

Precise Pulse Discrimination for Space-Based Timing Front Ends

By

Lydia Lee

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Kristofer S.J. Pister, Chair

Professor Vladimir Stojanovic

Professor Wenbin Lu

Summer 2023

Precise Pulse Discrimination for Space-Based Timing Front Ends

Copyright 2023
by
Lydia Lee

Abstract

Precise Pulse Discrimination for Space-Based Timing Front Ends

by

Lydia Lee

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Kristofer S.J. Pister, Chair

The conversion of trigger events to their digital equivalent is a central component of any timing-based front end, with applications found in mass spectrometry, single channel analyzers, and a huge variety of 3D mapping and ranging systems. At the same time, ever-tightening size, weight, and power budgets for space launches with a skyrocketing (no pun intended) number of launches in the last decade have made application-specific integrated circuit solutions increasingly appealing. However, conventional analog methods of pulse discrimination introduce timing walk or are limited to a narrow range of pulse shapes, while early-stage digitization requires impractically high sample rates for the events in question.

This work presents the analysis, design, and measurement of an integrated constant fraction discriminator with theoretically zero timing walk and a programmable, constant trigger fraction which does not depend on input pulse shape. The specific silicon presented here was designed for the Solar Probe Analyzer for Ions as part of its time-of-flight mass spectrometer to determine the ion composition of space plasmas. This dissertation discusses the front end requirements for a radiation hardened pulse discriminator in the context of SPAN-Ion. We then address the architectural modifications used to achieve a pulse shape-independent constant trigger fraction, as well as the analog and digital hardening techniques required to detect, correct, and/or mitigate radiation-induced effects. Finally, this work presents the first attempt at an integrated pulse-shaping front end for SPAN-Ion, concluding with simulation results from a more recent chip and a discussion of future work both for SPAN-Ion and for further code base development.

To my family and friends (pets included).

Contents

Contents	ii
List of Figures	v
List of Tables	xii
1 Introduction	1
1.1 The Solar Probe Analyzer for Ions	1
1.2 Radiation and Integrated Circuits	5
1.2.1 Effects	5
1.2.2 Radiation Hardening Electronics	14
1.2.3 Simulation and Testing	19
1.3 The Berkeley Analog Generator	23
1.4 Acronyms	25
2 Pulse Discrimination	26
2.1 Topology Overview	26
2.2 Timing Walk	29
2.3 Afterpulse Rejection	31
2.4 Constant Fraction Discrimination	32
2.5 SEE Watchdog	37
3 Chip V1	39
3.1 Chip Summary	39
3.2 Design and Measurements	41
3.2.1 Power and Reference Generation	41
3.2.2 No-Shape/Small Signal Chain	44
3.2.3 Full/Main Signal Chain	55
4 Chip V2	64
4.1 Chip Summary	64
4.2 Design and Measurements	65
4.2.1 No-Shape/Small Signal Chain	65

4.2.2	Full/Main Signal Chain	67
5	Conclusions and Future Work	72
5.1	Button-Press Design with BAG	73
	Bibliography	75
A	SCM3x	83
A.1	Digital LDO	83
A.2	ADC + PGA Sensor Interface	85
B	SCM3C Digital Flow Documentation	90
B.1	Code Base	90
B.2	Simulation	91
B.3	Scripts	94
B.4	Running the Flow	99
B.5	Known Idiosyncrasies	111
B.6	Signoff	112
C	Optical Receiver	115
C.1	Background	115
C.2	Chip Summary	116
C.2.1	I/O	118
C.2.2	Scan	120
C.3	Analog Front End	120
C.4	PWM Scheme	127
C.5	SPAD Test Structure	128
C.6	Code Base and Cadence Locations	130
D	BAG 2.0 Scripts	131
D.1	General Utilities	131
D.2	Scripts	131
E	Chip V1 Documentation	135
E.1	Infrastructure	135
E.2	IO	136
E.3	Scan Bits	138
E.4	Test Setup	139
E.5	Cadence Locations	143
E.6	Known Idiosyncrasies	144
F	Chip V2 Documentation	145
F.1	Infrastructure	145

F.2	IO	146
F.3	Scan Bits	148
F.4	Test Setup	148
F.5	Cadence Locations	150
F.6	Known Idiosyncrasies	150
G	Miscellaneous	152

List of Figures

1.1	[44] Block diagram of the SPAN-I sensor, including ESA, TOF, and individual components of the electronics box.	2
1.2	Operation of SPAN-I's time-of-flight mass spectrometer. (a) The ion is accelerated with a known potential to a speed defined by Equation 1.1. (b) Upon impact, the START carbon foil releases secondary electrons which are amplified then picked up by the START anode as a current pulse. (c) The ion travels through the START foil to the STOP foil, which is a fixed distance L away. This produces a STOP pulse by the same mechanism as the START pulse.	3
1.3	Time of flight calibration data versus theoretical time of flight with and without a $1.0\mu\text{g}/\text{cm}^2$ carbon foil. Each data point corresponds to the time of flight mode bin with 100ps resolution. Calibration data courtesy of SSL.	4
1.4	[49] An energy band diagram of a MOS structure for positive gate bias, indicating major physical processes underlying radiation response.	6
1.5	[21] Threshold voltage shift in commercial 130nm CMOS as a function of TID for different (a) NMOS and (b) PMOS transistor sizes, up to 136Mrads. PMOS threshold shift is absolute value. The last point refers to full annealing at 100°C	7
1.6	(a) [63] Ratio between the value of $1/f$ noise parameter K_f at 10Mrad absorbed dose of ^{60}Co γ -rays and the value before irradiation as a function of drain current I_D for NMOS and PMOS devices in the $0.13\mu\text{m}$ process. (b) [51] $1/f$ noise spectra S_v as a function of total dose. The device was under +6V bias during irradiation.	8
1.7	[35] $I_{DS}-V_{GS}$ characteristics before and after TID irradiation for a core NMOS device.	9
1.8	[76] Threshold voltage mismatch between identically designed pairs of a regular and b enclosed transistors before and after γ irradiation up to 100kGy.	10
1.9	[76] Standard deviation of the current factor mismatch between identically designed pairs of a regular and b enclosed transistors before and after γ irradiation up to 100kGy.	10
1.10	[7] Charge generation and collection phases in a reverse-biased junction and the resultant current pulse caused by the passage of a high-energy ion.	11

1.11	[46] (a) ITRS scaling of of first-level metal half pitch and progression into ultra-thin fully depleted silicon-on-insulator and multiple gate technologies, compared with a representative diameter of the ionized free charge distribution in silicon in the wake of a light ion interaction. Node indicators are approximate. (b) ITRS scaling of gate charge for a nominal 3X/6X NMOS/PMOS inverter.	13
1.12	Layers of radiation hardening.	14
1.13	Two variants of a MOSFET with $W_{\text{eff}} = 2W$. Purple is the gate with oxide—either for the channel or for isolation—underneath it, yellow is the source and drain of the device. The device with two fingers (a) has a greater total gate area with isolation oxide placed underneath it, due to the additional gate material required to electrically connect the two fingers.	15
1.14	(a) 3-to-1 majority voting and (b) 3-to-3 voting. 3-to-3 voting ensures the input of all Logic blocks B and C are correct in the event of one SET (assuming good layout practice).	17
1.15	The DICE memory cell, adapted from [13].	18
1.16	Leaky MOSFET used to emulate radiation-exacerbated leakage. The sign and to an extent the magnitude of I_{leak} depends on the voltage measured by the voltmeter, clipped at a fixed value. The leakage also scales with device aspect ratio.	20
1.17	Estimated charge collected per micron of collection depth for an SET produced by the ions available in the cocktails at Lawrence Berkeley National Laboratory [41].	21
1.18	SET emulator circuit. $\tau_r = R_r C_r$ and $\tau_f = R_f C_f$ for Equation 1.6. The input voltage source on the left provides a step function.	22
1.19	TID test flow. Adapted from [55].	23
1.20	The BAG design flow for a block.	24
2.1	MCP output, converted to a voltage and averaged over a periodic event trigger.	27
2.2	Block diagram of a common generalized CFD, implemented with two LTI operations $H_+(s)$ and $H_-(s)$ along with an ideal comparator.	28
2.3	Plot of timing walk given 2-20Me ⁻ pulses with $\tau_r = 500\text{ps}$, $\tau_f = 1\text{ns}$ to roughly match Figure 2.1, with a threshold set at $2 \times 3.3\text{V}/2^9 \approx 12.9\text{mV}$ for a transimpedance amplifier with a fixed gain-bandwidth product of 30GHz V/A. Increasing the transimpedance sees the timing walk asymptotically approach 611ps. This is the bleeding edge of what the process node can achieve under nominal operating conditions; decreasing the gain-bandwidth product to 10GHz V/A to account for variation in process, supply, and temperature makes it so even 1LSB of the 9-bit DAC is insufficient to meet walk requirements.	30
2.4	High level concept of CFD usage in parallel with an LED. Here, the LED determines if there is a pulse, and the CFD provides the timing of the pulse.	31

