSSIO supply rail. This current must then flow out of the chip through the VSSIO bondwires to board ground. After some time, the current will begin to flow more readily through the VDDIO bondwires and return to board ground via the decoupling capacitance. The change in current flowing through the bondwires associated with either the VDDIO or VSSIO pins will cause a voltage drop to form across the bondwires. This leads to a transient decoupling effect between the chip and board VDDIO/VSSIO lines due to the inductance of the bondwires. Figure 17(b) illustrates a simplified depiction of this decoupling. The potential on the VDDIO and VSSIO rails elevates w.r.t. board VSSIO causing the switching thresholds, V_{M+} and V_{M-}, of the Schmitt trigger IO to elevate as well. However, the Strobe signal potential, which is generated on the board, is not affected by this on-chip supply decoupling. The second mechanism is coupling of noise to neighboring IOs. In the example of Figure 17, the magnetic field generated by the initial edge of the positive discharge will cause a current, I_{VICTIM}, to flow out of the Strobe

line. This current will ultimately work to discharge the parasitic capacitance associated with the receiver circuitry. Figure 17(b) illustrates the resulting change in the on-chip Strobe signal level w.r.t. board VSS. As mentioned previously, for traces that run parallel to each other on-board, the ESD current inductively couples to neighboring traces as well, and can further increase the noise seen at the victim IO. Thus the drop in the Strobe signal and the increase of the switching thresholds both work together to cause a +(1→0) glitch for the leading edge of this example stress. Following a similar procedure for the other cases: +(0→1) during the trailing edge, -(0→1) during the leading edge, and –(1→0) during the trailing edge, all yield similar results.

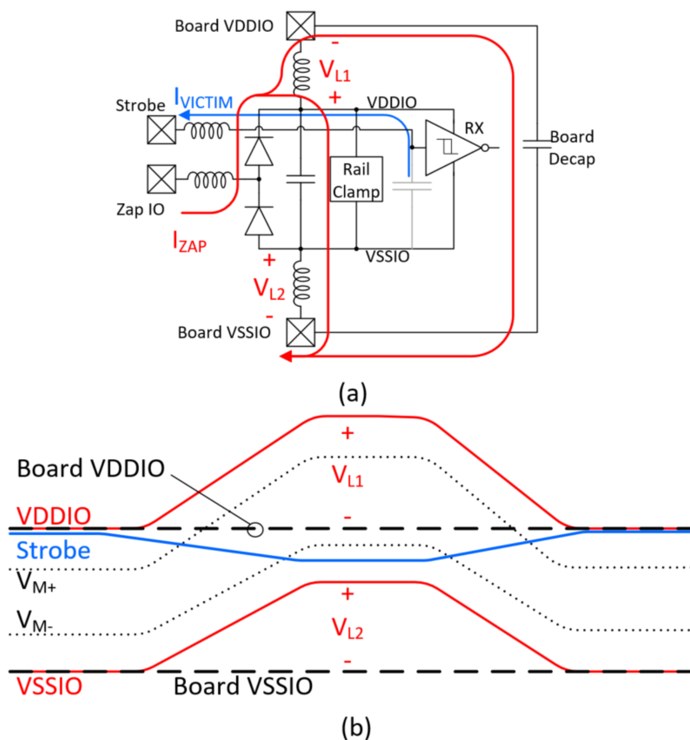

Figure 17: (a) Schematic representation of the aggressor and victim IO on chip. Due to the bondwire inductance, injected ESD current (I_{ZAP}) can cause chip VSSIO and VDDIO to rise w.r.t. their board levels. Further, I_{ZAP} can induce a current in nearby signal lines (I_{VICTIM}). (b) VDDIO and VSSIO are elevate w.r.t. to board ground due to the voltage drop across the bondwires. Strobe is pulled low due to the current exiting its pin. Together Strobe is pulled below the lower threshold (V_{M-}) of the receiver.

Gun zaps were next applied to the USB connector shield. (Note that metal connector shields undergo contact discharge during IEC61000-4-2 testing of a fully assembled, enclosed system.) The connector shield is tied to the ground plane of the system. The USB connector is not

nearby the glitch detector. Zaps were applied to a USB connector shield and the percentage of glitches observed are recorded in Table 7. A minimum of five zaps were conducted for each test case and the results from two test chips were tabulated. The failure threshold, i.e. minimum pre-charge voltage that induces a glitch, is much larger than in the previous experiments. As in the previous experiments, -(0→1) glitches have a lower threshold than -(1→0), and +(1→0) glitches have a lower threshold than +(0→1) glitches. In this experiment, the ESD current is injected to the board's ground plane, and thus the glitches are not the result of inductive coupling between parallel traces, as had been the case in the previous experiments. There are two possible mechanisms for the glitches. Noise may be induced on the signal line by the magnetic field radiated from the gun tip. Or, the zaps to the ground plane might be causing the ground potential to bounce. These potential fluctuations would be strongly coupled to the power plane by the large

Table 7: Occurrence of glitches when zaps were applied to a USB shield (board ground). The EUT was tested in both mobile and tethered configurations.

USB Shield Zaps					
$\lvert V_{pre} \rvert$ (kV)	Glitch Type	Mobile		Tethered	
		Positive Zap Upsets	Negative Zap Upsets	Positive Zap Upsets	Negative Zap Upsets
≤ 4	(0→1)	0%	0%	0%	0%
5	(0→1)	0%	0%	0%	0%
6	(0→1)	0%	33%	0%	0%
7	(0→1)	0%	66%	0%	100%
8	(0→1)	0%	100%	0%	100%
9	(0→1)	0%	100%	100%	100%
10	(0→1)	0%	100%	100%	100%
11	(0→1)	33%	100%	100%	100%
12	(0→1)	100%	100%	100%	100%
≤ 4	(1→0)	0%	0%	0%	0%
5	(1→0)	100%	0%	0%	0%
6	(1→0)	100%	0%	40%	0%
7	(1→0)	100%	0%	100%	0%
8	(1→0)	100%	0%	100%	0%
9	(1→0)	100%	0%	100%	40%
10	(1→0)	100%	0%	100%	100%
11	(1→0)	100%	0%	100%	100%
12	(1→0)	100%	100%	100%	100%

on-board decoupling capacitance, resulting in little change in the effective supply voltage given by $V_{DDIO} - V_{SSIO}$. There would be relatively little coupling to the IO1 signal trace and thus its potential would be stationary, at least initially. A positive zap would elevate the ground reference for VDDIO, causing a logic-high input signal to temporarily appear low, while for large amplitude negative zaps, a logic-low input signal could briefly appear to have switched high.

There is no systematic difference between the results for the tethered and mobile configurations of the EUT, except that the tether tends to reduce the margin between positive and negative glitch thresholds. The lack of a strong effect is not unexpected. The tether connects the board ground plane to the power supply ground through the series inductance of the supply cable; this cable has a high impedance on the short time scale during which the ground bounce is severe.

4.3.2 Static Latch Upsets

The static latch test circuits are programmed using signals DIN and Strobe. The noise on the corresponding on-board traces is minimized by the use of an RC filter and Zener diode; furthermore, these traces are not subjected to contact ESD. Zaps are applied to IOs 1 through 6 and the data are read out after each zap to see if any upsets (bit flips) have occurred. To eliminate differences between the IO circuits which might confound the data analysis, IO1 through IO4 were put into high-Z mode to make them similar to IO5 and IO6, which have no driver. A minimum of five zaps were applied to each IO at a given pre-charge voltage. The experiment was repeated on a second test chip. For these experiments, strobe was set to low. When the strobe is low, the latch is supposed to “hold” its logic value; this function is effected by the transmission gate shown in the schematic of Figure 13.

For gun precharge voltages up to ±8 kV, no upsets were observed when the stored data matched the input data. However, when the latch input data (D_{IN}; see Figure 13) and the stored

data (D_{HOLD}) were different, upsets occurred at precharge voltages as low as ± 0.5 kV, as indicated in Table 8 (system in mobile configuration) and Table 9 (system in tethered configuration). In these experiments, if even one static latch was upset, all were. Further, storing a logic high versus logic low had no impact on the likelihood of an upset taking place. The zap polarity does not strongly affect the upset percentages for zaps to the two IOs nearest the strobe line. However, for zaps to one of the two more distant IOs (3 and 6), the polarity strongly affects the likelihood of upset. Negative zaps to IO3 (DD protection) and IO6 (SCR protection) cause

Table 8: Zaps were applied to IOs 1-6 with the EUT in its mobile setup. The IO being zapped is in high-Z mode. The table lists the percentage of zaps that resulted in static latch upsets at each test condition.

	Zapped IO#					
V_{pre} (kV)	1 (DD)	2 (SCR)	3 (DD)	4 (SCR)	5 (DD)	6 (SCR)
-2	100%	100%	100%	100%	100%	100%
-1	0%	100%	100%	100%	100%	100%
-0.5	0%	0%	100%	100%	100%	100%
-0.25	0%	0%	0%	80%	80%	0%
0.25	0%	40%	20%	80%	100%	60%
0.5	0%	100%	0%	100%	100%	100%
1	0%	100%	100%	100%	100%	100%
2	100%	100%	100%	100%	100%	100%

Table 9: Same experiment as for Table 4 except that the EUT is in its tethered setup. Zaps were applied to each of the IOs.

	Zapped IO#				
V_{pre} (kV)	2 (SCR)	3 (DD)	4 (SCR)	5 (DD)	6 (SCR)
-2	100%	100%	100%	100%	100%
-1	100%	100%	100%	100%	100%
-0.5	40%	100%	100%	100%	100%
-0.25	0%	0%	100%	40%	0%
0.25	0%	20%	100%	60%	20%
0.5	40%	40%	80%	100%	100%
1	100%	100%	100%	100%	100%
2	0%	100%	100%	100%	100%

roughly equal numbers of upsets. However, positive zaps to IO6 caused more upsets to occur than positive zaps to IO3. This may indicate that for positive zaps, the top diode reduces the stress, and the noise coupling, more than the SCR does.

Based on the above, it is concluded that ESD is inducing noise on Strobe and this results in the inadvertent clocking-in of new data. Further support for this assertion may be found by considering the positional dependence of the upset threshold. The upset threshold is lowest when the zaps are applied to IO4 and IO5; at the IC, the Strobe pin is right between the pins for IO4 and IO5 (see Figure 12). Despite the robust on-board filtering of the Strobe signal line, logic level glitches appear to be causing upsets within the static latches. To further mitigate board-level noise coupling, the strobe line was shorted to ground very close to the socket and some experimentation was repeated. This short had no discernable effect on the occurrence of upsets, and thus it is assumed that glitches are induced on Strobe due to noise coupling inside the package.

Additionally, tests were performed to discern whether the state of the output driver at the zap pin affects the upset occurrences. IO3 and IO4 were configured as transmitters and zapped in both output high ($D_{OUT} = 1$) and output low ($D_{OUT} = 0$) states. Zaps were applied both with the output high and low. Table 10 shows the static latch upset percentage for each case, obtained at a precharge level of ±0.5 kV and ±1 kV, and with $D_{IN} \neq D_{HOLD}$. As in the previous experiment, zaps to IO4 cause upsets to occur more often than do zaps to IO3. Minor variations in the upset percentage are attributed to a different test chip being used for each dataset. It was found that the state of the output driver did not impact the occurrence of upsets in the static latches. To induce latch upset in the case that $D_{IN} = D_{HOLD}$, even higher precharge voltages had to be used. At +8 kV, two of the five chips tested had repeatable upsets. Upsets did not occur at lower precharge voltages, or with negative polarity.

Table 10: Static latch upsets from ESD gun zaps to IO3 and IO4. The output transmitters were enabled and the output level was varied. The latches' data inputs (D_{IN}) were set opposite to the stored data (D_{HOLD}).

	Zapped IO#					
	3			4		
V_{pre} (kV)	Hi-Z	Output Low	Output High	Hi-Z	Output Low	Output High
-1	100%	100%	100%	100%	100%	100%
-0.5	100%	100%	100%	100%	100%	100%
0.5	0%	0%	0%	100%	100%	100%
1	100%	100%	100%	100%	100%	100%

For completeness, the static latches were monitored for upsets following zaps to the USB shield (board ground-plane). These tests were performed with D_{IN} opposite D_{HOLD}. Only one upset was observed at 12 kV. Furthermore, this upset only changed the value of data stored in latches at the following addresses: 16, 20, 23-31, 32, 36, 39-47. These latches all contain the smallest feedback ("fb") inverter, as indicated in Table 3. When a glitch occurs on the Strobe line, the transmission gate is transparent for a short time allowing the D_{IN} signal to pass to the flip-flop. If this glitch is especially brief, the input to the flip-flops within the latches may not be sustained long enough for all of the latches to upset. It is not surprising that the set of latches with the smallest feedback inverter is most susceptible to this change since a smaller feedback inverter is less capable of sustaining the data when a transient is presented at the input of the flip-flop.

4.3.3 Upsets in Dynamic Logic

The dynamic OR gates are in the pre-charge state when Strobe = 0 and evaluate when Strobe = 1. During ESD testing, the dynamic node is first precharged (the gates' outputs will be low) and subsequently all of the inputs to the dynamic stage (refer to Figure 13(a)) are held low and the gate is placed in evaluate mode. This will place all of the dynamic nodes into their high-

impedance states. If the OR gate's output is observed to have changed state from low to high following ESD, it can be assumed that the charge stored on the dynamic node was erased. In this experiment, zaps were applied to IO3 and IO4. These IOs were chosen due to their proximity to the bank of dynamic logic gates within the core of the test chip.

For zaps to the mobile system, only gates that had a "driven input," i.e., an input signal that ultimately originated at pin DIN, were observed to upset (addresses 6, 7, 14 and 15). The gates that have only tied-low inputs do not upset. Thus when upsets are indicated, only the gates with driven inputs are considered. Once again, a minimum of five zaps were applied for each test case and the experiment was repeated on two test chips. Table 11 provides a summary of the upsets observed when the EUT was in its mobile configuration. If any single gate was observed to upset, all of the gates did so. Only positive polarity discharges caused upsets. These upsets occurred in a window between 1 kV and 6 kV. Only zaps to IO3 caused upsets; no upsets were observed following zaps to IO4. Further, the IOs' setting—high-Z, output-low, or output-high—was found to have little influence on the results.

As established earlier, for negative zaps, the leading edge will likely cause a –(0→1) glitch, while the trailing edge causes a -(1→0) glitch. This means that the leading edge will cause data erasure by pulling DIN high, but the trailing edge will cause Strobe to go low resulting in a recharging of the dynamic node. Thus, for this experiment negative zaps do not appear to cause upsets. However, upon reaching some threshold, the leading edge of positive zap is likely to cause a +(1→0) glitch while the trailing edge is likely to cause a +(0→1) glitch. The initial +(1→0) on Strobe will cause the dynamic node to precharge (note it was already charged), the subsequent +(0→1) on DIN will cause the dynamic node to discharge. This will result in a detectable upset. Further, noise has been observed to have an oscillatory behavior due to the

inductance of various bondwires. Despite seeing upsets for low-level positive zaps (+1 kV to +6 kV), further increasing the precharge voltage can cause subsequent oscillations to have a large enough magnitude to cause a secondary glitch on the Strobe line. If a (1→0) glitch subsequently appears on Strobe, the dynamic node to will be pulled high and the earlier upset would be lost. Since zaps to IO4 have already been established to cause more noise on Strobe, it can be assumed that IO4 zaps always causes the dynamic OR gates to enter the pre-charge state and return to the output low state.

It should be noted that while the Strobe pin is close to the zapped pins (IO3 and IO4), the DIN pin is significantly farther away. This indicates that while the glitches on DIN might be the result of noise coupling to its pin, other explanations, e.g. supply noise, must also be entertained.

The experiment used to generate the data for Table 11 was repeated with the EUT in its tethered mode. Tethering was known to lower the test chip's hard failure threshold [13] so, for this reason, zaps to IO3 were limited to ±6 kV and zaps to IO4 were limited to ±1 kV. Table 12 provides a summary of upsets that occurred using a tethered system. For this dataset the criteria for upset is defined as upsets to all driven gates and no others. Upsets did occur for non-driven gates or for a subset of driven gates. This indicates that a different mechanism, substrate current injection, was responsible [17]. This topic was investigated by another researcher and is outside of the scope of this document. Considering the data in Table 12, fewer upsets occur in the tethered configuration than in the mobile configuration, however, the upsets occur at lower pre-charge voltage. This might suggest that discharges to the tethered EUT cause as many or more upsets at the dynamic node than for the mobile EUT, but that the data upsets are being masked by the windowing effect. At 1 kV, upsets are observed for zaps to both IO3 and IO4. The current waveforms for discharges into tethered and mobile systems are different and thus the noise

spectra will differ; this may be why there are subtle differences between the results for the mobile and tethered systems.

Table 11: Fraction of experiments that cause data upset in the dynamic OR gates when IO3 is zapped at the level indicated. No upsets occurred for zaps to IO 4. The OR gates' outputs are low prior to ESD and the EUT is in its mobile configuration. Two test chips (TC) were used for this experiment.

Upsets to Dynamic Logic Gates When IO3 Is Zapped			
V_{PRE} (kV)	# of Upsets/# of Zaps		Combined %
	TC 1	TC 2	
-8 to -0.5	0/5	0/5	0%
0.5	0/5	0/5	0%
1	0/5	1/5	10%
2	5/5	5/5	100%
3	5/5	5/5	100%
4	5/5	0/5	50%
5	5/5	0/5	50%
6	0/5	1/20	4%
7	0/5	0/5	0%
8	0/5	0/5	0%
Upsets to Dynamic Logic Gates When IO4 Is Zapped			
V_{PRE} (kV)	# of Upsets/# of Zaps		Combined %
	TC 1	TC 2	
-8 to 8	0/5	0/5	0%

Table 12: Same experiment as for Table 11 except that the EUT is in its tethered configuration.

Upsets to Dynamic Logic Gates When IO3 Is Zapped			
V_{PRE} (kV)	# of Upsets/# of Zaps		Combined %
	TC 1	TC 2	
-6 to -0.5	0/5	0/5	0%
0.5	0/5	0/5	0%
1	5/5	5/5	100%
2 to 6	0/5 each	0/5 each	0%
Upsets to Dynamic Logic Gates When IO4 Is Zapped			
V_{PRE} (kV)	# of Upsets/# of Zaps		Combined %
	TC 1	TC 2	
-1	0/5	0/5	0%
-0.5	0/5	0/5	0%
0.5	0/5	0/5	0%
1	0/5	5/5	50%

4.3.4 Bit Flips in the Shift-Register

The SSTL circuitry was originally included for an unrelated study. The six functional SSTL transceivers in bank B are identical to those in bank A. All six transceivers on one side are

programmed to the same operating state. The code is stored in a 10-bit shift-register; one register bit rx_a [rx_b] is the RX enable for bank A [B], and four register bits ZQ_a [ZQ_b] are used to select the termination resistance for bank A [B] receivers. DC current is consumed only when a transceiver is in RX mode, because termination resistors are connected from the SSTL IO pin to V_{DD} and to V_{SS}. The DC current consumption is measured by setting bank A [B] in RX mode, bank B [A] in high-Z mode, and sweeping the ZQ_a [ZQ_b] values. The results are tabulated in Table 13. The current consumption is an increasing function of ZQ_a [ZQ_b], and can be approximated as a linear function as shown in Equation (1):

$$I_{DDQ} = rx_a(10 + 8.4 \cdot ZQ_a) + rx_b(10 + 8.4 \cdot ZQ_b) \quad (1)$$

Table 13: Measured DC current with one half (one bank) of the SSTL IO set to RX mode. The 4-bit termination resistance setting (ZQ_a/ZQ_b) is varied. Each side contains six working SSTL IOs, thus the I_{DDQ} values are the sum current consumption of six SSTL IOs on one side.

ZQ_a/ZQ_b	0	1	2	3	4	5	6	7
I_{DDQ} (mA)	10	22	30	38	47	55	63	70
ZQ_a/ZQ_b	8	9	10	11	12	13	14	15
I_{DDQ} (mA)	78	87	95	103	110	117	125	133

During system-level ESD testing, all the transceivers were programmed to the high-Z mode to minimize I_{DDQ}; this is achieved by setting all 10 bits of the shift-register to 0. After the ESD gun discharges to the system, the on-board I_{DDQ} monitor reported an increased current draw. The phenomenon was studied carefully for the case of ESD gun discharges to IOs 1-6 at a precharge voltage of ±2kV; the post-stress I_{DDQ} was different from the pre-stress value in more than 95% of cases and there was no discernible dependence on the zap pin, all of which were far from the SSTL block. By reprogramming the SSTL block to its original high-Z state, the current draw could be reduced back to its minimum level, showing that the chip had not been physically damaged, merely reprogrammed. Soft failures (i.e., SSTL state changes) were occasionally

observed at precharge voltage magnitudes less than 2 kV, but 2 kV is the level at which these became very frequent.

The test results are aggregated in Table 14. After each zap, it is possible to determine whether the transceivers are in RX, TX or High-Z mode simply by measuring the DC voltage at the IO pad. This information is given in the table. Using the data given in columns 2, 3 and 4 of the table, Equation (1) is used to find the post-ESD values of ZQ_a and ZQ_b; recall these were set to 0000 before the stress. After stress, certain shift-register settings occur more frequently than others; this is indicated in the rightmost column of the table, which indicates the percentage of soft failures that correspond to a given state. This suggests that the shift-register gets disturbed in a manner that is not entirely random and likely has some dependence on the physical design and setup of the test system.

Table 14: SSTL IO upsets. 96 zaps were performed at +2kV and 24 at -2kV. I_{DDQ} was changed after >95% of the zaps. In the table, soft failures are categorized according to post-stress I_{DDQ} and IO modes. The rightmost column indicates how often the given soft failure is observed.

V_{PRE} (kV)	I_{DDQ} (mA)	Bank A	Bank B	ZQ_a	ZQ_b	%
2	125	Hi-Z	RX	-	1110	21.5 %
2	133	Hi-Z	RX	-	1111	20.3 %
2	110	Hi-Z	RX	-	1100	16.5 %
2	78	Hi-Z	RX	-	1000	12.7 %
2	117	Hi-Z	RX	-	1101	5.1 %
2	103	Hi-Z	RX	-	1011	5.1 %
2	206	RX	RX	1011	1011	3.8 %
2	251	RX	RX	1110	1110	2.5 %
2	188	RX	RX	1010	1010	2.5 %
2	79	RX	RX	0010	0010	2.5 %
2	171	RX	RX	1001	1001	1.3 %
2	138	RX	RX	0111	0111	1.3 %
2	107	RX	TX	1100	-	1.3 %
2	88	Hi-Z	RX	-	1001	1.3 %
2	65	Hi-Z	RX	-	0110	1.3 %
2	10	RX	TX	0000	-	1.3 %
-2	133	Hi-Z	RX	-	1111	41.7 %
-2	125	Hi-Z	RX	-	1110	25.0 %
-2	110	Hi-Z	RX	-	1100	12.5 %
-2	227	RX	RX	1011	1110	4.2 %
-2	190	RX	RX	1010	1010	4.2 %
-2	129	RX	RX	0110	0110	4.2 %
-2	78	Hi-Z	RX	-	1000	4.2 %
-2	10	RX	TX	0000	-	4.2 %

In an attempt to ascertain the root cause of these upsets, the logic levels of the SCK, SDI and LE control signals were varied. SCK and SDI were individually left floating or tied low on the board. LE could be left floating, hard-tied low, or connected to a push-button and resistive pull-down. When pressed, this push-button connects LE to VDD. Note that data are written into the shift-register's output stage latches when LE is toggled high then low. It is this output stage that connects to the SSTL transceivers which, in turn, modulate the chip's power consumption. Thus, I_{DDQ} is expected to change **only** if a glitch appears on LE. Thus, when the configuration of the LE line is varied, the final state of the shift-register should be unaffected. Instead, the configuration of the LE line is expected to affect the occurrence of observed upsets. Further, it is possible that the output latches themselves become corrupted in addition to or instead of the shift-register contents. Such a distinction would be hard to observe. If the post-zap I_{DDQ} is dependent on the configuration of the SDI and SCK lines, then it is likely that the shift-register is being upset and subsequently passed into the latches. If there is no such dependence then it is possible that the contents of the output latches are being upset independent of the shift-register.

In the previous experiments, the SSTL shift-register was populated with all zeros prior to the zap. Since reprogramming involves overwriting bits within the shift-register, it is possible that zeros are being written to the shift-register in some cases. By resetting all bits to zeros, any zap resulting in all zeros being written into the shift-register would not be detected as an upset. For the following experiments, a different known code was written into the shift-register prior to each zap. This code was chosen to be 1010100000. This value was chosen to test several things. First, it is desirable to know if data are being clocked into the shift-register or if this register is being upset by some other means. By using the new code, if any ones are observed in the last five bits, or if five zeros appear to have shifted, than the number of bits clocked in can simply be counted.

Second, this code results in an I_{DDQ} of around 45 mA, which is not seen in Table 14. Thus, any upset is likely to be detected.

Positive 2 kV zaps were applied to IO3, which is located far from the SCK, SDI and LE pins (refer to Figure 12 and Figure 15). Many zaps were applied to the EUT for all possible combinations of the shift-register control signals. I_{DDQ} was measured before and after each zap. There did not appear to be any systematic change in the I_{DDQ} due to the new starting code, except that all-zero upsets could now be detected. This suggests that enough new bits are clocked in to completely overwrite the shift-register contents, or a different mechanism is at play which can overwrite bits without clocking in data.

In the histogram of Figure 18, the results have been grouped to highlight the effect of the SCK and SDI control signals. It is important to note that the shift-register contents were upset after **every** zap. This is in contrast to the previous data set of Table 14 where some zaps did not result in upset. This is attributed to the ability to detect all-zero upsets which was impossible before. Also, since upsets occurred after every zap, the configuration of the LE line had no discernable effect. From the histogram, the following conclusions can be drawn. If SCK is tied low, the upset states are more uniformly distributed. However, if SCK is floating, the value of SDI appears to have a significant effect. As seen in Equation (1), I_{DDQ} generally increases as the code is stepped up from 0000000000 to 1111111111; in general, a higher I_{DDQ} indicates more 1's are stored in the register. If SDI is tied low, many more low value I_{DDQ} values result suggesting that mostly zeros are written into the shift-register. Conversely, if the SDI pin is floating, mostly ones are written into the shift-register.

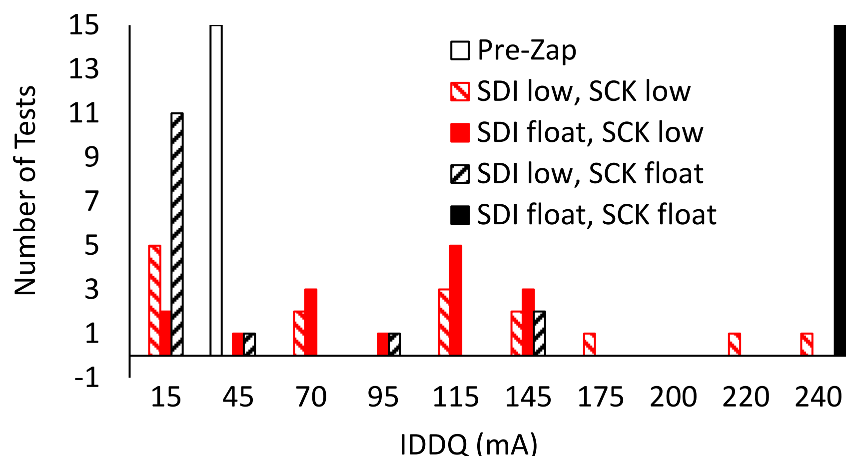

Figure 18: Histogram of the post-zap current draw. ESD zaps at +2 kV are applied to IO3 and the shift-register control signals are varied. The shift-register was initially set to a known code and the corresponding value of IDDQ is marked in the figures.

Based on these observations, the upset mechanisms can be surmised. The ESD current into IO3 returns to the board ground via the IC's VDDIO and VSSIO bond-wires. The ESD current has a large time-derivative which causes bounce of the on-chip power and ground nets. These nets are decoupled from the corresponding board-level nets due to the bond-wire inductance. The signals into pins SDI, SCK and LE are referenced to the board ground, rather than the chip ground. Thus, oscillatory noise is induced on the chip supply that can cause apparent glitches at the input pins. The noise seen at SCK and SDI is in phase thus when SCK goes high, SDI is also high and a one is written into the shift-register. Tying the SDI signal level low increases its glitch threshold, thus zeros are written into the shift-register when SCK goes high. When SCK is tied low, it has an improved immunity to glitches and the uniform distribution of post-ESD register codes may suggest that individual stages of the register are randomly upsetting due to supply noise.

Tethered-system experiments were planned, but due to a strange malfunction were impossible to conduct. While tethered, if any of the control signals (SDI, SCK, or LE) were tied to board ground, the I_{DDQ} of the board would vary between the known values associated with stored

codes in the shift register. This prevented a similar experiment to the above from being conducted. However, the experiment could be performed with all the control signals floating. Limited testing using the 1010100000 code showed that, in the tethered configuration, the shift-register upset every time to all zeros. The cause of the odd behavior has not yet been discerned.

4.3.5 Upsets to USB Functionality

The next experiment was to evaluate the USB signal integrity during system-level ESD. A USB transmitter only operates when the USB port is connected to another device. The USB ESD test setup was designed to emulate a realistic, worst-case scenario. The USB signal integrity was monitored by connecting the USB port to an oscilloscope via a custom cable that converts a USB cable into two SMA cables (one for each differential signal). Typically, an oscilloscope has a metal chassis that is shorted to earth-ground through the power cable. In this experiment, the oscilloscope's chassis was isolated from the earth-ground. This removes the low-frequency return path to the gun which would not exist if the USB TX was communicating with another mobile device. Gun zaps were applied to the USB shield since this may be the only point of an enclosed system that gets exposed to contact discharge. Figure 19 shows an example of the USB output signal and the current from the gun (measured using an F-65A current probe around the gun tip) during a -500V zap. The USB transmitter is seen to recover in a timely manner after the discharge event has passed. The differential signal resumes normal behavior immediately after the ESD event (around 225 ns). However, fluctuating dc offsets (> 0.5V) on the individual data-lines persist beyond this time, keeping the transmitter out of compliance with the USB specification. These offsets arise from residual charge that fluctuates between the oscilloscope (isolated from ground) and the board. For discharges up to ±5kV, the output waveform frequency and amplitude finally do recover after the zap, indicating that the USB transmitter was not damaged by the ESD.

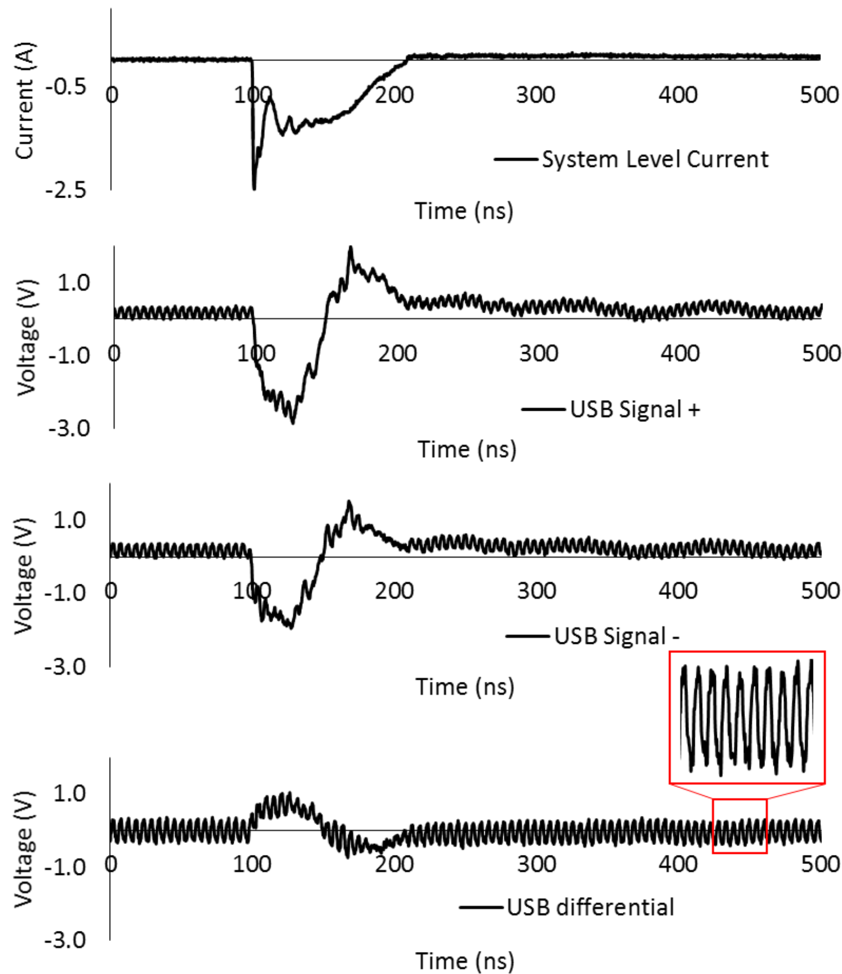

Figure 19: USB waveforms before, during, and after gun discharge.

4.3.5 Hard Failures

Several chips were taken to hard failure. Table 15 summarizes the observed failures. The IOs listed in the table were zapped with positive polarity discharges of increasing amplitude. To check for hard failure after each zap, the IO's I-V characteristic was measured in the high-Z, output high, and output low states (where applicable).

For test1, IO5 was zapped with the system in the floating (battery operated) configuration and for Test 2, IO5 was zapped with the system in the tethered (dc power supply) configuration. The tethered setup produces hard failure at a much lower stress level—8 kV rather than 21 kV. It

is clear that having a resistive ground return, as in the tethered system, dramatically increases the stress delivered to the system. In fact, for the mobile setup, hard failure only occurred once a spark formed, creating a resistive connection (the ionized air) between the board and the tabletop. The low-resistance spark causes there to be a greater current flow from the gun to the board; although the tabletop is not connected to the gun's ground, it acts as a large sink on the ESD time-scale. For the remaining tests, the tethered system configuration was utilized, as this is the worst case with respect to hard failures. A comparison of the results of tests 1, 3 and 4 suggests that the presence of an output driver does not significantly affect the hard failure level, and having the driver activated (as in test 4) at most only slightly improves the robustness. However, with the output driver in output-low state, some of the current from the gun zap may be shunted to V_{SSIO} through the driver NMOS, as suggested by the observed increased leakage to V_{SSIO}. A comparison of the results of tests 1 and 5 shows that the SCR is much less robust than the dual diodes and it fails at a much lower precharge voltage. This result was expected; the SCR's failure current is about half that of the top diode's during 100 ns TLP testing of stand-alone test structures.

Table 15: Hard failures observed during gun zapping. Electrical failure analysis is presented in the shaded boxes. Electrical failure analysis is performed by measuring the IO characteristics between the IO pin and V_{SSIO} and between the IO and V_{DDIO}.

Test	IO	Output Driver	Main Protection Device	Board Grounding	Failure Level (kV)
1	5	N.A.	Top Diode	Tethered	8
		Failure: Increased leakage from V_{DDIO} to IO			
2	5	N.A.	Top Diode	Mobile	21
		Spark formed during zap between board and tabletop. Failure: Short formed between IO and V_{SSIO}			
3	1	Hi-Z	Top Diode	Tethered	8
		Failure: Increased leakage from V_{DDIO} to IO and to V_{SSIO}			
4	1	OL	Top Diode	Tethered	7
		Failure: Increased leakage from IO to V_{SSIO}			
5	6	N.A.	SCR	Tethered	3
		Failure: Increased leakage from IO to V_{SSIO}			

CHAPTER 5:

ADVANCED NOISE MONITORS

The work presented so far has helped to identify several likely mechanisms for soft failures. However, these soft failures represent those found in a very specific chip architecture. It is also important to develop experiments that can confirm or refute the hypothesized mechanisms previously investigated. To achieve these goals and potentially yield ways of mitigating soft failures, another custom test chip has been designed and fabricated.

The primary goal of this new work is to more fully understand the correlation between power supply noise and the occurrence of soft failures. Novel supply voltage monitors have been developed to meet these needs. These monitors are intended to capture a quantitative measurement of the voltage deviation, above and below, the nominal supply-voltage for both the core power-domains and high-voltage domains. Sensitive logic elements will again be included within the core so that soft failures can be directly linked to supply noise levels.

On the previous test chip, the injected current from zaps applied to the chip were ultimately directed immediately into the VDDIO/VSSIO domain. This is because all of the IOs were placed in this domain and all of the ESD protection devices shunted current into this domain. The new test chip will attempt to improve upon the understanding of noise propagation and noise coupling between power domains. Dedicated zap IOs will be placed in the low-voltage domain as well as in the high-voltage domain. This study will allow for a comparison of the effects of noise that is directly injected into the core domain to the effects of noise that couples from the high-voltage domain into the core. Further, an internally generated supply is included on this test chip to determine if these types of power domains offer any robustness against soft failures.

5.1 New Test Chip Overview

The new test chip was designed to expand upon the experiments of the previous test chip. The test-chip layout is shown in Figure 20. There are three main power domains that power the monitors included on this test chip. VDD33 is the 3.3 V supply that powers the majority of the supply ring, eVDD is an externally supplied low-voltage domain with variable input (from 1.2 to 1.5 V) and iVDD1 is an internally generated supply that is powered from VDD33 and uses eVDD as the reference voltage. All 3.3 V pad cells were duplicated from the previous test chip, where applicable, and should have the same failure level. Minor changes were made to the rail clamps'

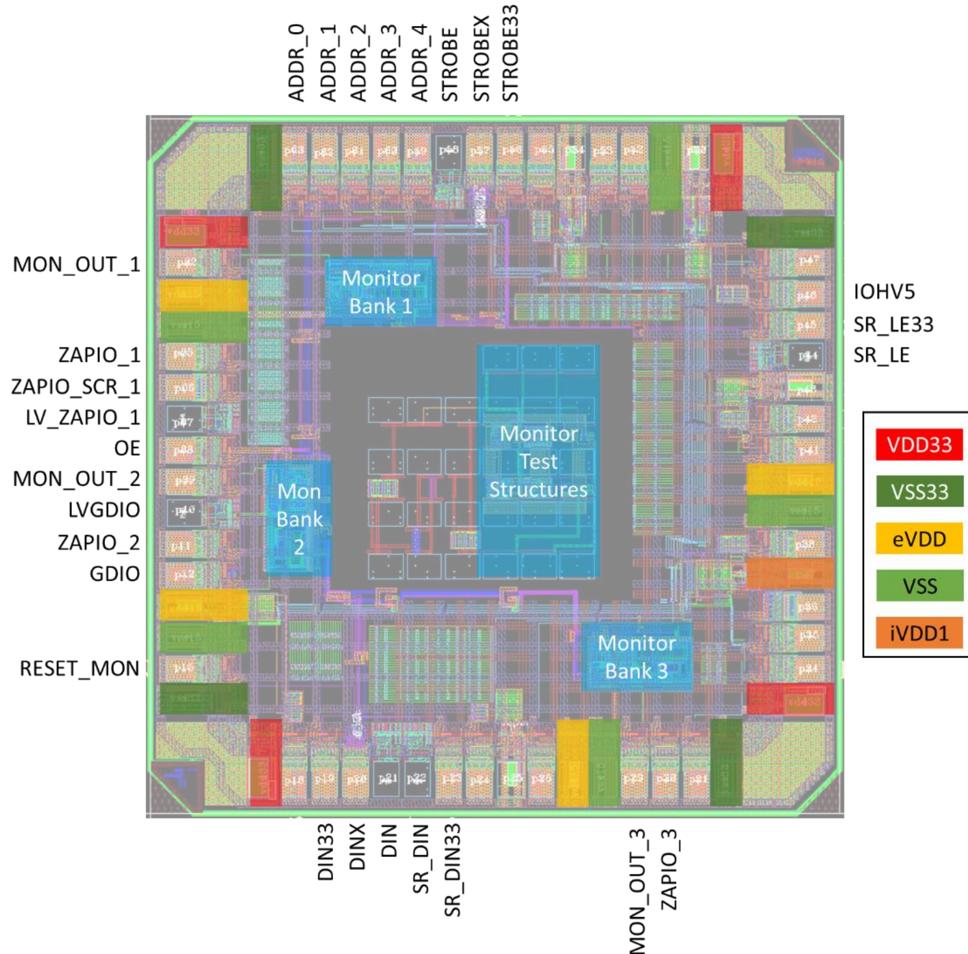

Figure 20: New test chip layout. Important pads and monitor blocks have been indicated.

trigger circuit to reduce leakage current, but this change should not affect the failure level. Low-voltage control signal receivers consist of a CMOS buffer and DD protection similar to the previous test system. There are several dedicated zap locations distributed near the monitor banks which are located in the core of the chip. ZAPIO_1, 2 and 3 as well as IOHV5 are 3.3 V cells with DD protection. ZAPIO_SCR_1 is a 3.3 V cell with SCR protection. These cells each contain a transmitter and receiver which are not connected to the pad. LV_ZAPIO_1 is a low voltage cell with DD protection and no transmitter or receiver.

The various control signals required for this work are indicated; these are not intended to be directly stressed, but all include the same DD protection as the dedicated zap cells. MON_OUT_1, 2 and 3 are the outputs of three distinct 32-bit multiplexers. One multiplexer is located in each of the monitor banks. The address pins ADDR_0 through ADDR_4 supply the address to the multiplexer. Two glitch detectors are included within the core of the test chip. One monitors the GDIO cell which is a 3.3 V TRX cell similar to IO1 from the previous test chip. The OE signal line determines whether the GDIO cell is receiving an off-chip signal or transmitting the glitch-detector output. The other glitch detector monitors the LVGDIO cell which is a low-voltage receiver only. The output from this glitch detector is read-out through the multiplexor in monitor bank 2. Further, two out-of-range error-detectors (ORED) [18] are included within the DINX, and STROBEX pad cells. The ORED circuit are used to detect when a signal line exceeds the supply voltage (VDD33) or falls below VSS.

The three monitor banks contain static latches, dynamic or-gates and shift-registers similar to those included in the previous test chip as well as novel supply-voltage monitors. Each of the following sections will focus on the types of monitors and sensitive logic elements included on the test chip. The three banks are each associated with a different supply. The supply monitors in

each bank will monitor the supply they are associated with. Monitor bank 1 is associated with the VDD33 domain (3.3 V). The core logic is supplied from the 1.5 V externally supplied eVDD domain, but the inputs will be supplied from standard 3.3 V receiver cells powered from the VDD33 domain. Monitor bank 2 is associated with the externally supplied low-voltage domain, eVDD. Inputs to logic within this bank will be supplied by low-voltage receiver cells that are powered from the eVDD domain. Monitor bank 3 is associated with the internally generated low-voltage domain, iVDD1. The core logic in this bank is powered by iVDD1. Inputs to logic elements in this bank come from the same IO cells that are used in monitor bank 2. By creating three monitor banks in this way, noise within the three power domains can be compared against each other. Also the effect of the IO cell supply voltage can be discerned. It is expected that the input circuits in the VDD33 domain will have a larger noise margin and thus experience fewer glitches.

5.1.1 Supply-Voltage Monitors

A prominent mechanism suspected of causing soft failures is supply noise. However, it is virtually impossible to monitor the on-chip supply voltages with external equipment, without introducing invasive measurement techniques. To meet this need, internal monitors were developed. These monitors are intended to sample the maximum and minimum potential differences between a supply net and its reference during ESD, and provide a readout of those values following the ESD event. This is a difficult task because, during ESD, it cannot be assumed that the power supply is adequate to power active devices within the core. Further, since various supplies are used, monitors have to be tailored to each supply domain. There are six supply monitors included on this test chip. The monitors fall into two categories: under-voltage and over-voltage. Under-voltage monitors are intended to sample and hold voltages that are lower

than the nominal supply voltage. Over-voltage monitors are intended to sample and hold voltages that are greater than the supply voltage. In both cases a capacitor is used to store an analog voltage during an ESD event and upon power-supply recovery, a flash analog-to-digital converter (ADC) is used to store that value digitally until readout can occur at a later time. There is a 3.3 V version and a low-voltage version of each monitor.

The 3.3 V under-voltage (HVUV) monitor is shown in Figure 21. This monitor front-end is a sampling circuit. All of these monitors utilize a simple diode-capacitor chain to sample the voltage on the supply and store the peak value for future readout. In this case the peak is a negative peak with a DC offset equal to the nominal supply voltage. The capacitor's voltage starts at the quiescent voltage of VDD33 due to leakage through the diode. During ESD, if a negative peak on VDD33 is large enough to forward-bias the diode, the capacitor will begin to discharge through the diode and the capacitor voltage will drop. If VDD33 subsequently begins to rise, the diode will become reverse biased and the lower voltage will be maintained on the capacitor. In order to let the capacitor track the negative peak of the supply voltage closely, the threshold voltage (V_t) of the diode must be minimum. Also, the diode must be able to respond quickly. In order to achieve both metrics, a MOS-based diode was used. For this application a thick-gate PMOS was most suitable. Using a low-voltage threshold (LVT) device would provide for a smaller V_t, but could not withstand the nominal 3.3 V of the supply. Using a PMOS diode leads to beneficial body-effect. As VDD33 falls below the capacitor voltage, V_C, the PMOS source (V_C) is at a larger potential than the PMOS body (VDD33) reducing the threshold voltage of the PMOS. Further, the parasitic devices associated with the PMOS will help to discharge the capacitance. The PN junction formed between the PMOS source and body will become forward-biased with a sufficient drop in VDD33 and conduct current to VDD33. At the same time the PNP formed

between the PMOS source, PMOS body, and chip substrate will turn on and discharge the capacitor into the chip's substrate.