2.5	The response of a retriggerable one shot and a non-retriggerable one shot to (a) a single pulse with duration shorter than $t_{1\text{shot}}$ (b) two pulses whose rising edges occur within $t_{1\text{shot}}$ of each other (c) a pulse with duration equal to $t_{1\text{shot}}$ (d) a pulse with duration longer than $t_{1\text{shot}}$. Each one shot produces HIGH pulses in response to rising edges.	32
2.6	The core methods behind prior implementations of constant fraction discriminators.	33
2.7	Example inputs to the CFD comparator, with (b) and without (a) the peak detector inserted in the shaping chain. The orange line is the delayed input pulse, the blue line is the input pulse, attenuated by a factor $f = 0.5$	34
2.8	The modified CFD with the peak detector added before the attenuator.	34
2.9	Block diagram of the front end and its operation with the CFD branch outlined in blue and the LED branch outlined in yellow. The one shot output is used to reset the peak detector	36
2.10	Diagram with one possible scenario for SEE-induced lockout. In general, lockout can occur if an SET on the peak detector raises the output of the attenuator to a level that real pulses can never reach. The CFD output remains low, and the system never resets the peak detector until the chip is reconfigured.	37
2.11	The peak detector's SEE detection and correction watchdog circuit and operation in the event of an otherwise lock-inducing transient. (1) The peak detector experiences an invalid transient which causes its output to trigger the LED, starting the LED_1shot timer. (2) After t_{stuck} , if the CFD has not registered an event, rst_stuck asserts, (3) resetting the peak detector (along with the LED and CFD outputs).	37
3.1	Chip die photo with main structures annotated. Photo courtesy of Hani Gomez.	40
3.2	Bandgap voltage versus temperature with the envelope of the standard deviation. Spikes at low temperatures were from condensation within the chamber.	42
3.3	The reference routing network for power distribution. There are three LDOs on the chip, all implemented in a similar fashion.	43
3.4	The low dropout regulator used within the chip. The enable device (green) is not included in the always-on regulator, and the reference voltage V_{REF} is produced using an on-chip bandgap reference circuit.	44
3.5	Block diagram of the small/no-shape signal chain.	44

3.6	Measurement setup and procedure for gathering timing statistics for the small signal chain. (a) Configure the DG535, TDC, and chip scan chain. (b) Trigger the DG535. The DG535's triggered output is then used as the START event to the TDC. A subsequent pulse nominally 1 μ s wide with an amplitude anywhere from 50mV to 600mV is routed down three paths: an attenuator (Kay Elemetrics 839); a coaxial cable roughly 60cm longer than that of the attenuator for an additional \approx 2ns delay; and directly to the positive input of the LED comparator. Each pulse amplitude test is performed 500 times with at least 100ns between pulses. (c) The chip output is latched and level shifted from the 1.8V core voltage to 3.3V to be fed into the TDC as a STOP event.	45
3.7	(a) Jitter and (b) time difference of arrival of the measured pulses. Worst case jitter was measured at 176.9ps _{rms} , and timing walk at 682ps.	46
3.8	(a) The resistive ladder DAC was chosen for its simplicity and guaranteed monotonicity. (b) The analog mux with N bits was constructed with $2^N - 1$ two-to-one muxes to enable direct feed of binary selection bits with no additional encoding.	47
3.9	(a) Voltage DAC transfer function with the supply voltage set to its lowest value. Each data point is the average of 100 measurements. The gain is 3.44mV/LSB for a full scale range of 877.22mV. (b) The DAC's RMS noise \leq 1.17mV, or 2.9 LSB. Measured with nominally 10nF of decoupling capacitance on the output of the DAC. (c) DNL min/max -0.17/0.23 LSB. (d) INL min/max -0.19/0.31 LSB.	48
3.10	The core of every comparator consisted of several cascaded fully differential stages (a) and one final stage for a single-ended conversion (b).	49
3.11	High level diagram of the autozeroed comparator. Because there is no clock, the sampling phase ϕ_1 occurs when the chip is reconfigured, i.e. the scan chain is LOADED into the rest of the chip. In our use, the nulling amplifier and the nulling pins are differential; we show the single ended variant here for clarity.	50
3.12	The core fully differential amplifier in grey, with its offset control in black and boxed on the left. The choice to burn additional current by adding R_{OS} was to more consistently define the gain of the offset control across corners; biasing resistors largely guaranteed the current source for I_{OS} behaved as such.	51
3.13	A standard scan chain circuit (a) and its operation (b). Buffers and the like have been removed for clarity.	52
3.14	A triple modular redundant scan cell as it was used in the chip. As a defense against timing violations, the clock was routed in reverse order relative to the input data signal.	53
3.15	The one shot pulse generator, with triplicated components outlined. Each outlined section was followed by a 3-to-3 majority vote on its output(s). The reset switch is necessary to ensure consistent output pulse widths when input events are closely spaced in time relative to the RC time constant.	53

3.16	Normal one shot operation without the reset ϕ_{rst} in the presence of a single short pulse from the fully settled stable state. The grey line is the inverter switching point. Note that the the RC node takes finite time to settle to $\approx 0V$ even after the output pulse has terminated.	54
3.17	Inconsistent output pulse widths due to the prolonged settling time of the RC node.	54
3.18	Simulated output pulse widths of a one shot pulse generator with (blue) and without (orange) the reset switch for $R = 20k\Omega$ and $C = 1pF$. Input pulses were 1ns wide with pulse spacing 1-100ns apart, randomly sampled. The median pulse width without the reset switch was 15.4ns and 16.8ns with the reset.	55
3.19	Measurement setup for gathering timing statistics for the small signal chain. (a) Configure the DG535, TDC, and chip scan chain. (b) Trigger the DG535. The DG535's triggered output is then used as the START event to the TDC. A subsequent pulse nominally 2ns wide with an amplitude ranging from 0.1V to 1V is connected with a 50Ω termination to the PCB and AC coupled with a 2pF capacitor for current pulses of 1.2mA to 12mA into the preamplifier. Each pulse amplitude test is performed 500 times with at least 100ns between pulses. (c) The chip output is latched and level shifted from the 1.8V core voltage to 3.3V to be fed into the TDC as a STOP event.	56
3.20	(a) Jitter and (b) time difference of arrival of the measured pulses for the main signal chain. Worst case jitter was measured at $743ps_{rms}$, and timing walk at 601ps.	56
3.21	The preamplifier. The referencing used for biasing is generated by a resistor ladder DAC identical to the one used in the LED (Section 3.2.2.1).	57
3.22	(a) The simulated and ideal peak of the preamplifier's output as the amplitude of the input pulse increases. The ideal values are calculated by linearly fitting the simulated data. The charge corresponds to $2Me^-$ to $20Me^-$. (b) The compression of the peak as the size of the pulse increases.	58
3.23	The peak detector. The switch (grey) was not included in this version of the chip.	58
3.24	The positive and negative inputs to the LED and CFD comparators, without a reset on the output of the preamplifier. If t_{1shot} is short relative to the settling time of the preamplifier, the peak detector output (red) will rise past the LED threshold and trigger the output once more, causing the appearance of afterpulsing.	59
3.25	(a) Static error of the peak detector. (b) The percent error of the measured output voltage relative to its target value. (c) The standard deviation of the measured output voltage. The error and statistics show noticeable spikes.	60
3.26	Tracking the spikes in voltage for (a) the largest and (b) second largest spikes in measurement variance.	61
3.27	Sallen-Key topology associated with the transfer function in Equation 3.11.	61
3.28	The resistive divider attenuator.	62
4.1	Chip die photo with main structures annotated. Photo courtesy of Alexander Alvara.	65