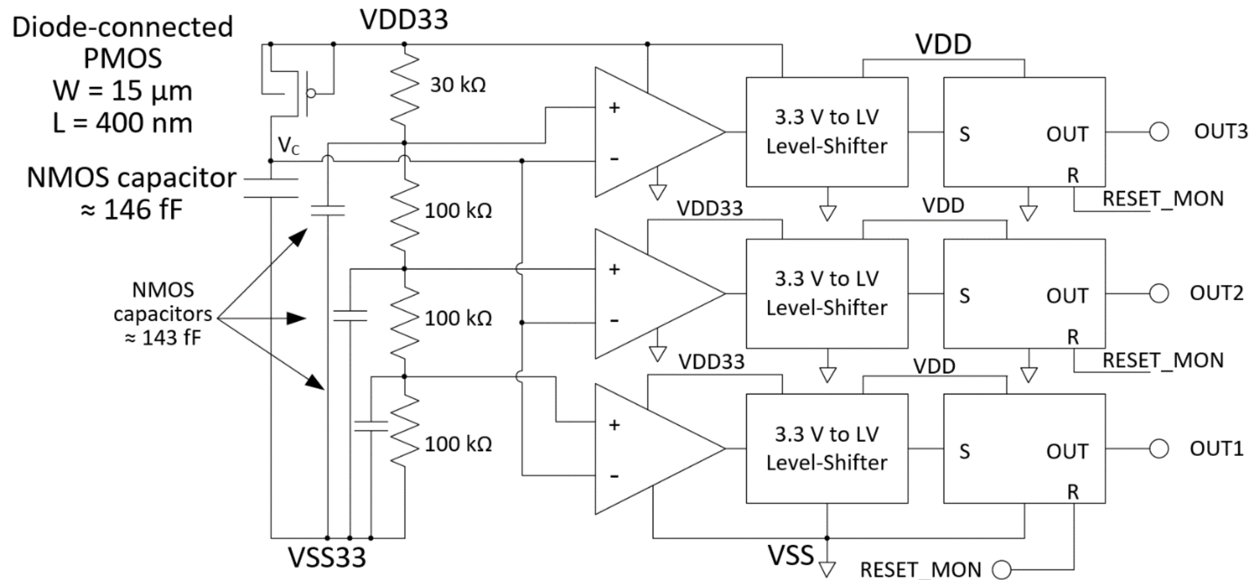

Figure 21: Under-voltage monitor for the VDD33 domain.

Due to carriers injected into the chip substrate, the capacitor design required special consideration. According to design kit of this process, the best capacitance for this application can be achieved from an NMOS layout placed in an N-well. This layout, referred to as NCAP, achieves higher capacitance than a standard NMOS since a positive voltage at the gate biases the device in accumulation. However, since this chip is intended for ESD experimentation, such a layout poses a potential risk for latch-up. The relatively large N-well area associated with the capacitor would need to be biased at the reference potential VSS33. This allows for the formation of a parasitic PNPN with large collector area. To remedy this, the capacitors were placed in an isolated N-well. The cross section of both scenarios is illustrated in Figure 22. By fully encapsulating the NCAPs inside a triple-well structure, the parasitic PNPs associated with the diode are isolated.

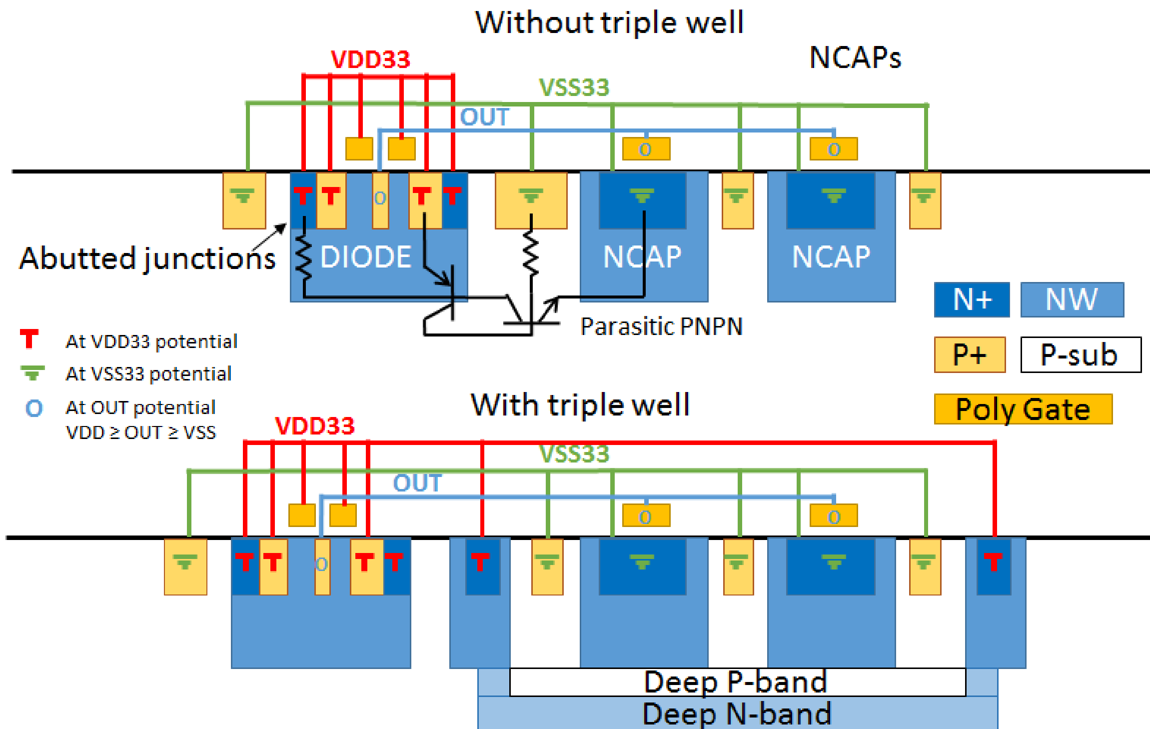

Figure 22: Cross section of PMOS diode and NCAP layout. The layout without triple-well isolation shows a parasitic PNPN with large NPN collector area. The layout with triple-well isolation removes this vulnerability.

Following the ESD event, once the supply has stabilized, the flash ADC will be functional and the voltage stored across the capacitor will be compared to the voltages formed in the resistor chain. Capacitors are included on the resistor chain to provide stable references during the comparison. The ADC provides three levels of resolution, nominally: 3 V, 2 V, and 1 V. If the voltage stored across the capacitance is less than one of these levels, that comparator will set a latch and hold for future readout. For under-voltage monitors, OUT1 indicates a larger under-voltage than OUT2 and OUT3. The three bits of the readout are logically combined into two bits with 00 representing no detected under-voltage and 11 representing the maximum detected under-votlage. These two bits are sent to the multiplexer within this monitor bank. All of the latches within the supply-voltage monitors are reset simultaneously using the RESET_MON signal which is generated off-chip. A simulation of the HVUV monitor is shown in Figure 23. The logical

combination of the three output bits into two has been omitted. The ESD current is produced from an ESD gun model. The simulation is setup to represent injection into an IO pin with DD protection and a rail clamp between VDD33 and VSS33. Decoupling capacitance is included both on and off-chip, and inductors are used to model the power, ground and IO bondwires. Because of the ESD protection, the injected current will first travel to VDD33 through the top-diode and then to VSS33 through the rail clamp and/or decoupling capacitance. The VSS and VSS33 domains are connected with APD. This connection allows noise on VSS33 to couple to eVSS. Several nanoseconds after VDD33 drops low enough, OUT3 goes high. The ripples on OUT3 are the result of the eVDD - eVSS noise mentioned previously. Despite the second reference level on the flash ADC being set to 2 V, the drop of VDD33 cannot trigger OUT2. This is because VDD33 must fall 2 V below the voltage drop across the diode-connected PMOS in order for the capacitor voltage to reach the 2 V level. Also, the voltage must be maintained at the low level long enough for the capacitor to discharge through the on-resistance of the diode.

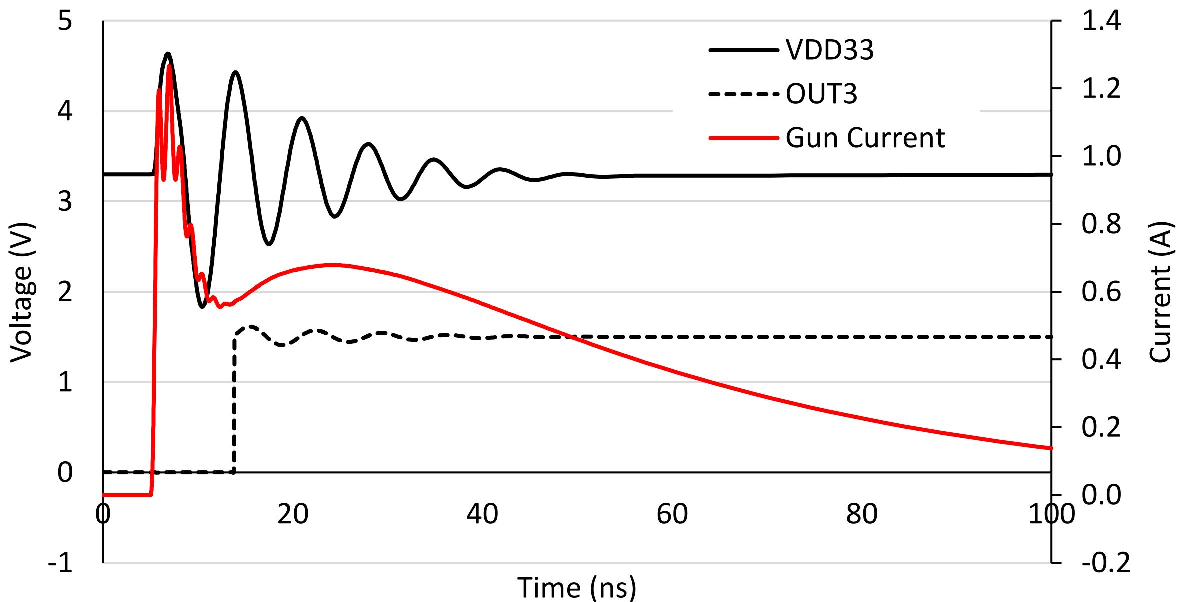

Figure 23: Simulation of the 3.3 V domain under-voltage monitor. The ESD current is injected into an IO pin with DD protection and a rail clamp. OUT3 goes high several nanoseconds after VDD33 begins to fall.

The under-voltage monitor for the low-voltage domain (LVUV) is shown in Figure 24. This design is largely similar to the previous design, but has been optimized for the low-voltage domain (eVDD shown). A LVT PMOS was used as the diode to reduce V_t and provide better supply voltage tracking. To reduce the design burden, the same flash ADC was used and placed in the 3.3 V domain. This time, however, the resistor chain is referenced to the low-voltage supply domain that is being monitored. The readout circuitry for this monitor is identical to the previous design. One of these monitors is placed in monitor bank 2 and monitors the eVDD supply and another is placed in monitor bank 3 to monitor the internally generated supply, iVDD1. A simulation of the LVUV monitor is shown in Figure 25. The simulation setup is similar to that of Figure 23. Since the ESD current is injected primarily into the 3.3 V domain, the current must be much larger to cause a large voltage fluctuation on the eVDD domain.

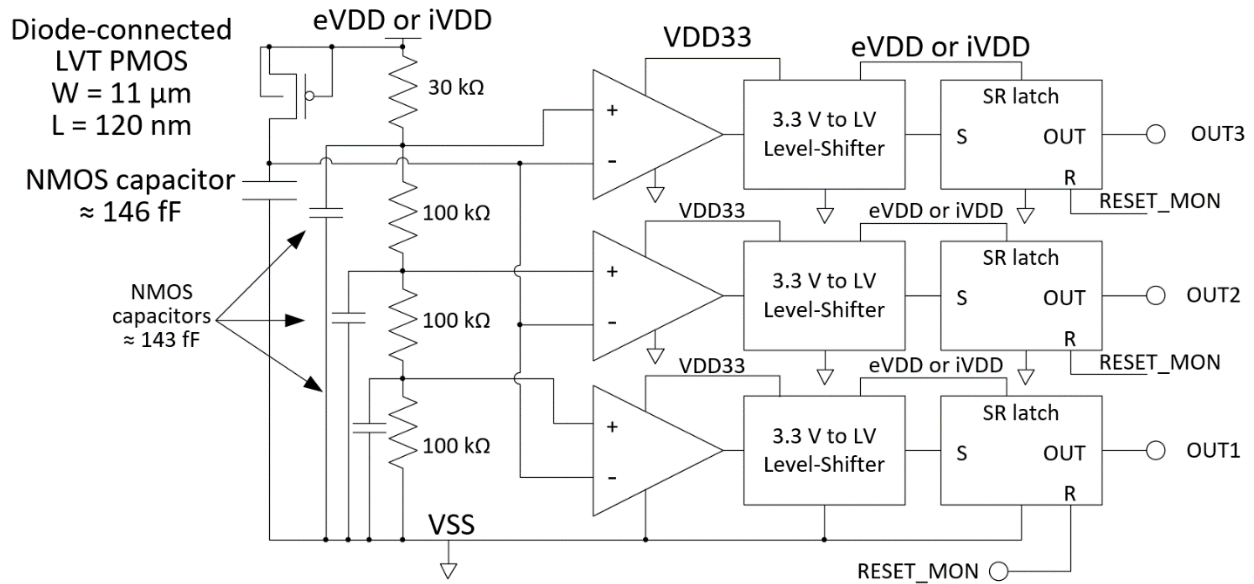

Figure 24: Under-voltage monitor for the externally and internally generated low-voltage domains.

The over-voltage monitor for the 3.3 V domain (HVOV) is shown in Figure 26. This topology is again similar to the under-voltage monitor, but the diode's polarity has been flipped. In this case the quiescent voltage across the capacitor is still 3.3 V, but if VDD33 rises high

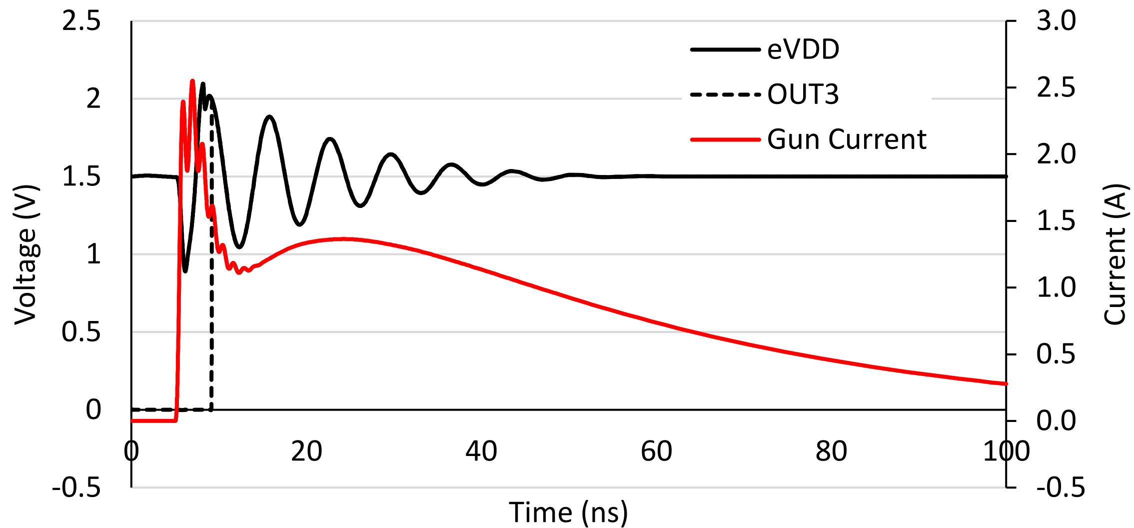

Figure 25: Simulation of the low-voltage domain under-voltage monitor.

enough during an ESD event, the diode will become forward biased and charge the capacitor to a higher voltage. A flash ADC that references VDD33 would not be able to provide a digital readout for the capacitance voltage which is greater than VDD33. Thus the voltage at the capacitor is first resistively divided. While this provides a shunt path to discharge the capacitor, the values of the resistive divider were chosen to be large enough to limit this effect and preserve the retention time. The resistor values for both this resistive divider and the flash ADC resistor chain were chosen through simulation such that OUT1 will be set if the capacitor voltage raises 0.5 V above nominal (3.8 V) and OUT2 will be set if the capacitor voltage raises 1 V above nominal (4.3 V). Also, the resistor chain has an NCAP placed at each node. This is done to prevent fluctuations on VDD33 from lowering these reference voltages during negative VDD33 swings. Even if the capacitor voltage is smaller than the designed thresholds for setting the latches, by reducing the reference voltages the comparator could cause the latches to set erroneously. A simulation of the HVOV monitor is shown in Figure 27. The VDD33 potential must reach 3.8 V plus the voltage drop across the diode-connected NMOS before OUT1 can be set high.

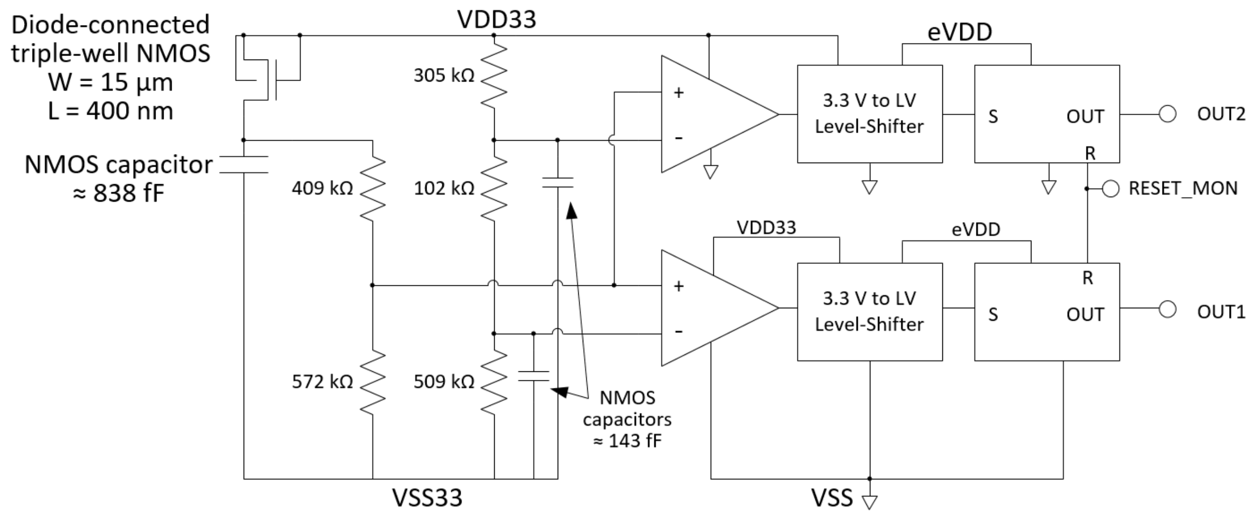

Figure 26: Over-voltage monitor for the VDD33 domain.

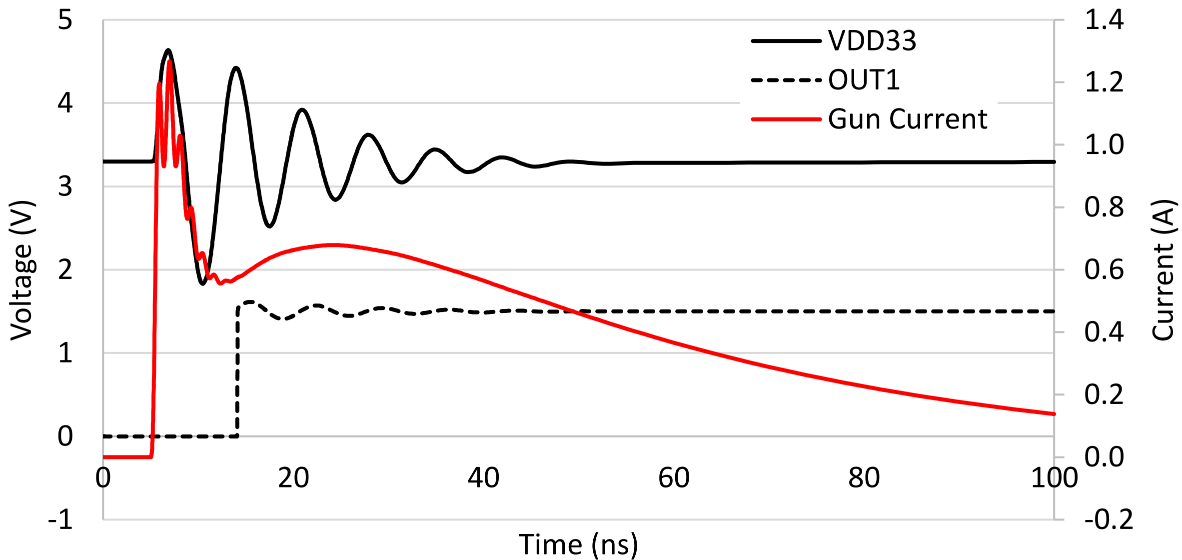

Figure 27: Simulation of the 3.3 V domain over-voltage monitor.

To accommodate the new diode polarity, a different diode layout must be used. Three scenarios are illustrated in the cross sections of Figure 28. Following and ESD stress, the capacitor is now storing a voltage larger than VDD33. If a diode-connected PMOS was used, the parasitic PNP between the P+ contact (connected to the capacitor), the N-well tie (connected to VDD33) and the substrate (connected to VSS33) would inject charge into the substrate. This current is directed away from charging the capacitor, stresses the PNP device which could lead to

breakdown, and further presents a latchup hazard. However, if a diode-connected NMOS was used, the body-effect would be severe since the substrate is biased to VSS33. The best solution for this application was to use a diode-connected NMOS inside a triple well. By placing the NMOS in an isolated P-well, the body can be safely biased to VDD33 reducing the body effect. The N-well isolation should be biased to VDD33 to prevent the parasitic PNP from turning on. During overvoltage, this PNP would inject charge into the substrate. Similar to the diode-connected PMOS case described above, this current potentially poses a latchup hazard. Unfortunately,

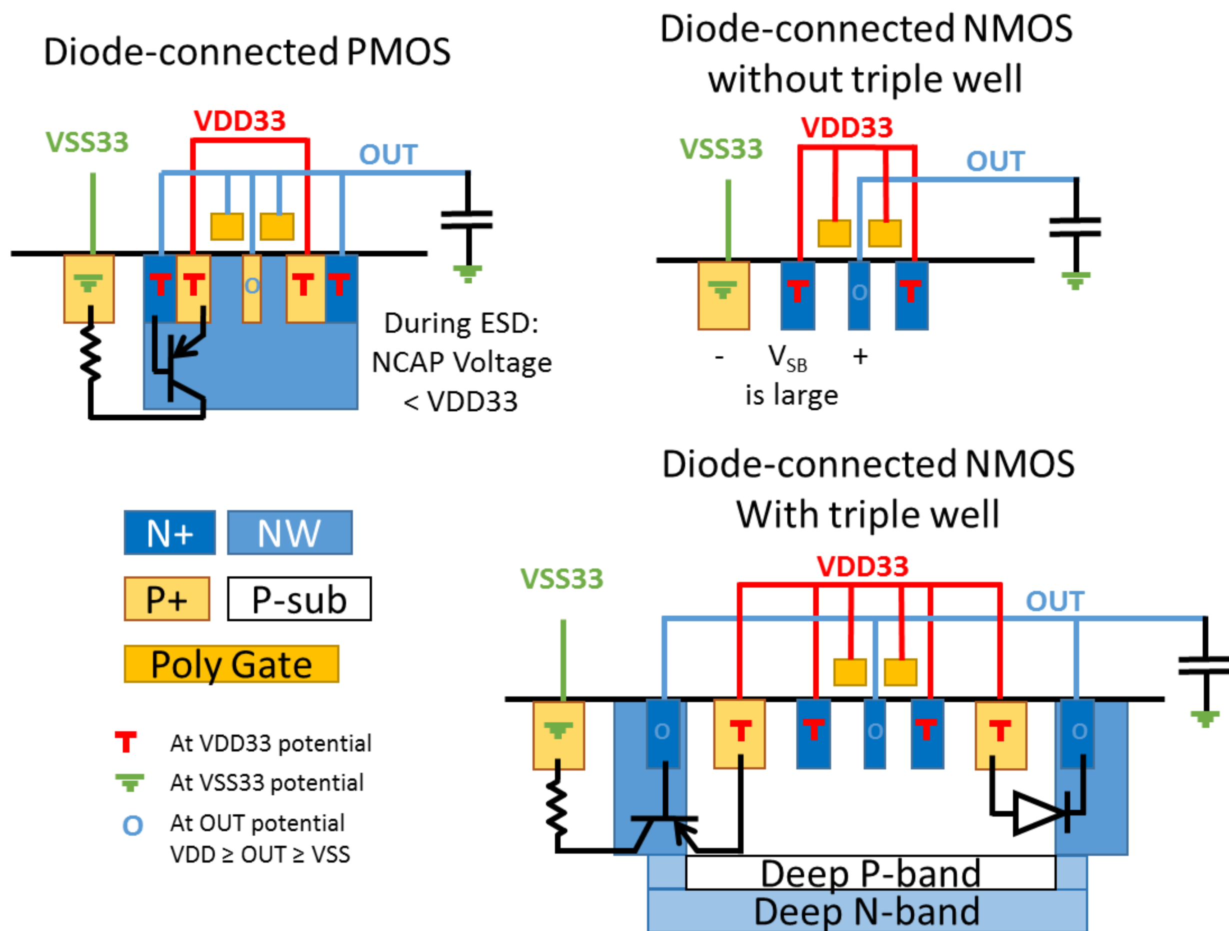

Figure 28: Cross section of three types of diode-connected MOS layouts. For the over-voltage monitor, the parasitic PNP of the PMOS would inject current into the substrate posing a latchup hazard. The NMOS without triple-well isolation would suffer from large body-effect. By utilizing an isolated P-well a diode-connected NMOS can be used. However, due to a design error the N-well isolation was mistakenly connected to the NCAP. This introduces a parasitic PNP device that poses a latchup hazard.

during the design this parasitic device was overlooked and the N-well isolation was connected the NCAP in order to insert a parasitic diode between the isolated P-well contact and the N-well isolation. This diode is forward biased when VDD33 is greater than the NCAP potential and would help to charge the NCAP during an ESD event.

Finally, the over-voltage monitor for the low-voltage domain (LVUV) is shown in Figure 29. The topology of this over-voltage monitor is similar to the previous except the diode-capacitor chain is placed in low-voltage domain (eVDD shown). Unlike the low-voltage under-voltage monitor, the resistor chain of this monitor is placed in the VDD33 domain along with the rest of the ADC circuitry. By doing this, the flash ADC can be used to detect voltages greater than eVDD without using a resistive divider on the NCAP. The values in the resistor chain have been set to produce reference voltages of 1.65 V and 1.8 V. A simulation of the LVUV monitor is shown in Figure 30. OUT1 is set high once the potential of eVDD is greater than 1.65 V plus the drop across the diode-connected NMOS.

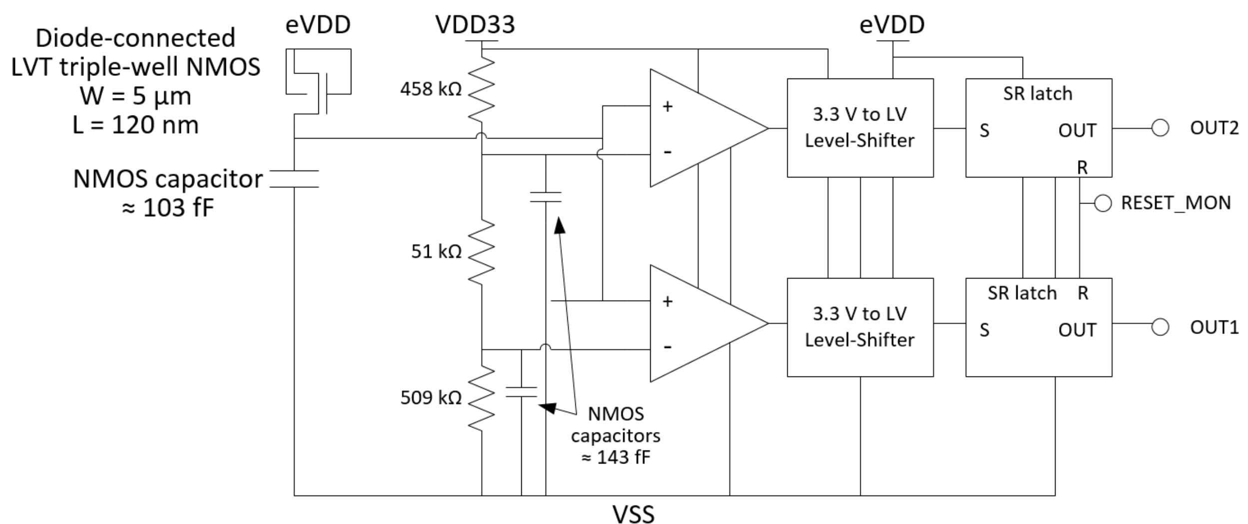

Figure 29: Over-voltage monitors for the internally and externally generated low-voltage domains.

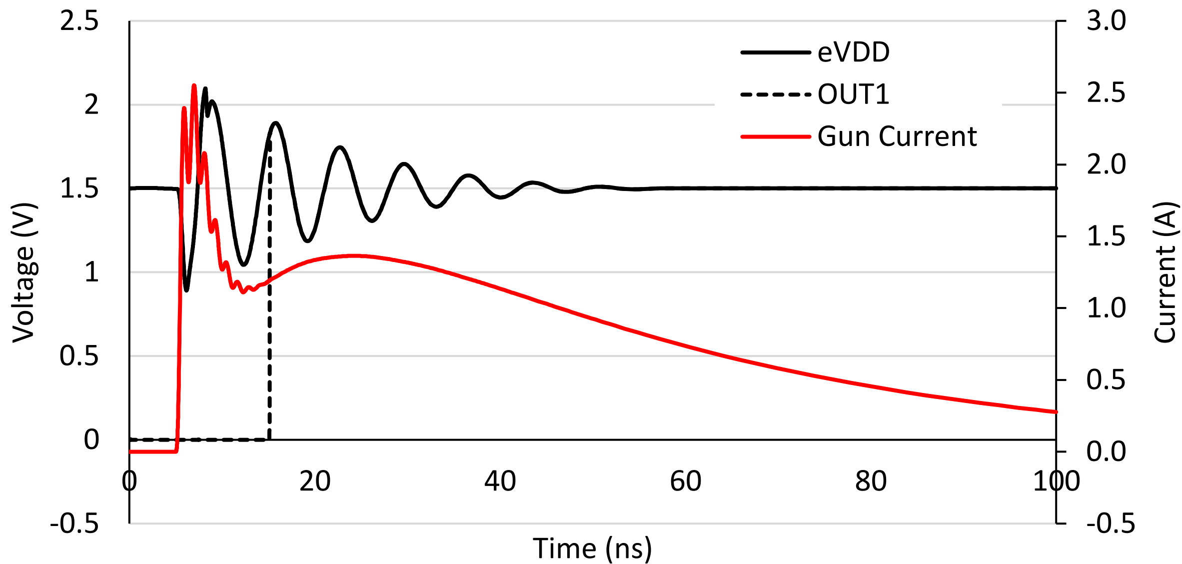

Figure 30: Simulation of the low-voltage domain over-voltage monitor.

An instance of each of these monitors has been placed as a test structure in the middle of the test chip along with sample-and-hold test structures. There are four sample-and-hold test structures which consist of a diode-capacitor sampling element whose output is connected to a source-follower amplifier. An example schematic (LVUV) is shown in Figure 31. The source-follower allows the stored voltage across the capacitor, V_C, to be measured by probing the output, V_{OUT}, without loading the diode during the stress. A MOS capacitor, C_G, is added to the node V_C to mimic the parasitic capacitance of the comparator found in the actual voltage monitor. A very fast transmission line pulse (VFTLP) generator is used to inject 4-ns-wide square pulses onto the power-bus being tested (VDD in this example), outside of the test system, as shown in Figure 32. The pulse is capacitively coupled to the power bus using a bias tee. The DC supply voltage is delivered through the choke found within the bias tee. During injection the power-bus voltage (with superimposed pulse) and the output are measured simultaneously using an oscilloscope. A high-impedance pick-off is used on the power-supply pad and the V_{OUT} pad to prevent the oscilloscope's low input impedance from loading the supply and test structure. An example

measurement is shown for each test structure in Figure 33. Each sample-and-hold circuit responds to the over- or under-voltage and the post-stress, steady-state value of the source-follower output, V_{out}, is seen to deviate from the pre-stress value. The change in V_{OUT} is less for the HVUV sample-and-hold than for the other three; thus it appears that the HVUV sample-and-hold circuit is not as sensitive as the others. The LVUV source-follower output takes a long time (500 ns) to settle. This is because the large capacitance associated with the probing pad must discharge through the large resistance R_S while, for the OV cases, the NMOS is responsible for charging this cap.

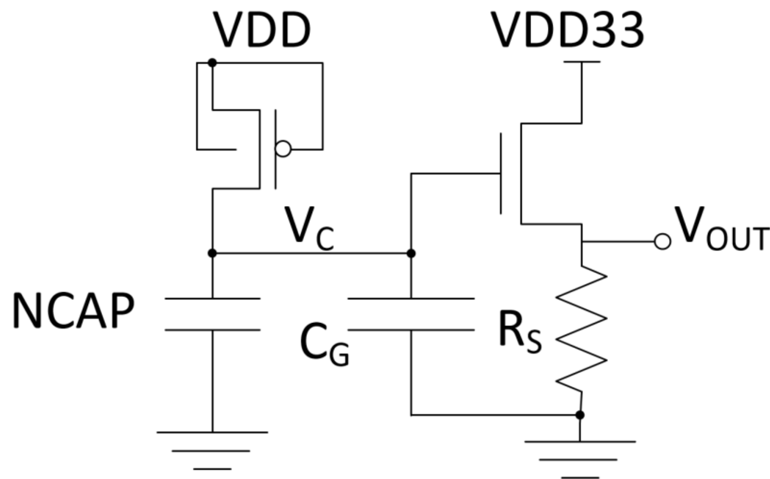

Figure 31: LVUV sample-and-hold test structure with source-follower.

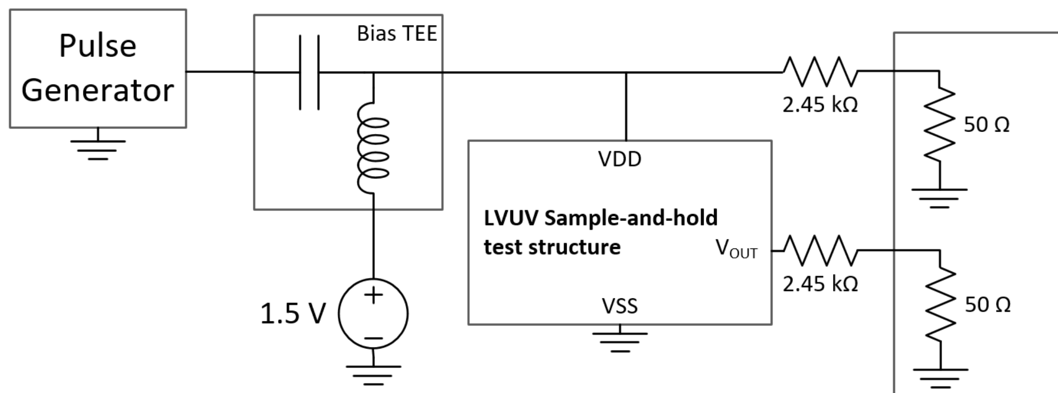

Figure 32: Measurement setup for testing the sample-and-hold test structure.

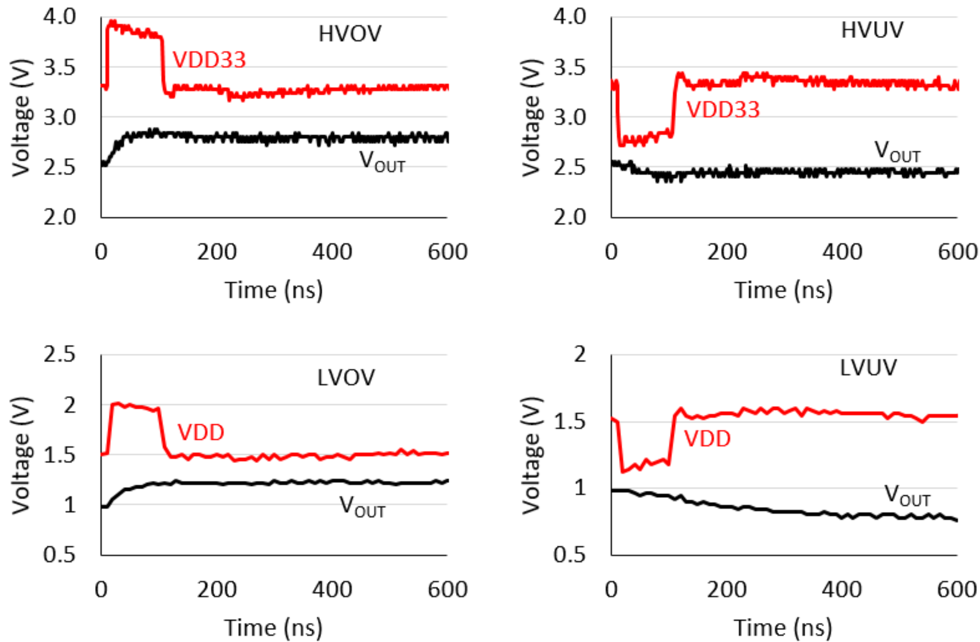

Figure 33: Output from the four sample-and-hold test structures. A VFTLP pulse was injected on the VDD33 and VDD pads.

The full test structure was probed and used to calibrate the output signals to known voltage levels. This was done by applying pulses to the power supply of the test structure of varying magnitude and pulse width to determine how the digital output signals respond to various stresses. A VFTLP generator is once again used to inject square pulses onto the power bus being tested within the test structure. The pulse width was set to either 4 ns or 100 ns; a 200 ps rise-time filter was used in both cases. As shown in Figure 34(a), the square pulse is coupled through a bias tee onto the VDD33 bus for testing the HVUV and HVOV monitors. Similarly, the pulse can be coupled onto the VDD domain for testing the LVUV and LVOV monitors as shown in Figure 34(b). The power bus voltage is monitored through a high-impedance pick-off using an oscilloscope. The test structure is reset using a toggle switch prior to the pulse injection. The reset signal resets the registers within all of the voltage monitors. While each voltage monitor has two unique outputs, only two are monitored at a time using digital voltmeters, all other monitors' outputs are left floating.

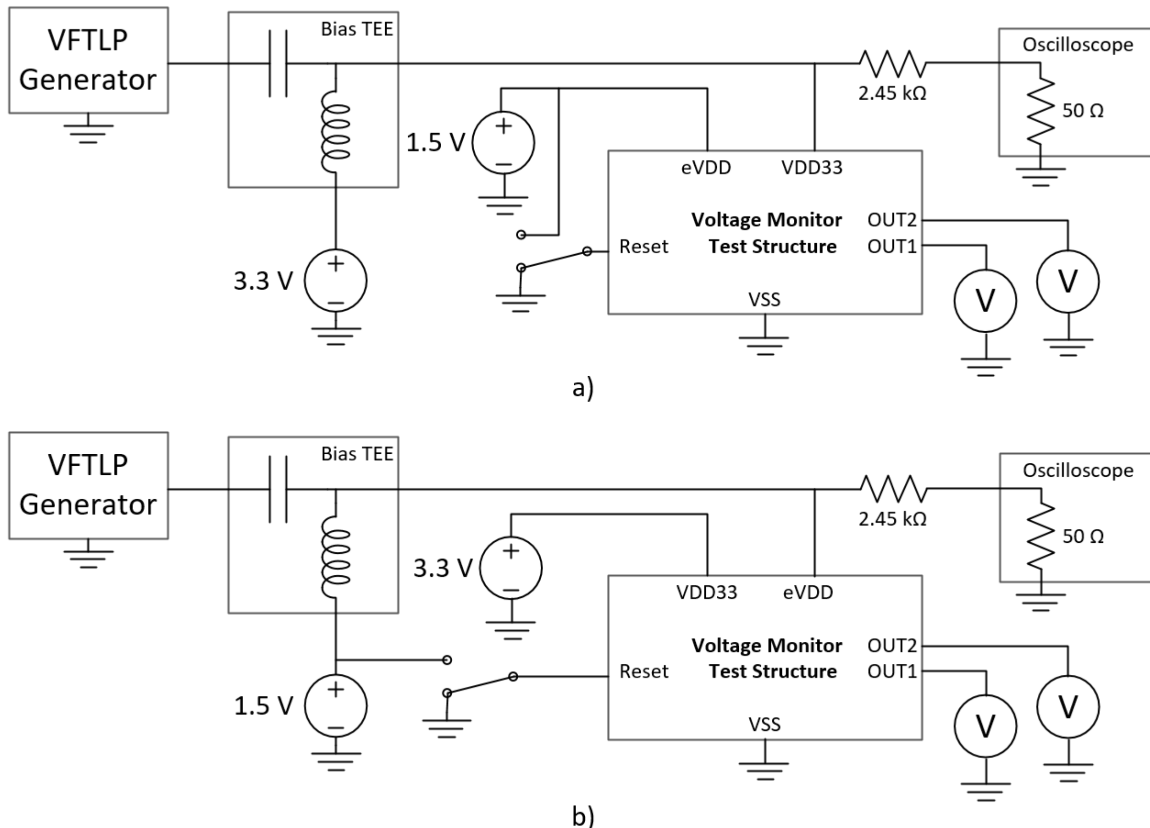

Figure 34: Measurement setup used to test voltage monitor test structures. The test structures share a common VDD33, VDD and VSS bus. VSS and VSS33 are shorted together within the test structures. A VFTLP is applied to the selected power bus using a bias tee. The pulse is applied to: (a) VDD33 for testing the HVUV and HVOV monitors and (b) VDD for testing the LVUV and LVOV monitors.

Pulses of varying amplitude were injected on the power bus under test. The voltage monitor's output levels are read from the voltmeters following each pulse, and the output code is tabulated along with the peak under- or over-voltage measured by the oscilloscope. It is important to note that the peak under- or over-voltage is measured on the power bus being tested and not at the capacitor node. There will be a voltage drop across the diode element that will not be reflected in the peak voltage values reported. The results from testing the HVUV monitor are shown in Table 16. The first under-voltage level appears to be less sensitive than originally designed for. It is likely that this loss in sensitivity results from the diode-capacitor sampling stage as it was shown to be less sensitive compared to the other three sampling stages (refer to Figure 33). The thresholds for the other monitor levels track the DC levels closely, if the threshold voltage, V_t, of

the 3.3 V PMOS (V_{t0} = -320 mV) is accounted for. Further, the monitor is more sensitive to under-voltage if the pulse-width is increased. This makes sense since as the diode-capacitor sampling element will low-pass filter the transient. If the under-voltage is applied for a longer duration, more of the charge can be removed from the capacitor and the capacitor voltage can fall to a lower level. However, the similarity between the monitor's response to the 100 ns pulse-width and 4 ns pulse-width reveals that the monitor is quick enough to capture transients on this time scale.

Table 16: HVUV monitor test structure testing. Square pulses with increasing negative amplitude are superimposed onto the VDD33 bus of the test structure. The HVUV monitor results are compared to the measured peak under-voltage as well as the ADC resistor chain's reference levels as found through DC simulation.

4 ns pulse-width			100 ns pulse-width			ADC resistor chain DC levels		
V_{MIN} (V)	HVUV output bit: 2	HVUV output bit: 1	V_{MIN} (V)	HVUV output bit: 2	HVUV output bit: 1	V_{MIN} (V)	HVUV output bit: 2	HVUV output bit: 1
3.11	0	0	-	-	-	> 3	0	0
2.92	0	0	2.89	0	0	< 3 > 2	0	1
2.72	0	0	2.69	0	0			
2.54	0	0	2.48	0	0			
2.35	0	0	2.26	0	0			
2.13	0	0	2.06	0	1			
1.94	0	1	1.86	0	1	< 2 > 1	1	0
1.75	0	1	1.66	1	0			
1.56	0	1	1.45	1	0			
1.36	1	0	1.23	1	0			
1.16	1	0	1.01	1	0			
0.95	1	0	0.81	1	0	< 1	1	1
0.78	1	0	0.58	1	1			
0.57	1	0	0.38	1	1			
0.36	1	1	0.14	1	1			
0.17	1	1	-	-	-			

The results from testing the LVUV monitor are shown in Table 17. Once again good agreement is observed between the results from both pulse-widths with the monitor showing slightly more sensitivity to the 100 ns pulse-width. The LVT PMOS used within the LVUV monitor has V_{t0} = -195 mV. Accounting for this voltage drop again yields very good agreement

between the pulse test results and the simulated DC thresholds. The LVUV monitor does not show diminished sensitivity for the first monitor level as the HVUV monitor did.

Table 17: LVUV monitor test structure testing. Square pulses with increasing negative amplitude are superimposed onto the eVDD bus of the test structure. The LVUV monitor results are compared to the measured peak under-voltage as well as the ADC resistor chain's reference levels as found through DC simulation.