4.2	Time difference of arrival of the simulated pulses with an emulated benchtop setup and microchannel plate.	66
4.3	67
4.4	Trimmed comparator stage.	68
4.5	Timing walk versus input common mode over pulses ranging from 2Me^- to 20Me^-	68
4.6	Group delays vs. frequency of a second order and fourth order Bessel filter, with the point of 10% and 50% change marked with a black horizontal line.	69
4.7	Two Sallen-Key filters to give the fourth order transfer function in Equation 4.2.	70
5.1	A typical design flow.	73
A.1	Annotated schematic of the digital LDO.	84
A.2	Block diagram of the sensor ADC subsystem.	86
A.3	(a) ADC code readout with a PGA gain setting of $2\text{V}/\text{V}$. The slope corresponds to approximately $1.2^\circ\text{C}/\text{LSB}$. (b) The number of measurements associated with each temperature measurement, taken with a TMP102 digital temperature sensor. Temperature precision was 0.01°C	87
A.4	Subblocks of the front end. (a) Regulator. (b) PTAT. Body connections are to ground. [45] (c) Programmable gain amplifier. (d) Successive approximation register analog to digital converter.	88
A.5	Sensor ADC (a) DNL and (b) INL. There are clear spikes at bit flips which are consistent with binary DACs, with some nonmonotonicity with the DAC.	89
A.6	The ADC output code versus input voltage. The nominal FSR is $V_{\text{DD,sensor}} = 1.2\text{V}$, and each data point is the result of 5 averaged samples. Unfortunately the raw data for this plot has been lost.	89
B.1	ModelSim window after launch, showing Library and Project tabs on the left and Wave tab on the right.	92
B.2	Creating a new work directory	93
B.3	The “run” button circled in red	94
C.1	Chip layout screenshot. The analog front end fits in $130\mu\text{m}\times 130\mu\text{m}$	117
C.2	A block diagram of the analog front end of the optical receiver.	121
C.3	General operation of the analog front end for a pulse width modulated input.	121
C.4	The self-biased transimpedance amplifier.	122
C.5	An eye diagram of the output of the TIA when it’s fed a 1.84Mbps OOK signal to program SCM with a 10nA signal. The eye opening is roughly 17.5mV	123
C.6	(a) The second order low pass filter with a corner frequency of $\omega_0 = \frac{G_m}{C}$. The G_m of each operational transconductance amplifier is controlled with the tunable constant- g_m circuit shown in (b), used for adjusting the tail current of each OTA.	124
C.7	(a) The main amplifier for the preamp. (b) The common mode feedback amplifier used in the preamp to maintain biasing.	125

C.8	An eye diagram of the output of the preamp when the TIA is fed a 1.84Mbps OOK signal to program SCM with a 10nA signal. The eye opening spans from 130mV to -15mV because of DC imbalance, suggesting 4b5b is necessary for any OOK.	126
C.9	The resistive ladder intended for a 1 μ A bias current.	126
C.10	(a) The strongarm latch and (b) subsequent SR latch used in the clocked comparator.	127
F.1	The L-resistor footprint. Only one direction (N/S or E/W) is populated.	149

List of Tables

1.1	Location and scale parameters of the Moyal distributions used to fit TOF data.	5
1.2	Some useful definitions for discussing radiation in electronics.	6
1.3	A summary of the various single event effects which can occur due to ionizing radiation.	12
2.1	Measured parameters for SPAN-I's previous chip designed by Johns Hopkins APL, and the target specification for the iteration of SPAN-I discussed in this dissertation. Area limits were set by available area on the BSAC shuttle. The 180nm process was chosen because BSAC does not need to pay for the area.	27
3.1	Chip V1 versus the target specifications.	41
4.1	Chip V2 simulated performance versus the target specifications.	66
4.2	Delay filter passband group delay and attenuation.	70
A.1	Simulated performance of the digital LDO.	84
A.2	Location of digital LDO libraries in Cadence Virtuoso	85
A.3	Location of sensor ADC libraries in Cadence Virtuoso	86
C.1	Optical receiver target specifications.	116
C.2	Simulated optical receiver versus the tested version in [81].	116
C.3	Optical receiver chip I/O and associated pad locations.	118
C.4	Description of optical receiver I/O.	119
C.5	Optical receiver scan bits.	120
C.6	129
C.7	Location of libraries in Cadence Virtuoso	130
D.1	SPAN-Ion specific schematic generators.	132
D.2	General analog schematic generators. There may be some overlap between other pre-existing BAG libraries due to accidental parallel or simultaneous creation. .	133
D.3	General digital schematic generators. There may be some overlap between other pre-existing BAG libraries due to accidental parallel or simultaneous creation. .	134
E.1	Chip V1 I/O and associated pad locations. Unconnected pads are marked as x.	136

E.2	Chip V1 I/O descriptions.	137
E.3	Chip V1 scan bits.	138
E.4	Top-level Cadence libraries for Chip V1.	143
F.1	Chip V2 I/O and associated pad locations. Unconnected pads are not included.	146
F.2	Chip V2 I/O descriptions.	147
F.3	Chip V2 scan bits.	148
F.4	Top-level Cadence libraries for Chip V2.	150
G.1	Locations of useful tribal knowledge.	152

Acknowledgments

As with any PhD, it would not have been possible without a veritable village of people. To my advisor, Kris Pister: without your encouragement, I wouldn't have even applied to a PhD program. Your endless supply of enthusiasm, patience, and Hawaiian shirts has made for both a fun and educational grad school experience.

To my collaborators at SSL—Robert, Roberto, Ken, Davin—and the rest of the wonderful people there: Thanks for putting up with all my derpy questions, and for sharing your experience and insight over the last few years. Your breadth and depth of expertise and knowledge have been an inspiration for an engineer who occasionally felt like things had gone stale.

I'd also like to thank Carl Grace and Professors Stojanovic and Pilawa-Podgurski for their encouragement and enthusiasm as members of my qualifying exam committee; and Professor Lu for agreeing on extremely short notice to be a part of my dissertation committee.