4 ns pulse-width			100 ns pulse-width			ADC resistor chain DC levels		
V_{MIN} (V)	LVUV output bit:		V_{MIN} (V)	LVUV output bit:		V_{MIN} (V)	LVUV output bit:	
	2	1		2	1		2	1
-	-	-	-	-	-	> 1.36	0	0
1.34	0	0	1.34	0	0	< 1.36 > 0.90	0	1
1.15	0	1	1.13	0	1			
0.94	0	1	0.90	0	1			
0.74	0	1	0.71	1	0	< 0.90 > 0.45	1	0
0.54	1	0	0.49	1	0			
0.34	1	0	0.28	1	1	< 0.45	1	1
0.13	1	1	0.10	1	1			
0.74	1	1	-0.13	1	1			

Test results from the HVOV monitor are shown in Table 18. Here, pulses with positive amplitude are applied to the power bus. Similar monitor response is seen for both pulse widths. The 4 ns pulse caused the second monitor level to latch at a lower over-voltage than the 100 ns pulse. This result is unexpected, but the difference is not large enough to warrant concern, especially with this limited dataset. The 3.3 V NMOS used in the HVOV monitor has V_{t0} = 380 mV. The offset between the DC simulated threshold levels and the level at which the monitors do set is approximately $V_{t0.}$ The parasitic PNP described above does not appear to be affecting the performance of the HVOV monitor, however, the PNP does still present a reliability concern.

The results from the LVOV monitor are presented in Table 19. The 4 ns pulse causes the monitor to set at a higher over-voltage than the 100 ns pulse, as expected. The LVT NMOS used in this monitor has V_{t0} = 245 mV. The monitor's response to pulse-testing tracks the DC threshold levels as expected.

Table 18: HVOV monitor test structure testing. Square pulses with increasing positive amplitude are superimposed onto the VDD33 bus of the test structure. The HVOV monitor results are compared to the measured peak over-voltage as well as the ADC resistor chain's reference levels as found through DC simulation.

<table>
<tr><th colspan="3">4 ns pulse-width</th><th colspan="3">100 ns pulse-width</th><th colspan="3">ADC resistor chain DC levels</th></tr>
<tr><th rowspan="2">V_{MAX} (V)</th><th colspan="2">HVOV output bit:</th><th rowspan="2">V_{MAX} (V)</th><th colspan="2">HVOV output bit:</th><th rowspan="2">V_{MAX} (V)</th><th colspan="2">HVOV output bit:</th></tr>
<tr><th>2</th><th>1</th><th>2</th><th>1</th><th>2</th><th>1</th></tr>
<tr><td>-</td><td>-</td><td>-</td><td>-</td><td>-</td><td>-</td><td>< 3.34</td><td>0</td><td>0</td></tr>
<tr><td>-</td><td>0</td><td>0</td><td>3.53</td><td>0</td><td>0</td><td rowspan="4">> 3.34
< 4.15</td><td rowspan="4">0</td><td rowspan="4">1</td></tr>
<tr><td>3.67</td><td>0</td><td>0</td><td>3.73</td><td>0</td><td>0</td></tr>
<tr><td>3.88</td><td>0</td><td>1</td><td>3.94</td><td>0</td><td>1</td></tr>
<tr><td>4.07</td><td>0</td><td>1</td><td>4.13</td><td>0</td><td>1</td></tr>
<tr><td>4.26</td><td>0</td><td>1</td><td>4.64</td><td>0</td><td>1</td><td rowspan="6">> 4.15</td><td rowspan="6">1</td><td rowspan="6">1</td></tr>
<tr><td>4.45</td><td>0</td><td>1</td><td>4.54</td><td>0</td><td>1</td></tr>
<tr><td>4.66</td><td>1</td><td>1</td><td>4.73</td><td>0</td><td>1</td></tr>
<tr><td>4.85</td><td>1</td><td>1</td><td>4.92</td><td>1</td><td>1</td></tr>
<tr><td>5.04</td><td>1</td><td>1</td><td>5.13</td><td>1</td><td>1</td></tr>
<tr><td>5.23</td><td>1</td><td>1</td><td>5.34</td><td>1</td><td>1</td></tr>
</table>

Table 19: LVOV monitor test structure testing. Square pulses with increasing positive amplitude are superimposed onto the eVDD bus of the test structure. The LVOV monitor results are compared to the measured peak over-voltage as well as the ADC resistor chain's reference levels as found through DC simulation.

<table>
<tr><th colspan="3">4 ns pulse-width</th><th colspan="3">100 ns pulse-width</th><th colspan="3">ADC resistor chain DC levels</th></tr>
<tr><th rowspan="2">V_{MAX} (V)</th><th colspan="2">LVOV output bit:</th><th rowspan="2">V_{MAX} (V)</th><th colspan="2">LVOV output bit:</th><th rowspan="2">V_{MAX} (V)</th><th colspan="2">LVOV output bit:</th></tr>
<tr><th>2</th><th>1</th><th>2</th><th>1</th><th>2</th><th>1</th></tr>
<tr><td>-</td><td>-</td><td>-</td><td>-</td><td>-</td><td>-</td><td>< 1.65</td><td>0</td><td>0</td></tr>
<tr><td>1.71</td><td>0</td><td>0</td><td>1.70</td><td>0</td><td>0</td><td>> 1.65</td><td>0</td><td>1</td></tr>
<tr><td>1.9</td><td>0</td><td>0</td><td>1.92</td><td>0</td><td>1</td><td rowspan="4">> 1.82</td><td rowspan="4">1</td><td rowspan="4">1</td></tr>
<tr><td>2.10</td><td>0</td><td>1</td><td>2.15</td><td>1</td><td>1</td></tr>
<tr><td>2.29</td><td>1</td><td>1</td><td>2.31</td><td>1</td><td>1</td></tr>
<tr><td>2.47</td><td>1</td><td>1</td><td>2.47</td><td>1</td><td>1</td></tr>
</table>

Overall, the performance of all four versions of voltage monitors has been verified for square-pulses injected onto the power bus under test. With the exception of the decreased sensitivity observed for the HVUV monitor, all monitors responded to the under- or over-voltage in the expected manner.

5.1.2 Static Latches

Static latches of the same topology as shown in Figure 13 are included on this test chip. The latches will have a balanced topology, i.e., the feed-forward and feed-back inverters are

identical. The inputs to these latches will be generated externally and supplied to the receiver in the IO ring. Buffers are placed along the path from the receiver to the latch. The latches placed in monitor bank 1 receive their data inputs from the DIN33 cell and the transmission gates are transparent when the STROBE33 signal is high. The IO cells supplying these inputs are standard receivers powered from the VDD33 domain. The latches placed in monitor banks 2 receive their data inputs from the DIN cell and the transmission gates are transparent when the STROBE signal is high. The IO cells supplying these signals are low-voltage receivers powered from the eVDD domain. A comparison of upsets between latches in monitor banks 1 and 2 will yield information about whether receivers designed for the 3.3 V domain are more or less susceptible to noise than those designed for the low-voltage domain. The latches placed in monitor bank 3 also receive their data input from the DIN cell and are controlled by the strobe signal. However, the latches themselves are powered from the internally generated iVDD1 domain. A comparison of upsets between latches in monitor banks 2 and 3 will yield information about whether ESD generates more or less noise on internally generated supplies than on externally generated supplies.

In addition to the previously mentioned latches, unbalanced latches are also included. These latches have feed-forward and feed-back inverters that, through simulation, have been tailored to favor one logic level over the other. These latches have hard-tied inputs that are opposite to the favored logic level and are written into the latch when the appropriate strobe signal goes high. Several versions of these latches, with varying levels of unbalance, are placed in each monitor bank. For lower-level stresses, it is expected that the most unbalanced latches will flip their states, referred to as an upset. For higher-level stresses, the more balanced latches are expected to upset. In this manner, a different type of power-supply noise-monitor is formed. The results obtained from these latches will be cross-referenced to the results from the over- and

under-voltage monitors. The same strobe signal used for the balanced latches are used for these latches. Once again, by comparing the results from the three monitor banks, differences in the noise generated within each of the power domains can be studied.

5.1.3 Dynamic OR Gates

Dynamic OR gates were placed within the three monitor banks. These gates have the same topology as those placed on the previous test chip (shown in Figure 13). Within each monitor bank there are three OR gates. All three OR gates are put into their precharge mode by setting the appropriate strobe signal high. The gates are in evaluation mode when that strobe signal is low. These gates all have four inputs, with three being hard-tied low using a tie-low cell near the gate. The three gates are presented in Figure 35. The tie-low cell uses a diode-connected PMOS to pull the gate of an NMOS high, while that NMOS is used to pull the OR-gate's inputs low (shown explicitly in Figure 35(a)). The tie-low cells are used to avoid reliability concerns that result from tying MOS gates directly to power nets. The gate in Figure 35(a) has the fourth input supplied by an off-chip signal that is buffered from the IO cell to the OR gate. The gate in Figure 35(b) has the fourth input supplied by a tie-low cell that is buffered twice before reaching the OR gate. The gate in Figure 35(c) has the fourth input directly tied low, along with the others, through a tie-low cell. The purpose of these gates is to discern how noise is manifest in the chip's core. The gate in Figure 35(c) is only likely to upset if power noise causes the dynamic node to become erased, as all of the inputs are directly tied to the reference supply, noise should not create a glitch on the input signal. The gate in Figure 35(b) could upset in the same manner as the gate in (c), but could also upset if localized supply noise cause the buffers that supply the low signal to upset. The gate in Figure 35(a) could upset as in the previous two cases or if an input glitch occurs in the chip's IO ring.

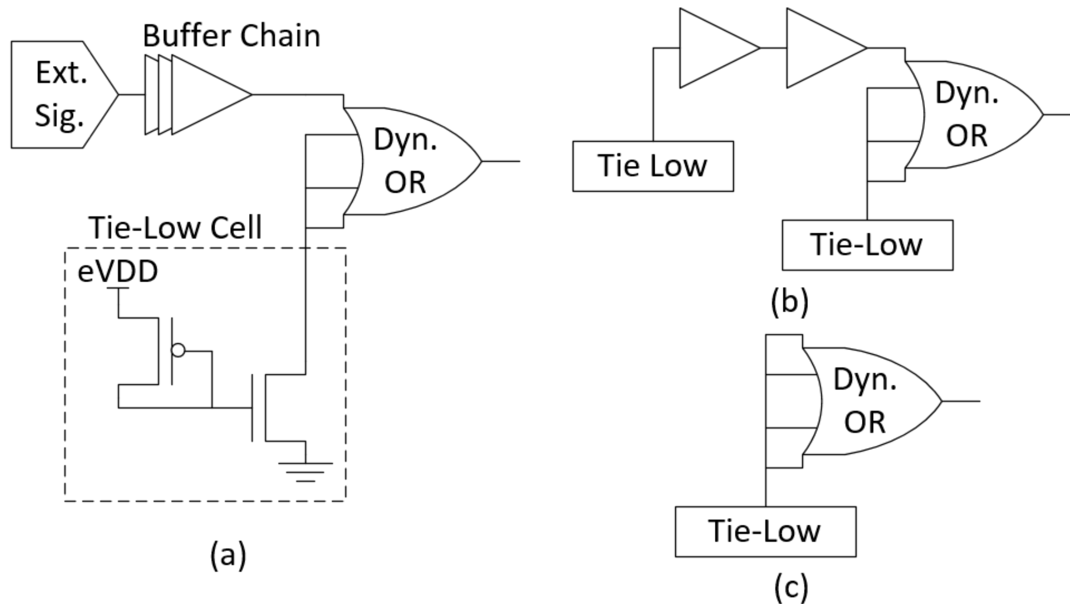

Figure 35: Dynamic OR gates included in each monitor bank. The gate in (a) has one input supplied via an off-chip signal. The gate in (b) has one input supplied via a tie-low cell through two buffers. The gate in (c) has all four inputs tied low directly.

The OR gates in monitor bank 1 have their strobe signal supplied by the signal STROBE33 and the off-chip input is supplied by DIN33. The OR gates in monitor banks 2 and 3 have their strobe signal supplied by STROBE and the off-chip input is supplied by DIN. Monitor bank 2 has one additional OR gate. This gate is identical to the gate with an externally supplied signal, except its strobe signal is supplied by STOBEX and its input is supplied by DINX. These two IOs are powered from the 3.3 V domain but have an out-of-range error detector (ORED) incorporated into each cell. The ORED circuit, shown in Figure 36, is designed to latch the OREDH bit high if the input to the chip exceeds the VDD33 supply, and latch the OREDL bit high if the input falls below VSS33 [18]. This experiment may yield some quantitative information about glitches seen at a chip input and allow that information to be correlated with upsets of internal logic.

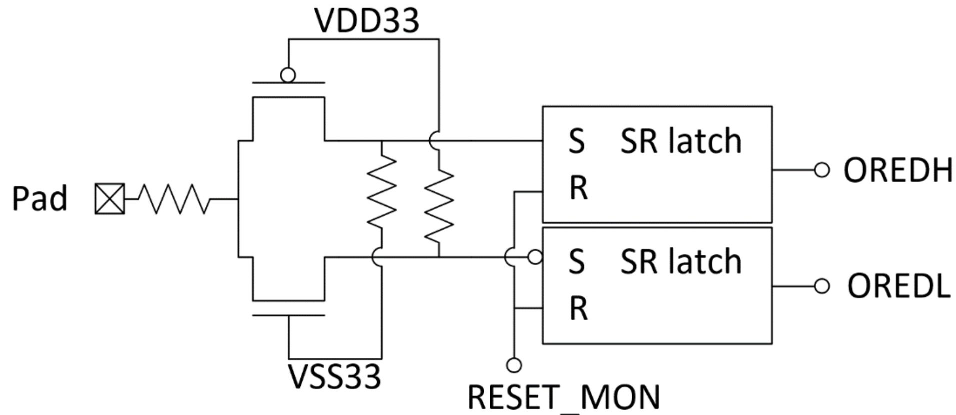

Figure 36: ORED circuit. If the voltage at the pad is greater than VDD33 than the top latch will set OREDH high, otherwise the pull-down resistor will keep the set signal low. Similarly if the voltage at the pad falls below VSS33, then the bottom latch will set OREDL, otherwise the pull-up resistor will keep the set signal high.

5.1.4 Shift-Registers

The final type of logic included in each of the monitor banks is a shift-register. Since a robust hypothesis for the cause upsets observed in the previous test chip's shift-register could not be developed, several variations of that shift-register have been placed on this test chip. All of the shift-registers on this test chip are 4 bits long. Figure 37 shows a schematic representation of the shift-registers. The first shift-register, shown in Figure 37(a), is a copy of the shift-register used within the previous test chip except only 4 bits long. Each bit of the shift-register is stored in a D flip-flop constructed using LVT transistors. The flip-flop was implemented with a set and reset terminal (both active low). The set and reset signals were both hard-tied to the low-voltage power supply (eVDD shown). In order to replicate what was done on the previous test chip, these two signals were tied high without utilizing tie-high cells, despite this being a bad practice. The inclusion of these two inputs could allow additional means for supply noise to upset the contents of the shift-register.

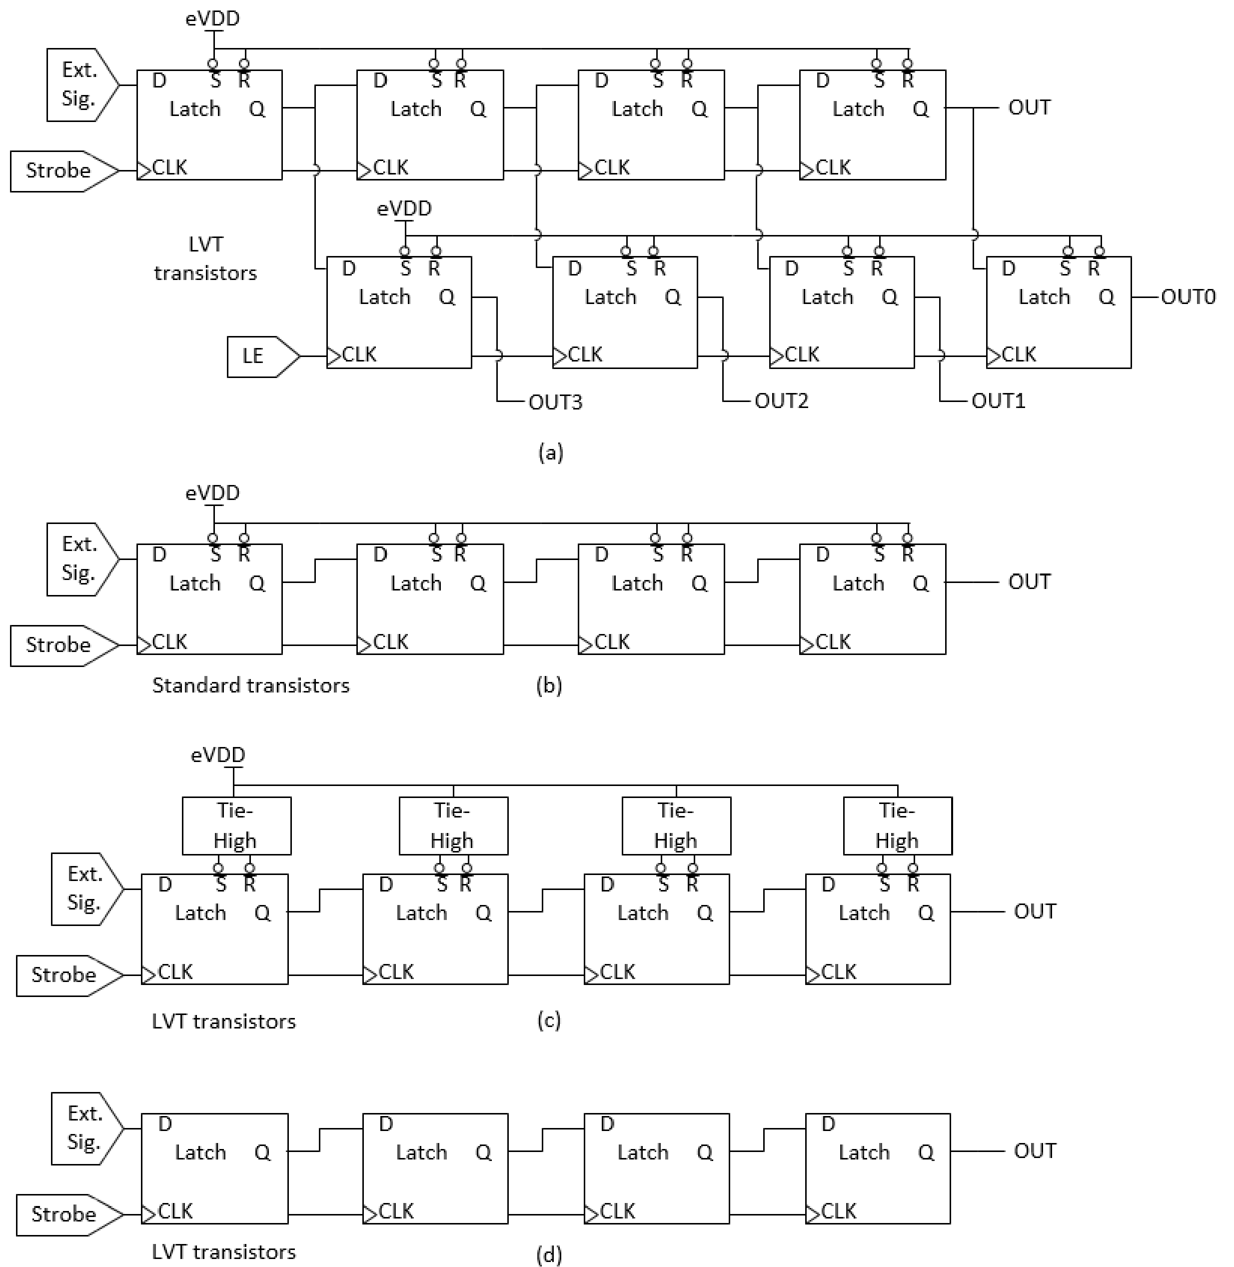

Figure 37: Several four-bit shift-registers are included on the new test chip.

Each bit of the shift-register in Figure 37(a) is stored in a separate output-latch when the LE signal is applied. Each of these outputs is routed to the multiplexer. This design is supposed to improve noise immunity. If data are clocked into the shift-register erroneously, e.g. due to noise on the data and strobe signals, the output of the latches would not change along with the shift-

register unless the LE signal is asserted. Additionally, the serial output of the shift-register was routed to the multiplexer so that the contents of the shift-register may be separately determined from what the output-latches indicate. This serial output was not available in the previous test chip. By including the serial output, it can be determined whether the contents of the shift-register are being upset, the contents of the output-latches are being upset, or both contents are being upset.

The second shift-register, shown in Figure 37(b), is a variation of the previous shift-register. Instead of using LVT transistors, the flip-flops were constructed with standard threshold-voltage (SVT) transistors. Further, to reduce the number of outputs (and therefore the required size of the multiplexer) the output-latches were omitted and only the serial output was routed to the multiplexer. By comparing the upsets found in this variation to those of the previous variation, it can be determined whether flip-flops made of SVT transistors are more immune to upset than flip-flops made of LVT transistors. The third shift-register, shown in Figure 37(c), is a variation of the shift-register in Figure 37(a) that uses tie-high cells to supply the active-low set and reset inputs. Again, only the serial output of this shift-register is routed to the multiplexer. This experiment should determine whether these tie-high cells offer any filtering of supply noise that could be causing the flip-flops to upset. The final shift-register, shown in Figure 37(d) is constructed using D flip-flops without the set and reset capability. By omitting these inputs entirely and comparing upsets observed in shift-register Figure 37(a) with those observed in shift-register Figure 37(c), it can be determined whether the set and reset lines are at all responsible for upsets within the shift-register contents.

Once more, this entire set of shift-registers is duplicated in monitor banks 1, 2, and 3. The shift-registers in monitor bank 1 use the STROBE33 signal to supply the strobe signal for

clocking in data. The data input to the first bit of the shift-register are supplied by SR_DIN33. These signals both come from receivers in the VDD33 domain. The shift-registers in monitor banks 2 and 3 both use STROBE and SR_DIN to supply the strobe and data input signals respectively. These signals come from receivers in the eVDD domain. This will again allow for comparison of noise susceptibility of different supply domains.

5.2 Circuit Board and Test Set-Up Designs

A custom test board was designed and fabricated to support the experimental work using the new test chip. A photograph of the board is shown in Figure 38. Similar to the previous test system, a four-layer FR4 board is used with dedicated supply and ground planes to represent common practices. The circuit board can be powered by a DC supply or by a battery pack; the chosen supply is connected to the DC jack shown in Figure 38. A fixed-voltage LDO regulator supplies 3.3 V for the VDD33 domain and peripheral components within the test system. A variable LDO provides 1.2 to 1.5 V for the eVDD domain. This supply also acts as the reference for the internally generated iVDD1 domain. Decoupling capacitance is included on each of the

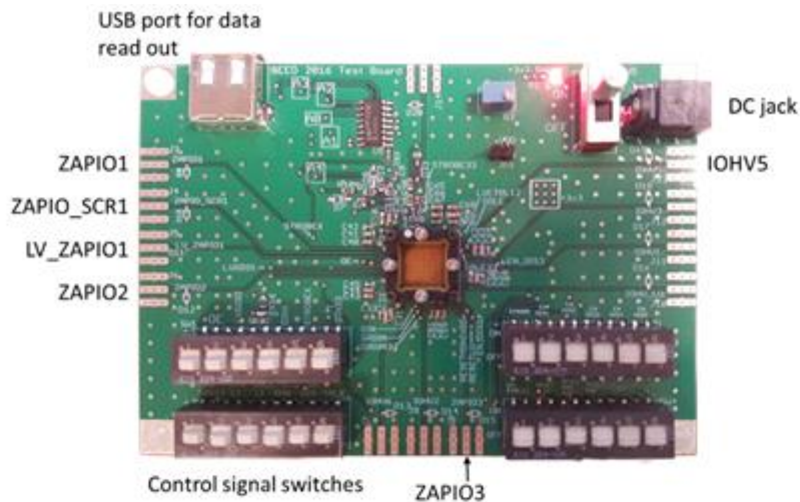

Figure 38: Test board photograph.

power nets, including a variety of capacitor values and sizes. Several values and sizes of SMD decoupling capacitors are placed on every supply pin, near the chip, following best practices.