Throughout grad school I've also had the privilege of being in the same group as some of the best people I know, and been lucky enough to call several of them my friends. Alex, your humor, warmth, resilience, and snappy wardrobe have kept me sane through the crazy-train of grad school more than you probably know. Mauricio, thank you for simultaneously being a great researcher, incredibly humble, and one of the most genuinely caring people I've had the pleasure of calling a friend. Daniel (Teal), you are the human embodiment of sunshine, if sunshine wore exclusively black; your technical acumen and unflappable optimism inspire the people around you to be better engineers, researchers, and people. Hani, thank you for showing me that there's more to happiness than productivity, and for teaching me how to be both a better communicator and researcher. Alexander, you're always there to root for people and help them out, and being your friend and colleague has taught me to be a better advocate for myself and for others. Nate, thanks for being a great desk buddy and giving me perspective for life beyond school—who'd have thought that we'd end up publishing together? Many thanks to the old(er) guard SC μ M crew of Fil, DB, Brad, and Osama for your mentorship and guidance in getting me started with grad school. Bhuvan, you've been the best undergrad a grad student could hope to work with, and there is zero doubt in anyone's mind that you're going to do amazing things. To the youngins of the Pister group—Yu-Chi, Yichen, Titan, Daniel (Lovell), Daniel (Finell), and generations to come—I look forward to the cool things y'all will accomplish in the coming years.

Beyond the Pister group, other members of the EECS department—Kevin (Zheng), Dalene, Shirley, Pat, Glenna, Finsen, Yifan, Shm, and countless more—have been invaluable for their insight, assistance, and endless generosity with time and energy throughout my studies at Berkeley.

Outside of grad student life, Berkeley has introduced me to Katherine, Raymond, Saavan, Divya, Jayss, Zeke, and many more people for whom I'd happily trade all my material possessions if it meant I could bring them with me. You all bring such joy and light into the world, and it makes leaving that much harder.

Finally, I'd like to thank my family. Lyndon, you've always been a better person than I, and you always know how to cheer me up no matter the circumstances. Mom and Dad, this would not have been possible without your unwavering support, advice, and recognition that "When are you going to graduate?" is a bad question to ask a PhD student.

Thank you all for believing in me, being there, and making this possible.

Chapter 1

Introduction

The utility of a circuit cannot be defined in a vacuum; it must be contextualized by the circuit's application. This chapter will describe the Solar Probe Analyzer for Ions as the specific target use case of this dissertation's hardware, with the understanding that the methods and analysis here can be broadly applied to similar systems. It will provide background on ionizing radiation in the context of integrated circuit design, as well as a discussion of the Berkeley Analog Generator as it was used for this dissertation.

1.1 The Solar Probe Analyzer for Ions

The Solar Probe Analyzer for Ions (SPAN-Ion) is an electrostatic analyzer (ESA) designed by the Berkeley Space Sciences Lab to measure the ion composition and 3D distribution function of the thermal corona and solar wind plasma [44]. It derives much of its design from the Mars MAVEN electrostatic analyzer, has a legacy which includes the Parker Solar Probe and the Lunar Gateway, and is slated for use in the Mars EscaPADE mission in mid-2024. Figure 1.1 shows a cross section of the instrument with the top hat ESA and time-of-flight (TOF) apparatus, as well as a block diagram of components of the electronics box below.





Figure 1.1: [44] Block diagram of the SPAN-I sensor, including ESA, TOF, and individual components of the electronics box.

At the highest level, its purpose is to advance our understanding of space weather to improve our nowcasting and forecasting capabilities. SPAN-Ion uses a time-of-flight mass spectrometer to determine the mass/charge ratios of the ions selected by the electrostatic analyzer (Figure 1.2).

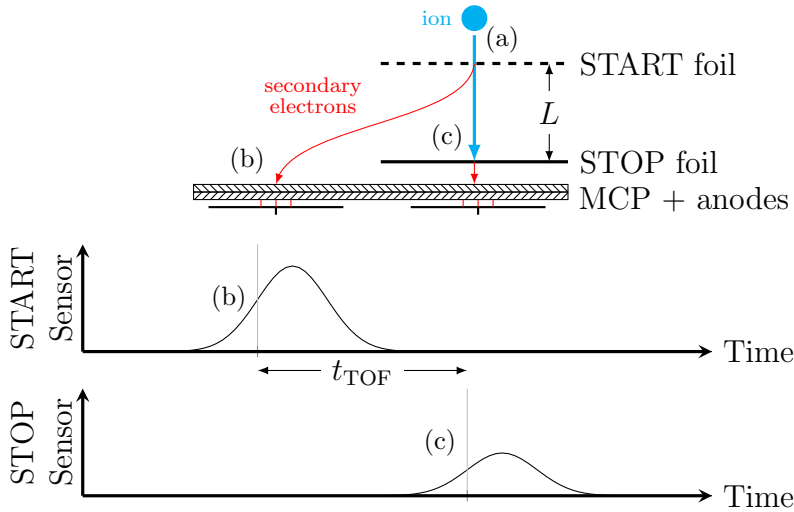


Figure 1.2: Operation of SPAN-I's time-of-flight mass spectrometer. (a) The ion is accelerated with a known potential to a speed defined by Equation 1.1. (b) Upon impact, the START carbon foil releases secondary electrons which are amplified then picked up by the START anode as a current pulse. (c) The ion travels through the START foil to the STOP foil, which is a fixed distance L away. This produces a STOP pulse by the same mechanism as the START pulse.

The mass spectrometer operates by accelerating the selected ion with a potential $U = -15\text{keV}$. Assuming the ion has negligible initial kinetic energy, the ion reaches a speed determined by its mass/charge ratio, described in Equation 1.1.

$$qU \approx \frac{1}{2}m\dot{x}^2 \quad (1.1)$$

The ion strikes a START carbon foil, triggering the release of secondary electrons which are then directed to a microchannel plate (MCP) to produce a START pulse. The ion continues through the START foil and travels $L = 2\text{cm}$ to the thicker STOP foil, and upon impact again produces secondary electrons which are amplified by the MCP to produce a STOP pulse. Assuming no energy is lost upon collision with the START foil, the time the ion takes to traverse the distance L is defined by Equation 1.2.

$$t_{\text{TOF,lossless}} = \frac{L}{\dot{x}} \approx \frac{L}{\sqrt{2U}} \sqrt{\frac{m}{q}} \propto \sqrt{\frac{m}{q}} \quad (1.2)$$

SSL provided calibration data, fitted to Moyal distributions, for the TOF of different ion species in Figure 1.3.

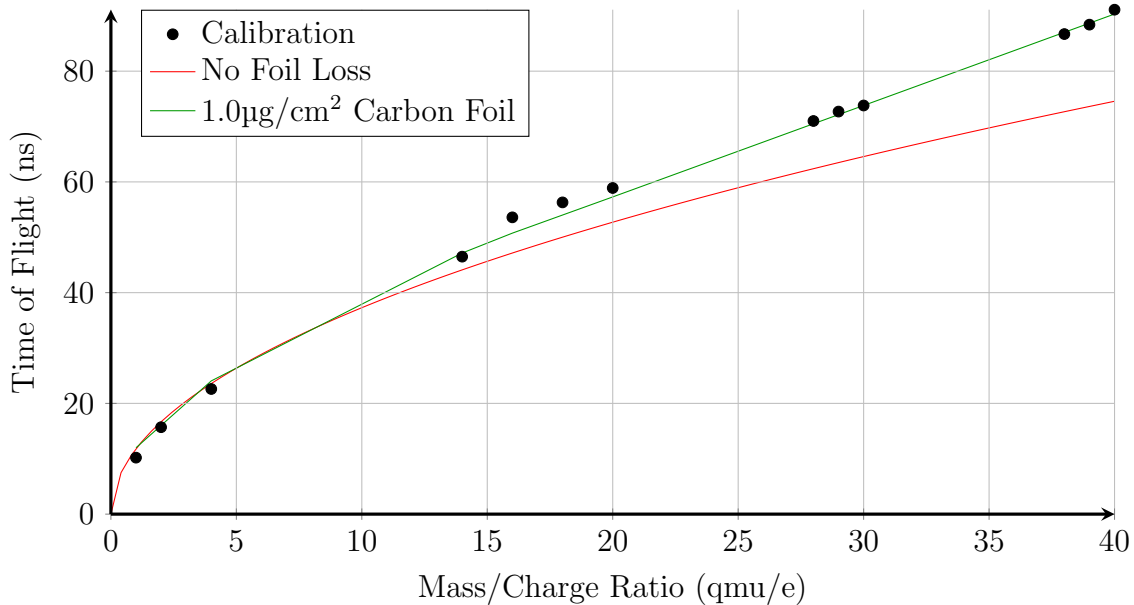


Figure 1.3: Time of flight calibration data versus theoretical time of flight with and without a $1.0\mu\text{g}/\text{cm}^2$ carbon foil. Each data point corresponds to the time of flight mode bin with 100ps resolution. Calibration data courtesy of SSL.

Comparing the ground calibration data and idealized expression for time of flight in Equation 1.2, we see in Figure 1.3 that there is non-negligible energy loss from the collision with the START carbon foil. Using TRIM [88], we estimate a carbon foil thickness of $\approx 1.0\mu\text{g}/\text{cm}^2$. More importantly for this dissertation, the timing accuracy and precision required to distinguish between the target mass/charge ratios is half the minimum difference between any two mass/charge ratios, or 600-800ps.

Both variation in energy loss from the carbon foil collision and initial kinetic energy cause spread within the TOF distribution—a phenomenon broadly categorized as straggling. The flexibility of the target timing precision is due to significant straggling within each TOF distribution, especially at higher mass/charge ratios. From Table 1.1, the Moyal scale parameters for TOF distributions for ions with mass/charge ratios of 29 and 30 are 5.6ns and 6.0ns—more than $4\times$ the difference in times of flight used to distinguish the two mass/charge ratios.

MQ Ratio	TOF Moyal σ (ns)	TOF Mode (ns)
1	0.3	10.2
2	0.4	15.7
4	0.3	22.6
14	1.9	46.5
16	2.8	53.6
18	3.9	56.3
20	3.2	58.9
28	6.4	71.0
29	5.6	72.7
30	6.0	73.8
38	8.8	86.7
39	8.6	88.4
40	9.7	91.1

Table 1.1: Location and scale parameters of the Moyal distributions used to fit TOF data.

1.2 Radiation and Integrated Circuits

Ionizing radiation refers to particles and photons with sufficient energy to detach electrons from atoms or molecules. This section will discuss its effects on integrated circuits, hardening techniques against it, and methods for validation in representative environments. We assume a basic understanding of energy bands and band diagrams.

1.2.1 Effects

Ionizing radiation can cause cumulative effects over periods of prolonged exposure, as well as more transient single event effects (SEEs). The former refers to changes in device characteristics due to total ionizing dose (TID), though it can also cause non-ionizing damage in the form of displacement damage (DD). The latter encompasses a host of phenomena which all result from the relatively short-lived injection and redistribution of charge as the ionizing radiation interacts with electronics.

Before continuing, a note about units: astronomers and astrophysicists often use CGS units over SI. One reason I've heard for this is so the energy density of an electromagnetic field can be written without needing the permittivity ϵ_0 and magnetic permeability μ_0 of free space (Equation 1.3).

$$U_{\text{CGS}} = \frac{1}{8\pi} (\mathcal{E}^2 + B^2) \quad (1.3a)$$

$$U_{\text{MKS}} = \frac{\epsilon_0}{2} \mathcal{E}^2 + \frac{1}{2\mu_0} B^2 \quad (1.3b)$$

The units in Table 1.2 are CGS for the most part only to give a sense of normalization quantities (e.g. per area, volume, etc.). This dissertation will always specify units.

Term	Definition	Common Units
Linear Energy Transfer (LET)	$\frac{dE}{dx} \frac{1}{\rho_{\text{medium}}}$	$\frac{\text{MeV} \times \text{cm}^2}{\text{mg}}$
flux	$\frac{\# \text{ particles}}{\text{area} \times \text{time}}$	$\frac{1}{\text{cm}^2 \times \text{s}}$
fluence	$\int (\text{flux}) dt = \frac{\# \text{ particles}}{\text{area}}$	$\frac{1}{\text{cm}^2}$
dose	$\text{LET} \times \text{fluence} \times 1.6(10^{-7} \text{Gy}) \times \frac{\text{mg}}{\text{MeV}} = \frac{E}{\text{mass}}$	1Gy = 100rad
SEE Cross-Section	$\frac{\# \text{ of errors}}{\text{fluence}}$	m^2

Table 1.2: Some useful definitions for discussing radiation in electronics.

1.2.1.1 Total Ionizing Dose

Total ionizing dose (TID) is used when describing the long-term degradation of electronics due to ionizing radiation. Figure 1.4 [49] shows how ionizing radiation causes this long-term damage.




Figure 1.4: [49] An energy band diagram of a MOS structure for positive gate bias, indicating major physical processes underlying radiation response.

First, the ionizing radiation generates electron hole pairs (EHPs). The quantity of EHPs depends on a number of variables including the bandgap of the medium and its density—the energy to generate an EHP in crystalline silicon is 3.6eV, while the energy required to generate an EHP in silicon dioxide is 17eV [70]. Second, some of the pairs recombine, but those which don't (usually holes due to lower mobility [6]) polaron hop via shallow traps into the dielectric. Finally, the unrecombined holes end up in oxide traps, often near the surface interface, though they can also appear in the bulk and at the interface [24].

Process improvements have generally worked in favor of robustness against total ionizing dose; thinner oxides and improved interface engineering allow for fewer defects and traps. However, dielectrics such as the buried oxide in silicon-on-insulator processes and field oxides used in trench isolation still experience the effects of total ionizing dose.

For MOS devices, holes trapped at interfaces do not anneal out like those in the dielectric can, resulting in a shift in threshold voltage defined by Equation 1.4 where C_{ox} is oxide capacitance, t_{ox} is oxide thickness, and ρ_{ox} is trap density.

$$\Delta V_t = -\frac{Q_{\text{interface}}}{C_{ox}} = -\frac{1}{\epsilon_{ox}} \int_0^{t_{ox}} x \rho_{ox}(x) dx \quad (1.4)$$

Note that this is an absolute threshold voltage shift, meaning NMOS devices will become faster and PMOS devices will become slower, e.g. [21] (Figure 1.5).




Figure 1.5: [21] Threshold voltage shift in commercial 130nm CMOS as a function of TID for different (a) NMOS and (b) PMOS transistor sizes, up to 136Mrads. PMOS threshold shift is absolute value. The last point refers to full annealing at 100°C.

Flicker or $1/f$ noise in MOS devices also degrades with total ionizing dose. One often-cited contributor to flicker noise in MOS devices [50] is carrier trap and release at or near the silicon-oxide interface. There is contention [34] surrounding the McWhorter model and

the source of electronic flicker noise, but physical origins of flicker aside, [51] measured the $1/f$ noise in MOS devices through irradiation and annealing, noting a strong correlation between flicker noise and oxide trapped charge, and no correlation between flicker noise and interface trap charge (up to 10kHz). [63] ran extensive device characterizations on $0.13\mu\text{m}$ CMOS devices and showed increases in the flicker noise coefficient K_f for n-channel devices, particularly at low drain currents, with significantly less of an increase for p-channel devices (Figure 1.6).