The three on-chip multiplexers are read out following ESD experimentation by attaching a separate microcontroller board via a USB port as shown in Figure 39. A USB port is used for convenience, yet the USB standard and protocols are not followed. Instead the four lines of the USB port are used for custom applications. One line of the USB port supplies a clock signal to the board, another supplies a reset, the third line is the data line used to carry data from the test system to the microcontroller board, and the last line is grounded on both boards. During testing the microcontroller board must be disconnected from the test system to remove the unwanted tethering and to avoid stressing the readout equipment. The connection between the microcontroller board and the test system is made using a magnetic quick-release USB cable. This cable allows the microcontroller board to be quickly disconnected and reconnected accurately to speed up testing.

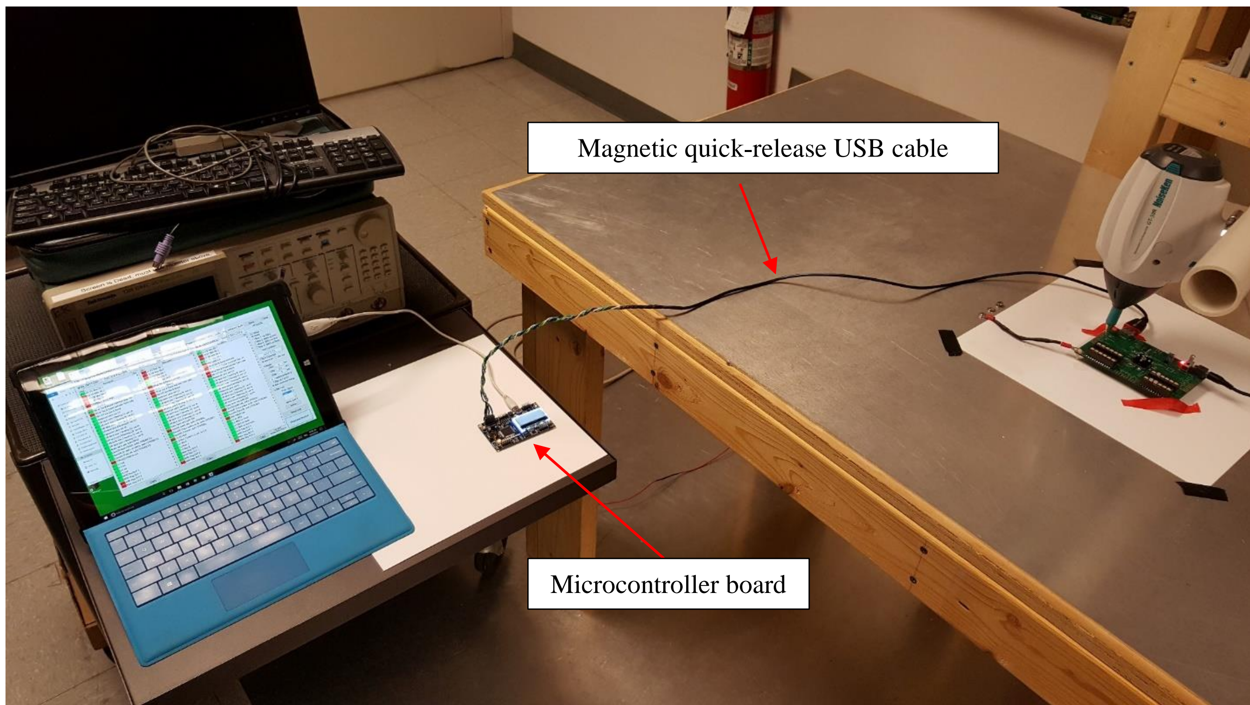

Figure 39: Data readout setup.

The microcontroller board controls the data read operation. A read operation is triggered when the computer sends the read command to the microcontroller board. The microcontroller asserts the reset first to set a counter, located on the test system, to output all zeros. One bit is read from the USB data line and then a clock signal is applied to the clock line causing the counter to increment. One bit is read from the data line for each clock cycle. The output lines from the counter are used to supply the address bits to the test chip as well as to two on-board multiplexers. One multiplexer has hard-tied inputs. As the address is incremented, a known progression of data bits is generated from this multiplexer. This string of bits is used as a header to validate the data that are recorded by the microcontroller board. If any errors in timing occur, for example if an extra clock edge is seen at the counter, the header would be wrong indicating an error. A second multiplexer is used to switch between transmitting the header and the contents of each monitor bank, sequentially. After 128 bits have been transmitted the microcontroller board will have received a 32-bit header and the 96 bits of data from the three on-chip monitor banks.

The data collected by the microcontroller board are converted to hexadecimal characters and transmitted to the computer. Custom readout software was written to accept the board data, display it in real-time and store it for future analysis alongside experiment details (e.g. stress level, test chip, and zap location). Figure 40 shows a screenshot of the data readout software. The software verifies the header and presents the data bitwise alongside description of each bit. The bits that have changed from the pre-zap conditions are colored red. The software can read the data and save it for later analysis or simply read the data and update the display. The latter case is used to verify that the bits within the test chip have been set to the proper states between zaps.

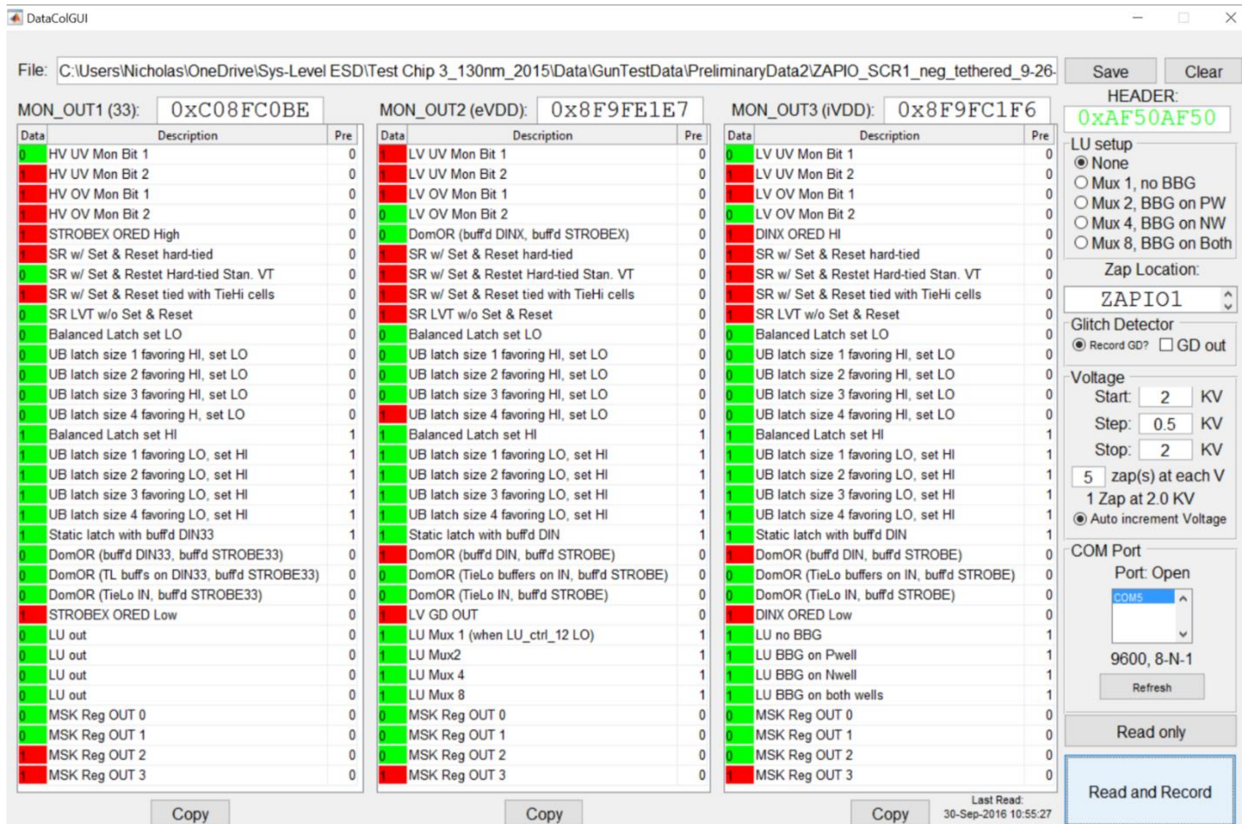

Figure 40: Screenshot of data readout software.

The chip IO pins that are intended to undergo ESD zapping are routed to the board edge to facilitate contact discharge testing. Landing pads for transient-voltage suppressors (TVS) are located one each zap line near the zap points. TVS devices can be included on these pads to study the effect external noise suppression has on the occurrence of soft failures. The various control signals are supplied via toggle switches. The signal lines from these control switches are heavily filtered as they were on the previous test system. The exception to this are the signal lines that drive the LVGDIO and GDIO pins. Since glitches are being studied on these lines, these filters, which could prevent glitches from forming, were omitted. Further, the LVGDIO and GDIO signal lines are driven by a low-voltage buffer and a 3.3 V buffer, respectively. The on-board traces that deliver the LVGDIO and GDIO signals to the test chip are routed in parallel to, and on both sides

of, the ZAPIO2 signal line. This was done so that the ZAPIO2 line could act as a comparable aggressor for each of the glitch-detector signals.

5.3 IEC 61000-4-2 Testing of New Test System

In order to explore the hypotheses developed with the first test system, ESD discharges were applied to the new test system according to the IEC 61000-4-2 standard. The primary goal of the new test system is to compare logic upsets in the core to supply noise. First, the voltage monitors' responses to zaps will be presented to verify their functionality during system-level ESD stresses to the test system. Next, logical upsets within the banks of static latches, dynamic or gates, and shift-registers will be compared to the readings from the voltage monitors, glitch detectors and OREDs. Unless otherwise noted, all tests were performed with eVDD (and by extension iVDD) set to 1.5 V and without any TVS protection on the test system.

The results from over 2500 zaps applied to five test chips have been aggregated to form a single dataset from which the following results derived. Zaps were applied to various dedicated zap IOs as well as to the USB shield which is connected to board ground. The test system was tested in both mobile and tethered configurations. Before every zap the data readout board was connected to the test system and the test system was set to the starting condition. The data readout software was used to read the system data to ensure the starting conditions are met. The data readout board was disconnected from the test system and the system was zapped. Following the zap the data readout board was reconnected and the system data were read and stored. This procedure was repeated for subsequent zaps. Every test was repeated at least five times before any changes were made to the setup.

5.3.1 Power Supply Noise Detected by Voltage Monitors

The results from the voltage monitors are reported as the largest monitor output level (converted from binary to decimal) that occurs for at least 40% of the applicable zaps as is indicated by Equation (2). Here L indicates the over- or under-voltage level, Z_i is the total number of zaps resulting in over- or under-voltage level i (from 0 to L_max), and Z_T is the total number of zaps. For example, if the monitor output is 1 for 60% of the zaps, 2 for 30%, and 3 for 10%, then the reported level will be 2, as this level was met or exceeded 40% of the time.

$$\max(L) \;\Big|\; \sum_{i=L}^{L_{max}} Z_i \;\geq 0.4 \cdot Z_T \tag{2}$$

Table 20 shows the recorded HVUV data for zaps to a mobile EUT and Table 21 shows the data for zaps to a tethered EUT. Due to the lower threshold of hard failures for tethered EUTs, some IOs were not stressed over the full range tested. Comparing the data within the two tables shows that the amount of under-voltage is very similar between zaps to a mobile EUT and zaps to a tethered EUT, only slight differences in thresholds can be discerned. A more interesting result is the difference in under-voltage between negative and positive zaps. Positive zaps to the 3.3 V IOs result in larger under-voltages than do negative zaps. As indicated in Section 4.3, a zap of either polarity is expected to induce both an over-voltage and an under-voltage since the ESD current pulse contains edges with +di/dt and –di/dt. Asymmetric results are not surprising given that the current return paths are different for negative and positive zaps, and that the return paths contain both linear and non-linear (passive and active) devices. Preliminary circuit simulations appear to be consistent with the measurement results, showing that a positive zap induces the more severe under-voltage [19]. Another interesting result is that zaps to distant IOs (ZAPIO3 and IOHV5) show similar results highlighting that chip-wide VDD33 under-voltage occurred. Zaps to ZAPIO1 (in close proximity to the HVUV monitor in monitor bank 1) and zaps to ZAPIO2 (slightly further

away) both cause more under-voltage indicating that there is a proximity effect on top of the global under-voltage. However, zaps to ZAPIO3 result in less under-voltage than zaps to the more distance IOHV5. This indicates that there may be an additional location dependency that is unrelated to the proximity of the zap IO to the voltage monitor. Zaps to ZAPIO_SCR1, which is located very near the HVUV monitor, results in less under-voltage than do zaps to ZAPIO1. For positive stress the SCR will direct ESD current to VSS33, while the DD protection will direct current to VDD33. Apparently this difference in current path results in more under-voltage in the latter case.

Table 20: IEC gun testing of the HVUV monitor on a mobile EUT.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	1	1	1	1	1	-
-4	0	1	1	1	1	1
-3	0	0	0	0	0	0
-2	0	0	0	0	0	0
-1.5	0	0	0	0	0	0
-1	0	0	0	0	0	0
-0.5	0	0	0	0	0	0
0	0	0	0	0	0	0
0.5	0	0	0	0	0	0
1	1	0	0	1	0	0
1.5	1	1	1	1	0	0
2	2	2	1	2	1	0
3	3	3	1	2	2	0
4	3	3	2	2	2	0
5	3	3	2	2	2	1

Additionally, zaps to low-voltage IOs result in under-voltage on the 3.3 V supply rails, although the magnitude is reduced relative to the case when a 3.3 V IO is zapped. ESD current is shared between the domains via the APD located between the VSS and VSS33 buses. As ESD current injected into a low-voltage IO tries to exit the chip, a portion of the current will flow through the APD to the VSS33 bus. Some current will exit through the VSS33 bondwires and some will flow to the VDD33 bus and exit through the VDD33 bondwires. In this manner zaps to

IOs in one power domain can influence the voltage in the other power domain. Data for -5 kV zaps to LV_ZAPIO1 are not shown since the test chip experienced hard failure at this stress level.

Table 21: IEC gun testing of the HVUV monitor on a tethered EUT.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	0	0	0	0	-	-
-4	0	0	0	0	-	-
-3	0	0	0	0	-	-
-2	0	0	0	0	0	0
-1.5	0	0	0	0	0	0
-1	0	0	0	0	0	0
-0.5	0	0	0	0	0	0
0	0	0	0	0	0	0
0.5	0	0	0	0	0	0
1	1	0	0	0	0	0
1.5	1	1	1	1	0	0
2	1	1	1	1	1	0
3	2	3	1	2	-	0
4	3	3	1	2	-	-
5	3	3	3	3	-	-

The results reported from the LVUV monitor are shown in Table 22. Due to the similarity between data from a mobile and a tethered EUT, the results have been combined. Once again positive polarity zaps yield more reported under-voltage than negative zaps. Zaps to all IOs yield very similar results again indicating global power noise on the eVDD domain. Zaps to ZAPIO1 and ZAPIO2 (both are close to monitor bank 2 where the LVUV monitor is located) yield more under-voltage than do zaps to the more distant ZAPIO3 and IOHV5, indicating a proximity dependence. Zaps to ZAPIO_SCR1 again result in less under-voltage than zaps to the DD-protected 3.3 V IOs. While negative zaps to LV_ZAPIO1 cause more power noise on the eVDD domain than do zaps to 3.3 V IOs, positive zaps do not strictly follow this trend. Comparing the results of zaps to ZAPIO1 and LV_ZAPIO1, we observe that positive zaps to the 3.3 V domain cause larger under-voltage than do zaps to the LV domain.

Table 22: IEC gun testing of the LVUV monitor in the eVDD domain.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	1	2	1	1	2	-
-4	1	2	1	1	1	3
-3	1	1	1	1	1	3
-2	1	1	1	1	1	1
-1.5	1	1	1	1	1	1
-1	1	1	1	1	1	1
-0.5	0	0	0	0	0	0
0	0	0	0	0	0	0
0.5	1	0	0	0	0	1
1	1	1	1	1	1	1
1.5	2	2	1	1	1	2
2	3	3	1	1	2	2
3	3	3	2	1	3	3
4	3	3	3	2	3	3
5	3	3	3	2	3	3

Table 23 shows the results reported by the LVUV monitor in the internally generated iVDD1 domain. Overall, larger under-voltages are observed in the internally generated domain relative to eVDD. In simulation, the voltage regulator on this test chip was observed to track the eVDD supply (to which it is referenced) very closely. Thus, it makes sense that the iVDD domain will experience at least as much noise as the eVDD domain. The polarity dependence and proximity effect highlighted for the previous two monitors are also observed here.

Table 24 shows the results from the HVOV monitor. Since the lowest threshold for the over-voltage monitor is set very close to 3.3 V, most zaps result in a detected over-voltage. The same polarity dependence observed in the under-voltage monitors is exhibited here for zaps to DD-protected IOs, i.e., positive zaps appear to cause larger amplitude supply voltage excursions than do negative zaps. Interestingly, zaps to the SCR-protected IO and the low-voltage IO both show more over-voltage for negative zaps. Again, more supply noise is observed for zaps to nearby IOs (ZAPIO1 and 2) than for more distant IOs (ZAPIO3 and IOHV5).

Table 23: IEC gun testing of the LVUV monitor in the iVDD1 domain.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	3	2	3	2	3	NaN
-4	3	2	3	2	3	3
-3	3	2	3	2	3	2
-2	3	2	2	2	3	2
-1.5	2	2	1	2	2	2
-1	1	1	1	1	2	1
-0.5	1	1	1	1	1	1
0	0	0	0	0	0	0
0.5	1	1	1	1	1	1
1	1	1	2	1	1	2
1.5	1	1	3	2	2	3
2	1	2	3	2	3	3
3	3	3	3	3	3	3
4	3	3	3	3	3	3
5	3	3	3	3	3	3

Table 24: IEC gun testing of the HVOV monitor.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	2	2	1	1	2	-
-4	2	1	1	1	2	2
-3	2	1	1	1	2	2
-2	1	1	1	1	1	1
-1.5	1	1	1	1	1	1
-1	1	1	1	1	1	1
-0.5	1	1	1	1	0	0
0	0	0	0	0	0	0
0.5	1	1	1	1	1	0
1	2	1	1	1	1	1
1.5	2	2	1	2	1	1
2	2	2	2	2	1	1
3	2	2	2	2	1	1
4	2	2	2	2	1	1
5	2	2	2	2	1	2

Table 25 shows the results of the LVOV monitor located in the eVDD domain. Here there is neither a stark polarity dependence nor a strong proximity dependence. Positive zaps to the low-voltage IO do result in much greater over-voltage on the eVDD domain than do zaps to the 3.3 V IOs.

Table 25: IEC gun testing of the LVOV monitor in the eVDD domain.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	2	2	2	2	2	-
-4	2	2	2	2	2	2
-3	2	2	2	2	2	2
-2	1	2	1	1	1	1
-1.5	1	1	1	1	1	1
-1	1	1	1	0	1	1
-0.5	0	1	0	0	0	0
0	0	0	0	0	0	0
0.5	0	0	0	0	0	2
1	1	1	1	0	0	2
1.5	1	1	1	1	1	2
2	2	2	1	1	1	2
3	2	2	2	2	2	2
4	2	2	2	2	2	2
5	2	2	2	2	2	2

Table 26 shows the results reported by the LVOV monitor in the iVDD1 domain. Results here are similar to the results found for the eVDD domain. This is not surprising since the iVDD supply tracks the eVDD supply closely.

Table 26: IEC gun testing of the LVOV monitor in the iVDD1 domain.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	2	2	2	2	2	2
-4	2	2	2	2	2	2
-3	2	2	2	2	2	2
-2	2	2	2	2	2	2
-1.5	1	1	1	2	2	2
-1	1	1	1	1	1	1
-0.5	0	1	1	1	1	1
0	0	0	0	0	0	0
0.5	1	1	1	1	1	1
1	1	1	1	1	1	1
1.5	1	2	2	2	1	2
2	2	2	2	2	2	2
3	2	2	2	2	2	2
4	2	2	2	2	2	2
5	2	2	2	2	2	2

In summary, reported voltage monitor output levels track the applied stress level in all cases. It was observed that a single zap will produce both an under-voltage and an over-voltage,

suggesting that the power supplies are ringing, a conjecture that is consistent with circuit simulation results [19]. Further, applying a zap to an IO in the 3.3 V [eVDD] domain will generate supply noise in the eVDD [3.3 V] domain, but will create much more supply noise in the 3.3 V [eVDD] domain. A polarity dependency was observed in which positive zaps generated larger supply voltage excursions than negative zaps in almost all cases. Additionally, the amount of generated supply noise will differ depending on where zaps are applied. Zaps applied very near a voltage monitor will cause that monitor to report more noise. Zaps to more distant IOs may be affected by other characteristics.

5.3.2 Static Latch Upsets

Static latches are included in each monitor bank. Some latches are balanced while others have been tailored to favor a high or low output. Further, some of the balanced latches have inputs driven by an off-chip buffer. Unfortunately, due to an undersized buffer chain, the latches cannot be programmed or reset by using control signals. Thus the state the latches assume at power up is the only state in which the latches can be tested. The balanced latches with driven inputs could not be tested since input glitches could not cause data erasure. The unbalanced latches favoring a low output default to output low and thus cannot be tested as they are already in their preferred state. Luckily, the unbalanced latches favoring output high default to output low upon power-up because, with all nodes starting low (while power is off), the contention caused by the feedback inverter within the latch is not present. These latches are still susceptible to upset. Following upsets, these latches can be reset by power-cycling the test system. Alongside the two sets of unbalanced latches there are also balanced latches with inputs tied high or low in the same fashion as the unbalanced latches they are located with.

The unbalanced latches (and corresponding balanced latch) are labeled 1 through 5. Latch 1 refers to the latch that is most unbalanced and tailored to have the smallest noise margin. This latch should upset with the lowest supply noise. As the latch number increases, the degree of unbalance decreases. The balanced latches are labeled 5, but were not observed to upset with the stress levels applied to the system. When a latch upsets, it is expected that all latches with more unbalance will also upset. For example, if latch 3 upsets, latch 2 and latch 1 should also upset. This is in fact the case found in all of the collected data. The largest numbered latch that upsets for at least 40% of the zaps are reported, as seen in Equation (3), similar to how the voltage monitor data are reported above. Here, Z_L refers to the number of zaps resulting in and upset of latch number L and Z_T refers to the total number of zaps. For example, and entry of 2 means that latch 2 upset for at least 40% of the applied zaps at that level, but latches 3, 4 and 5 did not.

$$\max(L) \;\Big|\; Z_L \;\geq 0.4 \cdot Z_T \tag{3}$$

Table 27 and Table 28 show the latch upsets in monitor banks 1 and 2 respectively. These two sets of latches are identical except for their location on the test chip (see Figure 20). Due to the large noise margin of static logic, the latches were not observed to upset for lower stress levels. In both sets of latches, a polarity dependence is clearly observed which corroborates the results reported by the voltage monitors. Also, the only occurrence of upset for negative zaps occurs for zaps applied to LV_ZAPIO1. This could indicate that more noise is generated on the eVDD domain for zaps to pad cells in the eVDD domain, which is consistent with the voltage monitor results. Comparing both tables reveals that the location of the latches within the test chip does not result in a systematic difference. Zaps to the distant IOHV5 pad did not result in latch upsets indicating that there is again a positional dependence.

Table 27: Unbalanced latches upsets within monitor bank 1.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	0	0	0	0	0	-
-4	0	0	0	0	0	3
-3	0	0	0	0	0	0
-2	0	0	0	0	0	0
-1.5	0	0	0	0	0	0
-1	0	0	0	0	0	0
-0.5	0	0	0	0	0	0
0	0	0	0	0	0	0
0.5	0	0	0	0	0	0
1	0	0	0	0	0	0
1.5	0	0	0	0	0	0
2	0	0	0	0	0	0
3	3	2	0	0	0	1
4	3	3	1	0	3	2
5	3	3	2	0	4	2

Table 28: Unbalanced latches upsets within monitor bank 2.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	0	0	0	0	0	-
-4	0	0	0	0	0	4
-3	0	0	0	0	0	0
-2	0	0	0	0	0	0
-1.5	0	0	0	0	0	0
-1	0	0	0	0	0	0
-0.5	0	0	0	0	0	0
0	0	0	0	0	0	0
0.5	0	0	0	0	0	0
1	0	0	0	0	0	0
1.5	0	0	0	0	0	0
2	1	0	0	0	0	0
3	3	3	0	0	2	3
4	3	3	2	0	4	3
5	4	3	3	0	4	3

Figure 41 shows the combined upsets for unbalanced latches in both monitor bank 1 and monitor bank 2 plotted against the level of under-voltage reported by the eVDD LVUV monitor. Recall that the larger-numbered latches are designed to have larger noise margins, and thus will be harder to upset. The bars indicate the number of zaps resulting in the upset (or lack of upset) of a particular latch at a particular LVUV level. The incidence of latch upsets increases when the

under-voltage reaches level 3. For the 4 ns pulse testing of the LVUV monitor test structure (refer to Table 17), level 3 corresponds to eVDD < 0.34 V. This suggests that the transient noise margin of these latches is extremely large. However, it also indicates that ESD stress can generate extremely large voltage excursions on the various power supplies.

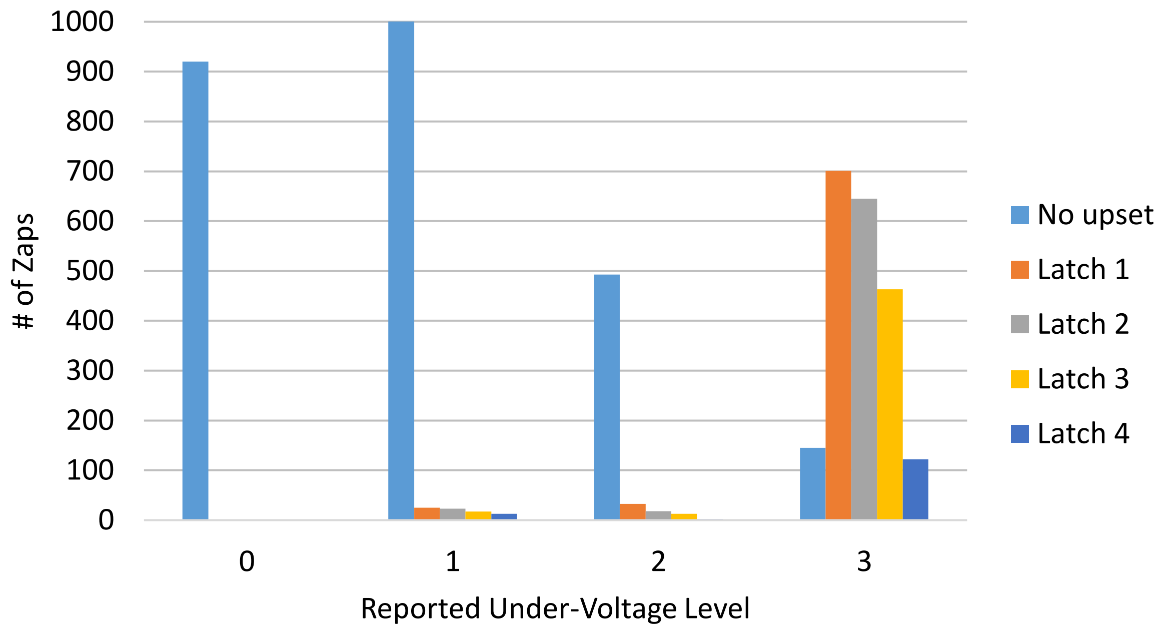

Figure 41: Unbalanced latch upsets compared with under-voltage observed on the eVDD supply. The bars indicate the number of zaps resulting in that particular latch upset for a given under-voltage level.

Latch upsets that occurred within the iVDD1 domain are reported in Table 29. Compared to the latch upsets in the eVDD domain, here upsets occurred within latches with higher noise margins and at lower stress levels. This confirms that the ESD stress causes more supply noise within the internally generated domain.

Table 29: Unbalanced latches upsets within monitor bank 3. These latches are powered from the internally generated iVDD1 supply.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	4	1	3	0	4	-
-4	4	1	2	0	4	2
-3	2	0	0	0	1	1
-2	0	0	0	0	0	0
-1.5	0	0	0	0	0	0
-1	0	0	0	0	0	0
-0.5	0	0	0	0	0	0
0	0	0	0	0	0	0
0.5	0	0	0	0	0	0
1	0	0	0	0	0	0
1.5	0	0	0	0	0	2
2	0	0	2	0	1	2
3	1	4	4	2	3	2
4	4	4	4	2	3	4
5	4	4	4	3	4	4

5.3.3 Bit Flips Due to Input Glitches

Various logic upsets observed within the test system can be ascribed to glitches at various IOs. Logic upsets within the dynamic OR gates and shift-registers fall within this category. These logic elements have inputs that ultimately originate off-chip and are generated by external buffer ICs. Some dynamic OR gates have all of their inputs hard-tied to VSS. These OR gates did not experience upsets for any zaps applied to the test system. The programmable dynamic OR gates located in monitor bank 1 have two 3.3 V inputs, DIN33 and STROBE33. Table 30 shows the upset percentages recorded for these OR gates. The percentage indicates the number of zaps resulting in an upset out of the number of applied zaps. Since upsets only occur within the programmable OR gates, these upsets are likely the result of input glitches. Glitches on the DIN33 pin would cause the dynamic OR gate to evaluate high resulting in an upset. Zaps to ZAPIO1 and ZAPIO2 result in upsets at lower stress levels despite ZAPIO3 being slightly closer to DIN33. To further test whether glitches are responsible for dynamic OR gate upsets, a pin adjacent to DIN33

pin was zapped. Zaps to this pin cause upsets for almost every zap at all stress levels. This finding is consistent with the glitch-based upset mechanism. It is possible that for zaps applied at an increased distance from the victim IO, the proximity effect is reduced to the point where other influences become dominant. For example, ZAPIO3 is in close proximity to a VSS33 pad cell which contains a large amount of decoupling capacitance as well as APD. When zaps are applied to ZAPIO3, current can flow to VSS33 through the nearby decap and then to VSS through the APD. Current can also flow to eVDD through decap located in the nearby VSS pad cell. It is possible that the availability of these many current paths will allow current to exit (or enter) the chip without generating voltage drops as large as those that occur for other IOs. The voltage monitor results are consistent with this conjecture. Figure 17 demonstrated that supply bounce is a factor in generating glitches; therefore, reducing the supply bounce may result in fewer observed glitches. Another possible factor is that DIN33 and ZAPIO2 are both nearby the same pair of VDD33 and VSS33 power pins. Perhaps as ZAPIO2 is stressed, the ESD current will flow to these nearby power pins and result in noise generated at the nearby DIN33 pin.

Table 30: Upset percentages for dynamic OR gates located in monitor bank 1. In addition to the dedicated zap pins that have previously been stressed, a DD protected 3.3 V IO directly adjacent to the DIN33 pin was also stressed.

V_{pre} (kV)	Zapped IO						
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1	Adjacent to DIN33
-5	12%	28%	4%	4%	67%	-	-
-4	32%	68%	16%	24%	40%	40%	-
-3	44%	60%	28%	0%	100%	10%	-
-2	30%	67%	0%	3%	27%	15%	100%
-1.5	0%	17%	0%	0%	0%	0%	100%
-1	0%	0%	0%	0%	0%	0%	100%
-0.5	0%	0%	0%	0%	0%	0%	100%
0	0%	0%	0%	0%	0%	0%	0%
0.5	0%	0%	0%	0%	0%	0%	80%
1	0%	0%	0%	0%	0%	0%	100%
1.5	0%	0%	0%	0%	0%	0%	100%
2	0%	0%	3%	0%	9%	0%	100%
3	57%	100%	0%	0%	33%	0%	-
4	14%	64%	0%	0%	33%	0%	-
5	43%	28%	0%	32%	7%	10%	-

The data in Table 30 does not exhibit a monotonic trend in which larger stress levels produce more upsets. This is not unexpected. The recorded upsets to dynamic OR gates on the previous test chip (see Table 11) also exhibited a window of stress levels that resulted in more upsets. Thus for the data presented in Table 30, the percentage of upsets decreases at higher precharge voltages because at these higher stress levels glitches appear on STROBE33 after appearing on DIN33. These glitches on STROBE33 cause the dynamic node to precharge and any upsets are eliminated.

To explore these dynamic upsets further, they can be compared with the glitch detector results. The glitch detector is located adjacent to ZAPIO2. Table 31 shows the results reported by the glitch detector circuit. These data were collected with the glitch detector's signal line driven low by an off-chip buffer IC. As expected, zaps to ZAPIO2 yield glitches at the lowest stress levels due to its close proximity. ZAPIO1 is slightly further away thus glitches begin at higher stress levels, around 1.5 kV. By comparison, the data in Table 30 suggest that the DIN33 pad cell begins to experience glitches around ±2 kV. However, ZAPIO3 and IOHV5 are both much farther away and thus zaps to these locations cause glitches only at high stress levels. When compared to IOHV5, ZAPIO3 (which is closer to the glitch detector line) appears to generate fewer glitches than expected. This finding is consistent with the dynamic OR gate upsets above. There is little difference between the DD protected IOs and the SCR protected IO and between low-voltage and 3.3 V IOs. Unfortunately, due to race conditions, while the glitch detector's signal line is driven high, the act of reading the glitch detector's output causes a glitch to be detected, preventing the glitch detector from being tested in this configuration.

Table 31: Glitches detected by the 3.3 V glitch detector. The glitch-detector's signal line was driven low.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	100%	100%	100%	35%	100%	-
-4	100%	100%	15%	0%	100%	100%
-3	100%	100%	0%	0%	100%	40%
-2	8%	100%	0%	0%	96%	33%
-1.5	0%	100%	0%	0%	0%	0%
-1	0%	100%	0%	0%	0%	0%
-0.5	0%	100%	0%	0%	0%	0%
0	0%	0%	0%	0%	0%	0%
0.5	4%	88%	0%	0%	0%	0%
1	0%	96%	0%	0%	0%	0%
1.5	8%	100%	0%	0%	44%	45%
2	83%	100%	0%	0%	90%	100%
3	100%	100%	0%	90%	100%	100%
4	100%	100%	0%	100%	100%	100%
5	80%	100%	0%	100%	100%	100%

Another dynamic OR gate with inputs generated from 3.3 V IOs is included on the test chip. This OR gate is located within monitor bank 2, but is programmed with DINX and STROBEX. The DINX and STROBEX inputs have 3.3 V receivers and OREDs. Table 32 shows the percentage of zaps resulting in bit-flips within this dynamic OR gate. Once again, despite being closer to the DINX pad cell, discharges to ZAPIO3 generated fewer upsets within the dynamic OR gate than ZAPIO1 and ZAPIO2. Unlike the other dynamic OR gate upsets, the data in Table 32 suggest that upset occurrence increases monotonically with stress level. The difference may arise from a difference in the strobe signal. The former dynamic OR gates (Table 30) use an active-high signal, STROBE33 (gates controlled by STROBE33 are precharged when STROBE33 is high). The additional OR gate (Table 30) uses an active-low signal, STROBEX. In both cases upsets begin at similar levels, consistent with the glitch detector. This is expected as DIN33 and DINX are both active-high and located adjacent to one another.

Table 32: Upset percentages for the dynamic OR gate located in monitor bank 2 with inputs DINX and STROBEX. The pad cells for these inputs contain ORED circuits.

Vpre (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	80%	100%	100%	100%	100%	-
-4	100%	100%	100%	100%	100%	100%
-3	100%	100%	100%	100%	100%	100%
-2	100%	100%	0%	40%	100%	100%
-1.5	60%	80%	0%	0%	60%	100%
-1	0%	0%	0%	0%	0%	40%
-0.5	0%	0%	0%	0%	0%	0%
0	0%	0%	0%	0%	0%	0%
0.5	0%	0%	0%	0%	0%	0%
1	0%	0%	0%	0%	40%	0%
1.5	100%	100%	20%	0%	100%	0%
2	100%	100%	20%	60%	100%	0%
3	100%	100%	100%	100%	100%	20%
4	100%	100%	100%	100%	100%	10%
5	100%	100%	100%	100%	100%	0%

Since DINX and STROBEX also contain OREDs, the bit-flips within the OR gates can be compared directly to the ORED results. As mentioned previously, the STROBEX signal line must be driven high to put the gate into evaluate mode. DINX must remain low to avoid erasing the dynamic node. Table 33 reports the percentage of zaps that were able to drive the DINX signal above VDD33 (ORED-high). Comparing these ORED results to the OR gate upsets reveals very close agreement. Bit-flips within the OR gate occur at slightly lower stress levels than are required to achieve an ORED-high output. This is unsurprising as the signal level need not be driven all the way above VDD33 to register as a logic high by the CMOS logic. Figure 42 shows the correlation between dynamic OR gate upsets and the ORED-high results for the DINX pad cell for 2 kV zaps applied to all 3.3 V IOs. When ORED-high is detected, 100% of the applied zaps result in an upset to the dynamic OR gate. When ORED-high is not detected, only 20% of zaps resulted in an upset of the OR gate.

Table 33: ORED-high results from the DINX pad cell. DINX was driven low during these tests. The ORED-high signal indicates that the DINX pad voltage was elevated above VDD33 during the zap.

Vpre (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	96%	100%	100%	80%	100%	-
-4	100%	100%	20%	76%	100%	50%
-3	96%	100%	0%	0%	100%	0%
-2	13%	30%	0%	0%	13%	0%
-1.5	0%	0%	0%	0%	0%	0%
-1	0%	0%	0%	0%	0%	0%
-0.5	0%	0%	0%	0%	0%	0%
0	0%	0%	0%	0%	0%	0%
0.5	0%	0%	0%	0%	0%	0%
1	0%	0%	0%	0%	3%	0%
1.5	0%	50%	0%	0%	0%	0%
2	90%	93%	0%	0%	60%	0%
3	100%	100%	0%	100%	100%	75%
4	100%	100%	44%	100%	100%	100%
5	100%	100%	56%	100%	100%	100%

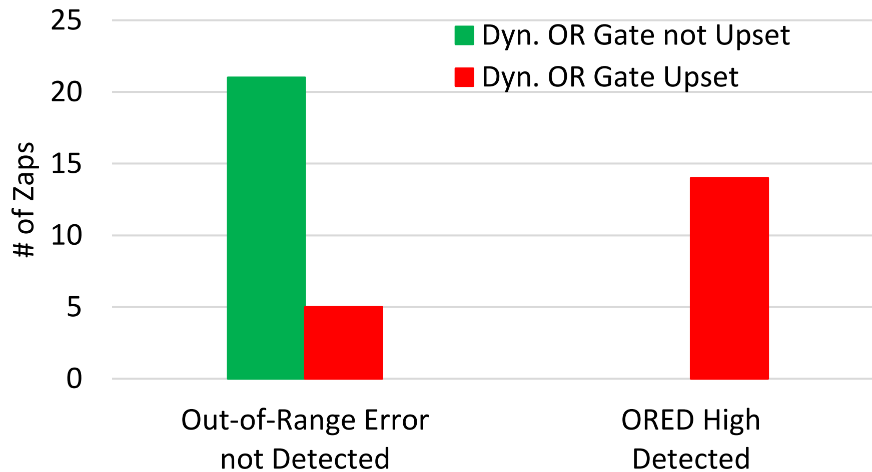

Figure 42: Correlation between ORED-high on the DINX pad cell and upsets within the dynamic OR gate.

The ORED can also detect when the DINX signal line is pulled below VSS33 (ORED-low). Table 34 shows that the DINX falls below VSS33 for almost all zaps. Since DINX was initially driven low, a much smaller voltage deviation is required to produce an ORED-low detection compared to ORED-high. The results from the ORED within the STROBEX pad cell show very similar results to these. STROBEX is initially driven high and thus almost all zaps

result in ORED-high. ORED-low results resemble Table 33 as the STROBEX signal level must be pulled from logic-high to below VSS33.

Table 34: ORED-low results from the DINX pad cell. DINX was driven low during these tests. The ORED-low signal indicates that the DINX pad voltage was pulled below VSS33 during the zap.

Vpre (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	96%	100%	100%	100%	100%	100%
-4	100%	100%	100%	100%	100%	100%
-3	100%	100%	100%	100%	100%	100%
-2	100%	100%	100%	100%	100%	100%
-1.5	100%	100%	100%	100%	100%	100%
-1	97%	100%	100%	100%	97%	100%
-0.5	67%	90%	74%	63%	77%	50%
0	0%	0%	0%	0%	0%	0%
0.5	100%	100%	87%	100%	100%	53%
1	100%	100%	100%	100%	100%	100%
1.5	100%	100%	100%	100%	100%	96%
2	100%	100%	100%	100%	100%	100%
3	100%	100%	100%	100%	100%	100%
4	100%	100%	100%	100%	100%	100%
5	100%	100%	100%	100%	100%	100%

Table 35 shows the percentage of zaps resulting in upsets within the dynamic OR gates located in monitor bank 2. Upsets occur for almost all zaps. Glitches on the DIN signal line would result in data upset within these dynamic OR gates. Table 36 shows the occurrence of glitches detected by the low-voltage glitch detector. Comparing these two tables reveals the OR gate upsets correspond with low-voltage IO glitches. While the low-voltage glitch detector is not monitoring the DIN signal line, it is clear that glitches on these low-voltage receivers occur for zaps to nearby and distant IOs alike. Other than the supply voltage, the main difference between the 3.3 V IOs and the low-voltage IOs is that the 3.3 V IOs have a Schmitt trigger incorporated, while the low-voltage IOs do not. In some cases, performance requirements could prohibit using low-voltage receivers with hysteresis. Without the Schmitt trigger, these IOs will be extremely vulnerable to glitches regardless of their location relative to the zapped IO.

Table 35: Upset percentages for dynamic OR gates located in monitor bank 2.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	96%	100%	100%	100%	100%	-
-4	100%	100%	100%	100%	100%	100%
-3	100%	100%	100%	100%	100%	100%
-2	100%	100%	100%	100%	100%	100%
-1.5	100%	100%	100%	100%	100%	100%
-1	100%	100%	100%	100%	100%	95%
-0.5	87%	90%	52%	90%	93%	90%
0	0%	0%	0%	0%	0%	0%
0.5	97%	100%	77%	70%	53%	83%
1	97%	97%	100%	100%	97%	93%
1.5	100%	100%	100%	100%	100%	94%
2	100%	100%	100%	100%	100%	100%
3	100%	100%	100%	100%	100%	100%
4	100%	100%	100%	100%	100%	100%
5	100%	100%	100%	100%	100%	100%

Table 36: Glitches detected by the low-voltage glitch detector located in monitor bank 2. The glitch detector's signal line is driven low.

V_{pre} (kV)	Zapped IO					
	ZAPIO1	ZAPIO2	ZAPIO3	IOHV5	ZAPIO_SCR1	LV_ZAPIO1
-5	80%	100%	100%	100%	100%	-
-4	100%	100%	100%	100%	100%	100%
-3	100%	100%	100%	100%	100%	100%
-2	100%	100%	100%	100%	100%	100%
-1.5	100%	100%	100%	100%	100%	100%
-1	100%	100%	67%	100%	100%	100%
-0.5	100%	100%	17%	0%	100%	100%
0	0%	0%	0%	0%	0%	0%
0.5	100%	100%	0%	0%	100%	100%
1	100%	100%	0%	80%	100%	100%
1.5	100%	100%	100%	100%	100%	91%
2	100%	100%	100%	100%	100%	100%
3	100%	100%	100%	100%	100%	100%
4	100%	100%	100%	100%	100%	100%
5	100%	100%	100%	100%	100%	100%

Due to the complexity of the test procedure, the shift-registers within the three monitor banks were tested for a small number of zaps. Table 37 reports the percentage of zaps resulting in data upset within the shift-registers. Five zaps were applied to ZAPIO1 for each test voltage indicated. All of the shift-registers were programmed to "1010" before the zap was applied.

Following the zap, the data stored within the shift-registers are read out from the test system. The shift-registers located in monitor bank 1 are programmed via the 3.3 V inputs SR_DIN33, STROBE33, and SR_LE33. The shift-registers located in monitor bank 2 and 3 are programmed via the low-voltage inputs SR_DIN, STROBE, and SR_LE. Similar to the dynamic OR gate results, upsets within the shift-registers located in monitor banks 2 and 3 upset at lower stress levels than those located in monitor bank 1. This is also consistent with the low-voltage and 3.3 V glitch-detector results.

Table 37: Percentage of zaps resulting in bit flip data in the shift-registers within each monitor bank.