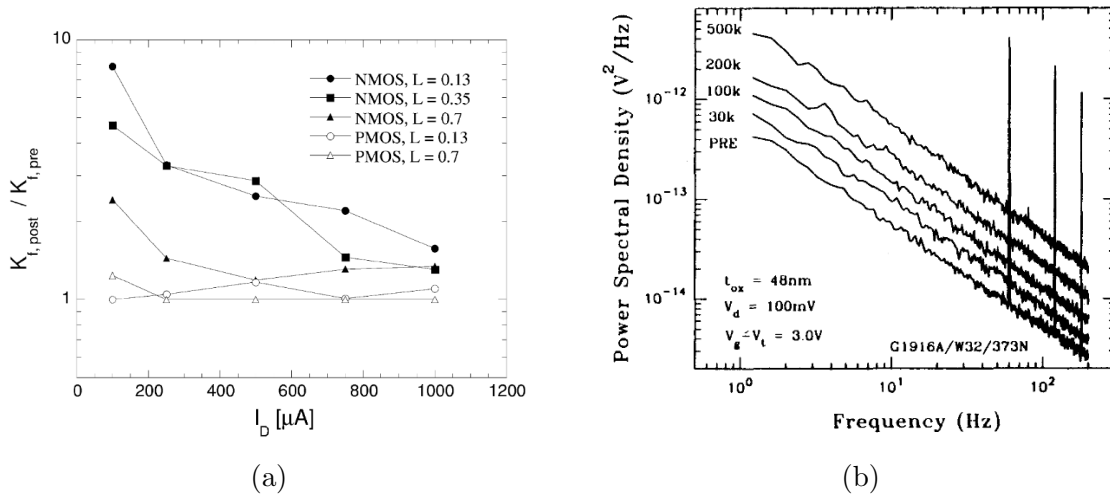


Figure 1.6: (a) [63] Ratio between the value of $1/f$ noise parameter K_f at 10Mrad absorbed dose of ^{60}Co γ -rays and the value before irradiation as a function of drain current I_D for NMOS and PMOS devices in the $0.13\mu\text{m}$ process. (b) [51] $1/f$ noise spectra S_v as a function of total dose. The device was under +6V bias during irradiation.

For MOS device leakage, [35] showed significant increases in off-state leakage current for HV, I/O, and core devices, particularly under subthreshold bias conditions (Figure 1.7). They also observed a “hump” in the I_D - V_{GS} curves for several devices, attributable to oxide trapped charge trapping in the corners of devices’ shallow trench isolation (STI).

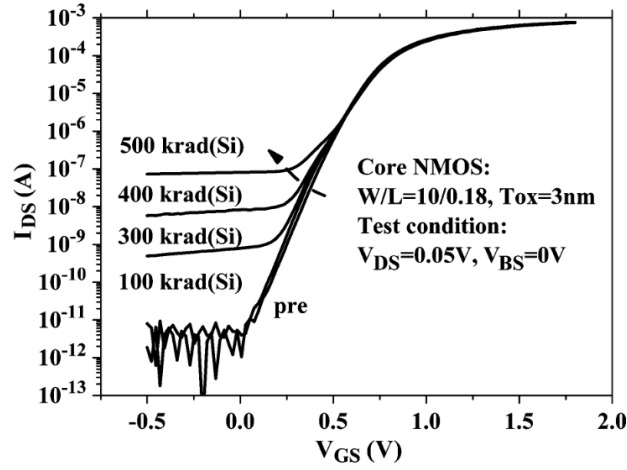


Figure 1.7: [35] I_{DS} - V_{GS} characteristics before and after TID irradiation for a core NMOS device.

Trapped charge in isolation oxide can also form a parasitic transistor, exacerbating drain induced barrier lowering (DIBL) by providing a charge path through what should have been isolation oxide on the sides of the transistor [86]

Ionizing radiation also tends to exacerbate mismatch by degrading devices more quickly, but at different rates. [76] measured variation in the mismatch parameters ΔV_t and $\frac{\Delta\beta}{\beta}$ based on Equation 1.5.

$$\frac{\Delta I_D}{I_D} = \frac{\Delta\beta}{\beta} + \Delta V_t \left(\frac{g_m}{I_D} \right) \quad (1.5)$$

Examining threshold voltage and current factor mismatch before and after irradiation, they found marginal increases in the mismatch parameters for standard MOSFETs in their given process, and smaller increases in these parameters for dedicated enclosed layout devices (Figures 1.8 and 1.9).

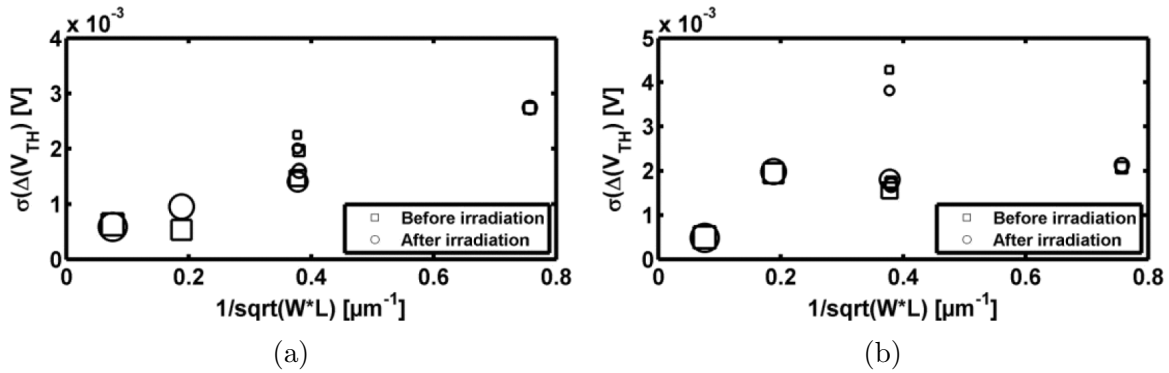


Figure 1.8: [76] Threshold voltage mismatch between identically designed pairs of a regular and b enclosed transistors before and after γ irradiation up to 100kGy.

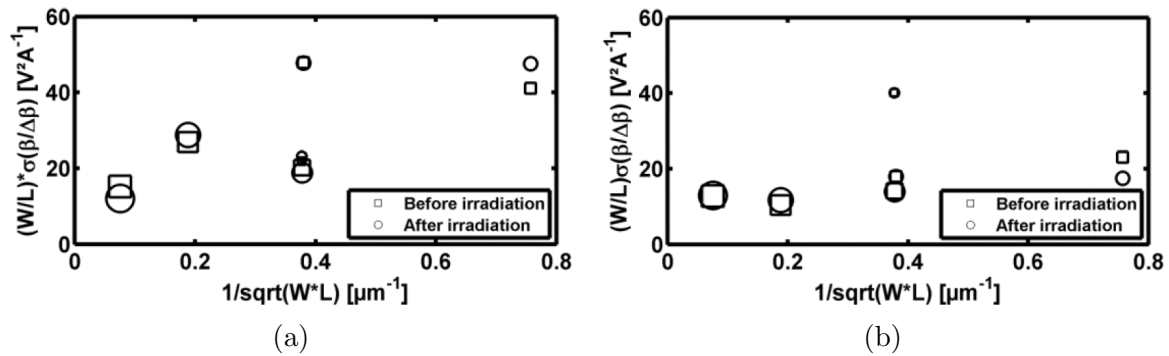


Figure 1.9: [76] Standard deviation of the current factor mismatch between identically designed pairs of a regular and b enclosed transistors before and after γ irradiation up to 100kGy.

Similarly, bipolar devices experience increased base current for a given base-emitter voltage, corresponding to decreased gain. Interestingly, [19] found that gain degradation is worst at the lowest dose rate, with the gain recovering after annealing, suggesting that the previously standard method for testing bipolar circuits for space applications was not valid for all bipolar circuits. [38, 39, 66, 67, 68, 83] developed analytical expressions for bipolar base current, distinguishing between the primary modes of npn gain degradation—oxide trapped charge and interface traps—and pnp device degradation—increased surface recombination velocity from interface traps.