V_{pre} (kV)	Mon. Bank 1	Mon. Bank 2	Mon. Bank 3
-5	100%	-	-
-4	100%	-	-
-3	100%	-	-
-2	15%	100%	100%
-1.5	0%	100%	100%
-1	0%	20%	20%
-0.5	0%	0%	0%
0	0%	0%	0%
0.5	0%	0%	0%
1	0%	80%	100%
1.5	0%	100%	100%
2	0%	-	-
3	100%	-	-
4	100%	-	-
5	100%	-	-

There are four topologies (“a” through “d”) of shift-registers used within each monitor bank. These are described by Figure 37. Table 38 reports the pre- and post-zap data (MSB first) stored within the shift-registers for a subset of the zaps presented in Table 37. Shift-registers with the topology shown in Figure 37(a) have dedicated output registers, the code stored within these output registers is shown separately for the applicable shift-registers. The data clearly indicate that for each zap, all shift-registers within the same monitor bank contain identical post-zap codes despite differences in circuit topology. This trend was verified across all data, not only this subset.

Table 38: Subset of shift-register bit flips. Two zaps are shown for four zap voltage levels. "SR topology" refers to the style of shift-register as described in Figure 37. Only style "a" has dedicated output registers. The code stored in these output registers is indicated in the table for these registers. All shift-registers were programmed to the same starting code prior to zapping.

V_{pre} (kV)	SR topology	Monitor Bank 1			Monitor Bank 2			Monitor Bank 3		
		Starting Code	Post-Zap Code	Post-Zap Output Registers	Starting Code	Post-Zap Code	Post-Zap Output Registers	Starting Code	Post-Zap Code	Post-Zap Output Registers
1	a	1 0 1 0	1 0 1 0	1 0 1 0	1 0 1 0	1 0 1 1	0 1 0 1	1 0 1 0	1 0 1 1	0 1 0 1
1	b	1 0 1 0	1 0 1 0		1 0 1 0	1 0 1 1		1 0 1 0	1 0 1 1	
1	c	1 0 1 0	1 0 1 0		1 0 1 0	1 0 1 1		1 0 1 0	1 0 1 1	
1	d	1 0 1 0	1 0 1 0		1 0 1 0	1 0 1 1		1 0 1 0	1 0 1 1	
1	a	1 0 1 0	1 0 1 0	1 0 1 0	1 0 1 0	1 0 0 1	0 1 0 0	1 0 1 0	1 0 0 0	0 1 0 0
1	b	1 0 1 0	1 0 1 0		1 0 1 0	1 0 0 1		1 0 1 0	1 0 0 0	
1	c	1 0 1 0	1 0 1 0		1 0 1 0	1 0 0 1		1 0 1 0	1 0 0 0	
1	d	1 0 1 0	1 0 1 0		1 0 1 0	1 0 0 1		1 0 1 0	1 0 0 0	
1.5	a	1 0 1 0	1 0 1 0	1 0 1 0	1 0 1 0	1 0 1 1	1 0 1 1	1 0 1 0	1 0 0 1	1 0 0 1
1.5	b	1 0 1 0	1 0 1 0		1 0 1 0	1 0 1 1		1 0 1 0	1 0 0 1	
1.5	c	1 0 1 0	1 0 1 0		1 0 1 0	1 0 1 1		1 0 1 0	1 0 0 1	
1.5	d	1 0 1 0	1 0 1 0		1 0 1 0	1 0 1 1		1 0 1 0	1 0 0 1	
1.5	a	1 0 1 0	1 0 1 0	1 0 1 0	1 0 1 0	1 0 0 1	0 1 0 0	1 0 1 0	1 0 0 1	0 1 0 0
1.5	b	1 0 1 0	1 0 1 0		1 0 1 0	1 0 0 1		1 0 1 0	1 0 0 1	
1.5	c	1 0 1 0	1 0 1 0		1 0 1 0	1 0 0 1		1 0 1 0	1 0 0 1	
1.5	d	1 0 1 0	1 0 1 0		1 0 1 0	1 0 0 1		1 0 1 0	1 0 0 1	
2	a	1 0 1 0	1 0 1 0	1 0 1 0	1 0 1 0	1 1 0 1	1 1 1 1	1 0 1 0	1 0 0 0	1 1 1 0
2	b	1 0 1 0	1 0 1 0		1 0 1 0	1 1 0 1		1 0 1 0	1 0 0 0	
2	c	1 0 1 0	1 0 1 0		1 0 1 0	1 1 0 1		1 0 1 0	1 0 0 0	
2	d	1 0 1 0	1 0 1 0		1 0 1 0	1 1 0 1		1 0 1 0	1 0 0 0	
2	a	1 0 1 0	1 0 1 0	1 0 1 0	1 0 1 0	1 1 0 1	1 1 1 0	1 0 1 0	1 1 0 0	1 1 1 0
2	b	1 0 1 0	1 0 1 0		1 0 1 0	1 1 0 1		1 0 1 0	1 1 0 0	
2	c	1 0 1 0	1 0 1 0		1 0 1 0	1 1 0 1		1 0 1 0	1 1 0 0	
2	d	1 0 1 0	1 0 1 0		1 0 1 0	1 1 0 1		1 0 1 0	1 1 0 0	
3	a	1 0 1 0	1 0 0 0	1 0 1 0						
3	b	1 0 1 0	1 0 0 0							
3	c	1 0 1 0	1 0 0 0							
3	d	1 0 1 0	1 0 0 0							
3	a	1 0 1 0	1 0 0 0	1 0 1 0						
3	b	1 0 1 0	1 0 0 0							
3	c	1 0 1 0	1 0 0 0							
3	d	1 0 1 0	1 0 0 0							

This suggests that the upsets seen here are the result of glitches at the signal lines and not due to any random isolated bit flips within the shift-registers or output registers. Further, since the shift-register shown in Figure 37(a) is identical (except for length) to the shift-register used in the previous test system, the cause for those upsets should be similar. The susceptibility of the shift-registers in monitor banks 2 and 3 (and those in the previous test system) can again be ascribed to

the lack of a Schmitt trigger within the receiver and the reduced input receiver supply voltage. The contents of the shift-registers contain both ones and zeros following an upset. The position of the ones and zeros reveal that both spurious ones and zeros are written into the shift-registers. For example, an entry of 1001 could only result if at minimum a zero was shifted into the register followed by a one. Further, the data stored within the output registers often differ from the data stored within the shift-register. This indicates that the three control signals experience glitches independently. This is expected as they are located on three different sides of the test chip.

5.3.4 Effects of Reduced Supply Voltage

As semiconductor technologies advance, supply voltages continue to shrink. With a reduced supply voltage, there is a possibility that smaller noise margins will result in an increase susceptibility to soft failures. The test chip designed for this work can be powered by either a 1.5 V core supply (eVDD) or a 1.2 V core supply. Recall that the internally generated domain, iVDD1, is referenced to eVDD, and thus when the supply voltage is changed, iVDD1 will track it. Five zaps were applied to ZAPIO1 for each stress level tested, with the core supply voltage of the test chip set to 1.2 V. The test was repeated with the supply voltage set to 1.5 V. A comparison of the voltage monitor results from both experiments are shown in Table 39. The supply voltage is indicated by the voltage in the column header. Since the under-voltage monitors' reference levels are provided by a resistor chain in the eVDD (or iVDD) domain, the reference levels will scale with the supply voltage. However, the over-voltage monitors use a resistor chain in the 3.3 V domain. Thus, when eVDD is varied, the reference levels will not change. There is not a large difference observed between the relative under-voltage monitor readings. This is expected since the ADC reference levels are scaled along with the supply voltage. However, the LVOV monitor results do vary greatly because the reference levels are constant despite a reduction in supply

voltage. A comparison of the unbalanced latch results are shown in Table 40. The latch upset data do not show a strong dependence on supply voltage. This is consistent with Table 39 as unbalanced latch upsets were shown to correlate with under-voltage which did not vary with supply voltage.

Table 39: Comparison of voltage monitor results for zaps to the system with eVDD powered from a 1.5 V source and a 1.2 V source. Zaps were applied to ZAPIO1 (3.3 V IO).

V_{PRE} (kV)	HVUV		eVDD LVUV		iVDD LVUV		HVOV		eVDD LVOV		iVDD LVOV	
	1.5 V	1.2 V	1.5 V	1.2 V	1.5 V	1.2 V	1.5 V	1.2 V	1.5 V	1.2 V	1.5 V	1.2 V
0	0	0	0	0	0	0	0	0	0	0	0	0
0.25	0	0	0	0	1	1	1	1	0	0	0	0
0.5	0	0	0	0	1	1	1	1	0	0	1	0
0.75	0	0	1	1	1	1	2	1	1	0	1	0
1	0	1	1	1	1	1	2	2	1	0	1	0
1.5	1	1	2	2	1	1	2	2	1	0	1	1
2	1	1	3	3	1	2	2	2	1	0	2	1
2.5	2	2	3	3	2	2	2	2	2	0	2	1
3	2	2	3	3	3	3	2	2	2	2	2	1
4	3	3	3	3	3	3	2	2	2	2	2	2
5	3	3	3	3	3	3	2	2	2	2	2	2

Table 40: Comparison of unbalanced latch results for zaps to the system with eVDD powered from a 1.5 V source and a 1.2 V source. Zaps were applied to ZAPIO1 (3.3 V IO).

V_{PRE} (kV)	Monitor Bank 1		Monitor Bank 2		Monitor Bank 3	
	1.5	1.2	1.5	1.2	1.5	1.2
0	0	0	0	0	0	0
0.25	0	0	0	0	0	0
0.5	0	0	0	0	0	0
0.75	0	0	0	0	0	0
1	0	0	0	0	0	0
1.5	0	0	0	0	0	0
2	0	2	1	2	0	1
2.5	3	3	4	3	0	0
3	3	3	4	3	1	0
4	3	3	4	4	4	4
5	3	3	4	4	4	4

Table 41 shows the results reported from the low-voltage glitch detector and the dynamic OR gates with low-voltage IOs. Due to the susceptibility of the low-voltage IOs to noise, glitches and upsets were observed for almost all applied stresses. Glitches and resulting upsets occur at a

lower stress level when the supply is reduced. This suggests that the reduced noise margin may be leading to increased noise susceptibility, however, this result is derived from a small dataset.

Table 41: Low-voltage glitch-detector and dynamic OR gate upset results for zaps to the system with eVDD powered from a 1.5 V source vs. a 1.2 V source. Zaps were applied to ZAPIO1 (3.3 V IO).

V_{PRE} (kV)	Low-Voltage Glitch Detector		Dynamic OR Gates with Low-voltage IOs (eVDD)	
	1.5	1.2	1.5	1.2
0	0	0	0	0
0.25	0	0.2	0	0.2
0.5	1	1	1	1
0.75	1	1	1	1
1	1	1	1	1
1.5	1	1	1	1
2	1	1	1	1
2.5	1	1	1	1
3	1	1	1	1
4	1	1	1	1
5	1	1	1	1

5.3.5 Zaps to a Protected System

Commercial systems must be designed to reduce to the occurrence of soft failure. One way to mitigate soft failures is to prevent signal lines from being exposed to stress. For example, a system enclosure may be designed to ensure signal lines cannot be contacted during routine handling. The IEC 61000-4-2 standard requires that any exposed conductors, such as connector shields, be stressed. Another way to mitigate soft failures would be to protect exposed signal lines through the use of a transient voltage suppressor (TVS). TVS devices are designed to turn on quickly at a voltage level above the nominal supply voltage. Once on, the device forms a low impedance between the signal line and the board ground and/or power planes. While some residual current will likely enter the signal pin, the majority of the ESD current will be shunted away from the protected component.

A bi-directional TVS was placed between the ZAPIO1 signal line and board ground near the zap location. The TVS is rated for ±30 kV IEC 61000-4-2 discharges with a reverse standoff voltage of 3.3 V. Zaps were then applied to the ZAPIO1 signal line. In a separate experiment,

zaps to a system within an enclosure were emulated. Zaps were applied to the USB connector shield which is well connected to board ground. The power noise detected by the supply voltage monitors as well as the unbalanced latches for zaps to both protected systems are shown in Table 42. These results are juxtaposed with the results obtained from zapping ZAPIO1 on an unprotected system. It is clear that on-chip power noise is greatly reduced when zaps are applied to the protected systems. All four supply voltage monitors show reduced levels of noise and the unbalanced latches did not show any upset for zaps to the TVS protected IO or for zaps to the USB connector. Further, preventing stress on internal signal pins appears to be more effective at reducing on-chip power noise than adding a TVS to an IO. However, despite the added protection, noise was still generated inside the IC at relatively low stress levels. This is striking as, especially for the USB shield zaps, the main current path from the discharge point to AC ground does not include the IC.

Table 42: Comparison of supply noise generated by zaps to a protected system compared to noise generated by zaps to an unprotected system. Two types of protected systems were considered. Data in columns marked “TVS” come from zaps to an IO protected with a TVS diode. Data in columns marked “USB Shield” come from zaps applied to the (board-grounded) USB connector shield. Data in columns marked “None” come from zaps to an unprotected system.

	HVUV			eVDD LVUV			HVOV			eVDD LVOV			Unbalanced Latches in eVDD Domain		
V_{pre} (kV)	None	TVS	USB Shield	None	TVS	USB Shield	None	TVS	USB Shield	None	TVS	USB Shield	None	TVS	USB Shield
-12	-	0	0	-	1	1	-	2	1	-	2	2	-	0	0
-10	-	0	0	-	1	1	-	2	1	-	2	2	-	0	0
-8	-	0	0	-	1	1	-	1	1	-	2	2	-	0	0
-6	-	0	0	-	1	1	-	1	1	-	2	1	-	0	0
-4	0	0	0	1	1	1	2	1	0	2	2	0	0	0	0
-2	0	0	0	1	0	0	1	1	0	1	1	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	2	0	0	3	1	1	2	1	0	2	1	1	1	0	0
4	3	0	0	3	1	1	2	2	0	2	1	1	3	0	0
6	-	0	0	-	1	1	-	2	1	-	2	2	-	0	0
8	-	0	0	-	1	1	-	2	1	-	2	2	-	0	0
10	-	0	0	-	1	1	-	2	1	-	2	2	-	0	0
12	-	0	1	-	1	1	-	2	1	-	2	2	-	0	0

Upsets within the dynamic OR gates were also considered for zaps to protected systems. These results are shown in Table 43. The two protection methodologies examined do not appear to offer as much benefit concerning IO noise (shown to be the primary mechanism for OR gate upset) as it did for power noise. Dynamic OR gate upsets occurred for relatively low stress levels, especially to gates with low-voltage IOs. This disparity results because on-chip power noise requires that ESD current flow into the IC, the protection methods applied to this system provide a great reduction in current flow into the IC. However, IO noise does not require ESD current to flow on chip. Board-level ESD current can cause the signal-lines potential to fluctuate, generating glitches, as can radiation from the ESD gun itself.

Table 43: Comparison of dynamic OR gate upsets generated by zaps to a protected system compared to upsets generated by zaps to an unprotected system. Two types of protected systems were considered. Data in columns marked "TVS" come from zaps to an IO protected with a TVS diode. Data in columns marked "USB Shield" come from zaps applied to the (board-grounded) USB connector shield. Data in columns marked "None" come from zaps to an unprotected system.

	Dyn. OR gates (3.3 V IOs)			Dyn. OR gates (LV IOs)		
V_{pre} (kV)	None	TVS	USB Shield	None	TVS	USB Shield
-12	-	100%	100%	-	100%	100%
-10	-	100%	80%	-	100%	93%
-8	-	100%	67%	-	100%	100%
-6	-	100%	67%	-	100%	100%
-4	100%	100%	0%	100%	100%	100%
-2	100%	0%	0%	100%	100%	100%
0	0%	0%	0%	0%	0%	0%
2	100%	0%	0%	100%	100%	100%
4	100%	0%	50%	100%	100%	100%
6	-	100%	100%	-	100%	100%
8	-	100%	100%	-	100%	100%
10	-	80%	100%	-	100%	100%
12	-	100%	100%	-	100%	100%

To determine whether OR gate upsets are the result of ESD gun radiation, an additional experiment was designed. In this experiment, gun zaps were applied to the HCP nearby the USB connector. In this manner the conducted stress into the test system would be diminished, but the radiated stress from the ESD gun would be largely the same. Table 44 shows the supply noise

monitor results along with the dynamic OR gate upsets. It is clear that power supply noise is the results of ESD current from the gun being injected into the test system. However, the dynamic OR gates upset similarly whether the USB connector shield is zapped or the HCP is zapped. Thus these glitch-based upsets are likely the result of ESD gun radiation rather than conducted stress into the test system.

Table 44: Supply noise and dynamic OR gate upsets resulting from zaps to the USB shield compared to those resulting from zaps to the HCP nearby the shield.

V_{pre} (kV)	HVUV		eLVUV		HVOV		eLVOV		Dyn. OR gate	
	USB	HCP	USB	HCP	USB	HCP	USB	HCP	USB	HCP
-12	0	0	1	1	1	0	2	1	1	1
-10	0	0	1	0	1	0	2	1	1	1
-8	0	0	1	0	1	0	2	0	1	1
-6	0	0	1	0	1	0	2	0	0.5	1
-4	0	0	0	0	0	0	1	0	0.2	0
-2	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0
4	0	0	1	0	0	0	1	0	0	0.5
6	0	0	1	0	1	0	2	0	1	0.1
8	0	0	1	1	1	0	2	1	1	0.8
10	0	0	1	1	1	0	2	1	1	1
12	0	0	1	1	1	1	2	1	1	1

CHAPTER 6:

FUTURE WORK

A primary goal of the new test system was to develop a method for quantitatively identifying supply noise during ESD injection and to relate logical upsets to supply noise. While that goal was met reasonably well, new questions arose. Further exploration is required, through the use of circuit simulation, to fully understand the cause of the stark polarity dependence observed in the reported under- and over-voltage data. Several design flaws within the voltage monitors and logic elements have been identified. These elements must be redesigned and fabricated to verify improved performance.

The three banks of monitors placed in the new system offer comparisons between three common power domains included on modern products: externally generated high-voltage domains, externally generated low-voltage domains, and internally generated low-voltage domains. Studying upsets associated with these different types of domains, within the confines of one test chip, has yielded results pertaining to relative susceptibilities each type of domain exhibits. However, further work must be done, through simulation, to understand why the reported level of under-voltage observed within the internally generated supply does not increase monotonically with the applied stress.

Glitches were observed to occur within standard 3.3 V IO cells when ESD was applied to other signals traces that did or did not terminate at the test chip. These results have been augmented by the addition of a glitch detector placed in a low-voltage IO cell. The data reported from the low-voltage glitch detector reveal that low-voltage IOs did experience many more glitches than 3.3 V IOs. However, the omission of a Schmitt trigger within the low-voltage receiver may lead to an unfair comparison between the low-voltage and 3.3 V glitch detector

results. A glitch detector should be placed on IOs including 3.3 V IOs with and without a Schmitt trigger and a low-voltage IOs with and without a Schmitt trigger. In this manner input glitches can be fully understood. Further, glitch-detector circuits could be placed on control-signal lines instead of dedicated glitch-detector IOs. This would allow upsets caused by glitches on signal lines to be compared directly to glitches detected on those signal lines.

Modifications to the overall system could also yield a deeper understanding of the causes of soft failure. One planned experiment involves placing the system in a metallic enclosure. Such an enclosure, if implemented correctly, would act as a Faraday cage and reduce the amount of energy coupled from the ESD gun to the internals of the system. Zaps would be applied to external connectors (e.g. the USB shield). This type of experiment would distinguish between upsets that result from ESD that is injected into the system and upsets that arise from noise produced within the gun.

REFERENCES

[1] ESD Association, "Fundamentals of electrostatic discharge, part one: An introduction to ESD," ESD Association, Rome, NY, 2013.

[2] *System Level ESD Part 1: Common Misconception and Recommended Basic Approaches,* JEDEC Standard JEP161, 2011.

[3] *Road Vehicles — Test Methods for Electrical Disturbance from Electrostatic Discharge,* ISO Standard 10605, 2011.

[4] *For Electrostatic Discharge Sensitivity Testing Human Body Model (HBM) - Device Level,* ANSI/ESDA/JEDEC Standard JS-001-2014.

[5] *For Electrostatic Discharge Sensitivity Testing Charge Device Model (CDM) - Device Level,* ANSI/ESDA/JEDEC Standard JS-002-2014.

[6] G. Muchaidze, J. Koo, Q. Cai, T. Li, L. Han, A. Martwick, K. Wang, J. Min, J. Drewniak and D. Pommerenke, "Susceptibility scanning as a failure analysis tool for system-level electrostatic discharge (ESD) problems," *IEEE Transactions on Electromagnetic Compatibility,* vol. 50, no. 2, pp. 268-276, 2008.

[7] *Electromagnetic Compatibility (EMC) -- Part 4-2: Testing and Measurement Discharge Immunity Test,* IEC Standard 61000-4-2, 2009.

[8] T. Schwingshackl, B. Orr, J. Willemen, W. Simbürger, H. Gossner, W. Bösch and D. Pommerenke, "Powered system-level conductive TLP probing method for ESD/EMI hard fail and soft fail threshold evaluation," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2013.

[9] D. Johnsson and H. Gossner, "Study of system ESD codesign of a realistic mobile board," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2011.

[10] P. Besse, F. Lafon, N. Monnerear, F. Caignet, J. P. Laine, A. Salles, S. Rigour, M. Bafleur, N. Nolhier and D. Trémouilles, "ESD system level characterization and modeling methods applied to a LIN tranceiver," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2011.

[11] P. Besse, K. Abouda and C. Abouda, "Identifying electrical mechanisms responsible for functional failures during harsh external ESD and EMC aggression," in *European Symposium on the Reliability of Electron Devices, Failure Physics and Analysis Proceedings*, 2011.

[12] F. Caignet, N. Nolhier, A. Wang and N. Mauran, "20GHz on-chip measurement of ESD waveform for system level analysis," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2013.

[13] N. Thomson, Y. Xiu, R. Mertens, M. S. Keel and E. Rosenbaum, "Custom test chip for system-level ESD investigations," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2014.

[14] R. Mertens, N. Thomson, Y. Xiu and E. Rosenbaum, "Theory of active clamp response to power-on ESD and implications for power supply integrity," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2014.

[15] M. Stockinger, Z. Wenzhong, K. Mason and J. Feddeler, "An active MOSFET rail clamp network for comonent and system level protection," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2013.

[16] M. S. Keel, N. Jack and E. Rosenbaum, "ESD-resilient active biasing scheme for high-speed SSTL I/Os," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2013.

[17] Y. Xiu, N. Thomson, R. Mertens and E. Rosenbaum, "A mechanism for logic upset induced by power-on ESD," in *Electrical Overstress/Electrostatic Discharge Symposium*, 2014.

[18] R. Merterns, "Understanding, modeling, and mitigating system-level ESD in integrated circuits," Ph.D. dissertation, Univ. of Illinois at Urbana-Champaign, 2015.

[19] Y. Xiu, Interviewee, *Private conversation.* [Interview].