1.2.1.2 Single Event Effects

Figure 1.10 shows the phases of a single event transient. When ionizing radiation interacts with electrically sensitive regions such as pn junctions, a track of EHPs with high carrier

concentration forms in its wake. In the case of structures like pn junctions, this causes the depletion region to extend deeper into the substrate, causing a drift-induced current spike. The errant charge then diffuses into the depletion region, and the current eventually settles to its initial value [7]. The result is a current pulse which can be approximated as Equation 1.6, where Q is the total charge collected during the event, τ_f is the time constant for the initial drift and τ_r is the time constant for diffusion [53].

$$I(t) = \frac{Q}{\tau_f - \tau_r} (e^{-t/\tau_f} - e^{-t/\tau_r}) \quad (1.6)$$

A caveat to Equation 1.6: [10] concluded that a dual double exponential was a more accurate representation of long SETs. The vast majority of the SEU- and SET-prone logic in Chapters 3 and 4 is edge- rather than level-based, and so the added complexity in modeling was unnecessary.

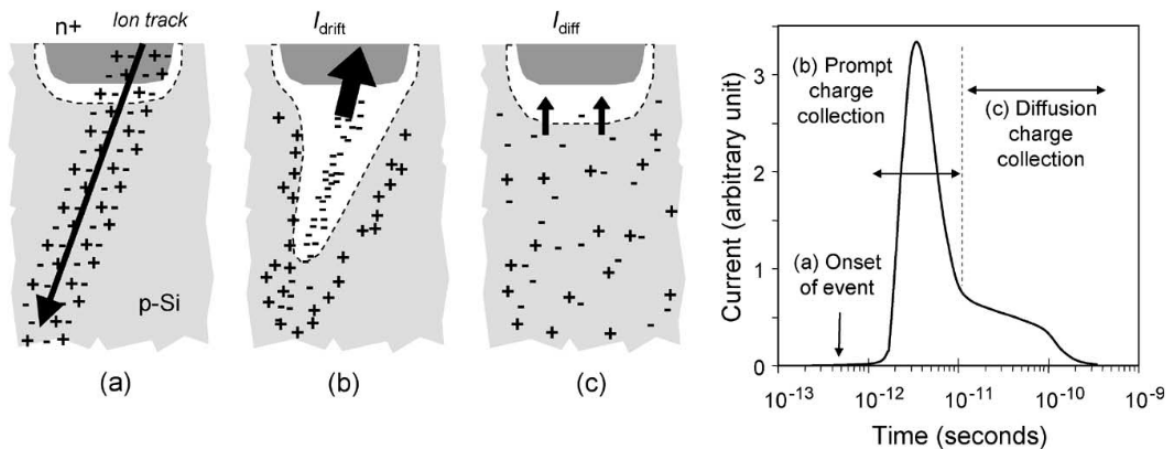


Figure 1.10: [7] Charge generation and collection phases in a reverse-biased junction and the resultant current pulse caused by the passage of a high-energy ion.

This phenomenon of a single event transient (SET) is the foundation of and one of the large variety of single event effects (SEEs) listed in Table 1.3, but ultimately the severity of the effect depends on whether the error is destructive and downstream hardware. Unlike total ionizing dose, process node scaling has generally worsened processes' inherent robustness against SEEs [46]. Figure 1.11a shows how increasing device density compares with the diameter of the charge track induced by a light ion; more tightly packed devices in a given area makes it so a single event can affect multiple devices, producing errors such as multi-bit upsets. Figure 1.11b estimates the charge used to store a bit on an inverter's gate; the push for smaller, lower power devices and their correspondingly lower operating voltages requires less charge per operation, meaning SEEs—which unfortunately do not scale with process nodes—produce larger transients and are more likely to affect information storage.

Single Event Effect	Short Description	Devices	References
single event burnout (SEB)	Destructive triggering of a parasitic BJT, followed by regenerative feedback.	power transistors	[33]
single event dielectric rupture (SEDR)	Formation of a conducting path through dielectric in a high field region of the dielectric.	CMOS	[4, 75, 79]
single event gate rupture (SEGR)	Gate dielectric rupture due to high \mathcal{E} field from SEE charge.	CMOS	
single event functional interrupt (SEFI)	Data path corruption which leads to the disruption of normal operation. This is an SEU in a register critical to device function.	state machines, control	[37]
single event hard error (SEHE)	Permanent and unalterable state change due to damage to a memory cell.	memory	[17, 74]
single event latchup (SEL)	Self-sustaining current caused by a parasitic pnpn kicked into regenerative forward bias.	CMOS	[8]
single event snapback (SESB)	Amplification of avalanche current from activation of the parasitic BJT of a MOSFET.	MOSFETs, SOI	[70]
single event transient (SET)	SEE impulse response.	all	[7]
single event upset (SEU)	Corruption of a bit stored in memory.	memory	[47]
multi-bit upset (MBU)	Corruption of multiple bits in memory from a single event.	memory	[15, 59, 60]

Table 1.3: A summary of the various single event effects which can occur due to ionizing radiation.

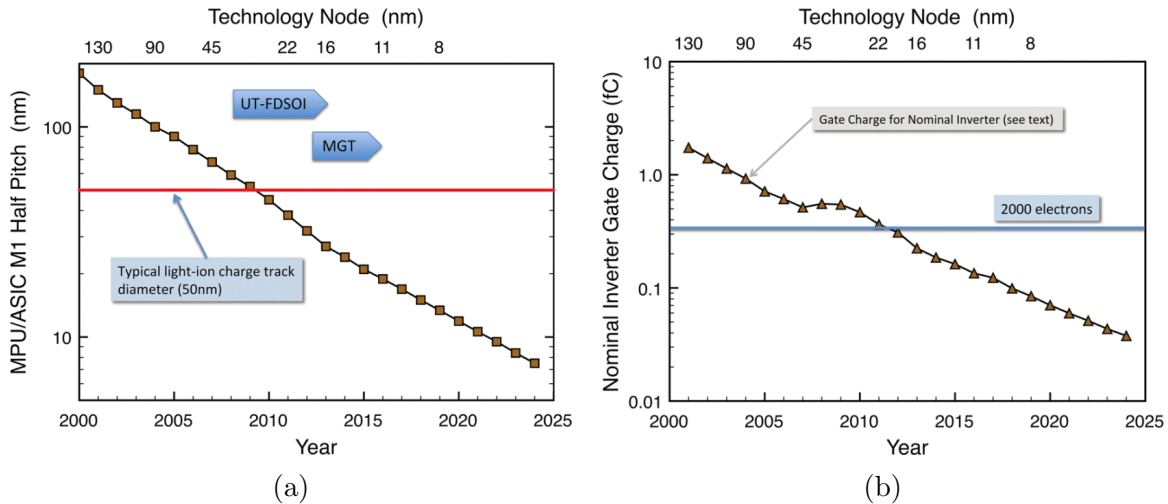


Figure 1.11: [46] (a) ITRS scaling of first-level metal half pitch and progression into ultra-thin fully depleted silicon-on-insulator and multiple gate technologies, compared with a representative diameter of the ionized free charge distribution in silicon in the wake of a light ion interaction. Node indicators are approximate. (b) ITRS scaling of gate charge for a nominal 3X/6X NMOS/PMOS inverter.

For a process with $\approx 1\mu\text{m}$ of sensitive depth, a single 1.5MeV α particle would produce roughly 75,000 electrons which appear in a single event transient. If dumped in its entirety on a 1fF capacitor, this corresponds to a voltage spike of 12V —sufficient to damage core device gates in many sub-micron technologies. With respect to transients, if drift and diffusion time constants in Equation 1.6 are on the order of 100ps to 1ns [9, 27], this corresponds to current spikes on the order of hundreds of μA . Section 1.2.3 will address specific values used for validation in this dissertation.

1.2.1.3 Displacement Damage Dose

Displacement damage refers to atoms displacing from their lattice sites, where vacancies and interstitials then migrate to either recombine or form stable Frenkel defects. The primary knock-on atoms (PKAs) displaced by the initial event can then proceed to create additional collisions and defect cascades [82]—a 1MeV neutron or high energy proton can produce a 50keV recoil PKA, which goes on to produce additional defects. The energy per mass associated with displacement damage is referred to as displacement damage dose (DDD) or total non-ionizing dose (TNID). In silicon, the threshold displacement energy is 21eV , and a relatively small fraction of energy deposition goes into displacement damage [2, 73]. As such, this dissertation will not address displacement damage during the circuit design process; [2] includes an overview of the effects and severity of displacement damage for a host of electronic devices.

1.2.2 Radiation Hardening Electronics

“Radiation hardening” refers to the myriad techniques used to make hardware robust against the effects described in Section 1.2.1. Figure 1.12 shows the levels where an electronics system designer can make adjustments to account for radiation effects. This section will *not* include a discussion of radiation shielding.

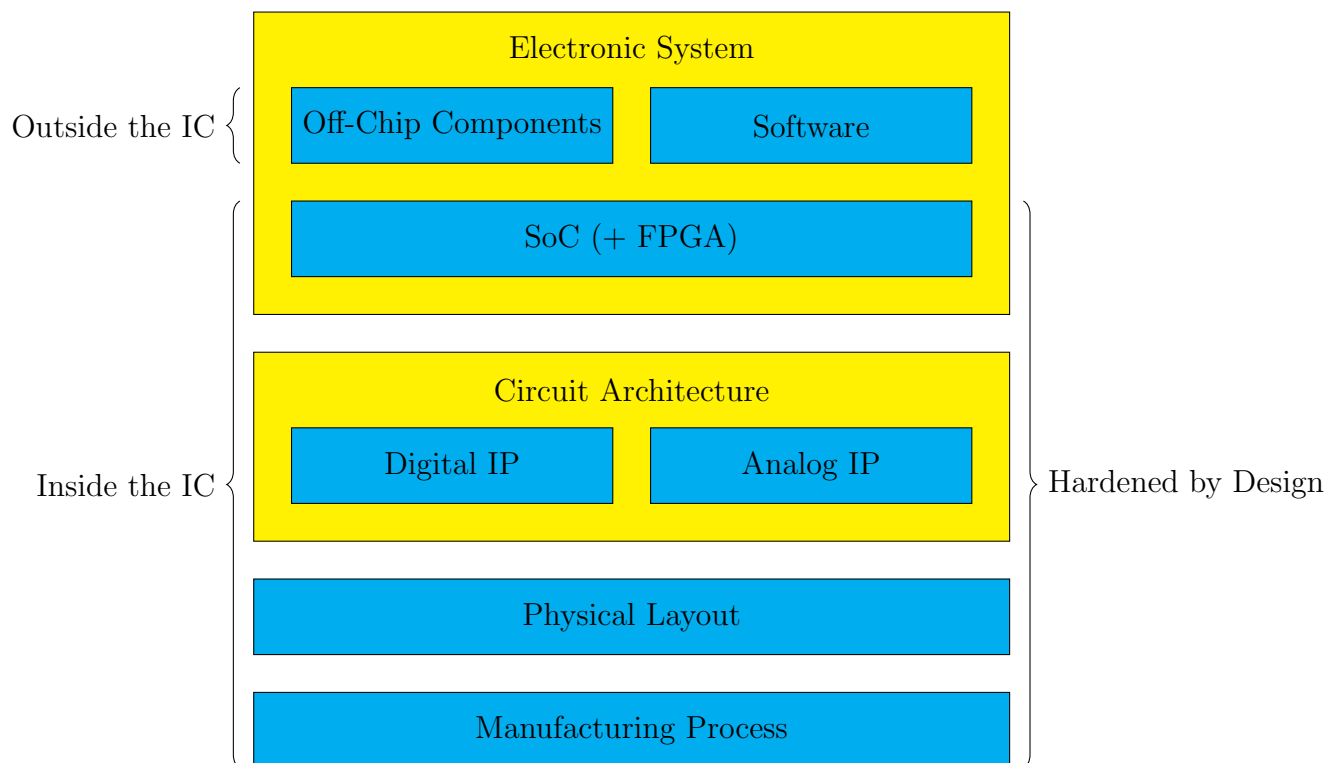


Figure 1.12: Layers of radiation hardening.

1.2.2.1 Foundry and Process

At the foundry level, minimizing the amount of oxide near device channels aids against total ionizing dose. Substrate choice, process node, and stackup selection can have a significant effect—[69] observed reduced sensitivity to SEEs for silicon-on-insulator ICs; devices with more physically compacted memory cells can tip the balance between a single event upset and a multi-bit upset.

Some process development kits (PDKs) include specialized devices to reduce the amount of oxide near channels [42, 5]. Enclosed and edgeless devices can also provide vast improvements over standard devices for post-irradiation device leakage [42]. These specialized devices are nonstandard for most PDKs, however, and can introduce additional difficulty: the two-dimensional current profile makes traditional tactics of effective device width and

length more complicated, which in turn makes the modeling of and laying out for good device matching difficult. The physically larger devices also leads to larger node capacitance relative to standard devices. While this can be beneficial against single event transients for filtering and increased charge motion, by the same token it's detrimental to power consumption.

1.2.2.2 Circuit Designer

Once a process node has been chosen, there are still a huge variety of methods a circuit designer can employ for radiation hardening. Figure 1.13 shows how increasing a single device's width rather than splitting it into fingers can reduce the total amount of silicon-oxide interface. For devices of equal effective width, the multi-fingered device in (a) requires more isolation oxide than (b) because of the additional gate material used to connect the two fingers. This assists against TID effects by reducing the oxide interface area.




Figure 1.13: Two variants of a MOSFET with $W_{\text{eff}} = 2W$. Purple is the gate with oxide—either for the channel or for isolation—underneath it, yellow is the source and drain of the device. The device with two fingers (a) has a greater total gate area with isolation oxide placed underneath it, due to the additional gate material required to electrically connect the two fingers.

[3] demonstrated how a common n-well for PMOS devices in inverter chains can quench SET-induced pulses. Simulating and measuring the pulse width and number of SETs measured at the output of an inverter chain, they noted a decrease in pulse width as well as a significantly reduced quantity of pulses in the chain with a common n-well. With separated n-wells comes greater spacing between devices, with the intermediate p-well doubling as a doped barrier to charge diffusion.

Guard rings are a common technique to prevent latchup in all integrated circuit design by reducing the impedance of wells which are nominally shorted. For single event latchup, [42] observed up to an order of magnitude increase in the charge threshold to trigger latchup

with the insertion of a dual n+ and p+ guard ring. This is because the guard ring provides a low impedance path between the bases of the parasitic BJTs which can initiate latchup, as well as providing a heavily doped region to act as a carrier sink while separating well edges and diffusion regions to reduce lateral transistor bipolar gain [20].

Differential signals reject common mode spikes on their output, so laying out a device where differential nodes are closely spaced increases the likelihood that ionizing radiation will affect both nodes and so be rejected as a common mode perturbation. Strategic resistor and capacitor placement and adjustment can low pass filter SETs. The addition of an RC low pass filter at the gates of the cross-coupled inverters of an SRAM cell can reduce the probability of an SEU, at the expense of speed [58], and similarly asymmetric variants can reduce the probability of bit flips [30]. Selectively increasing device size also has the benefit of additional filtering capacitance, though the benefit is slightly offset by the larger sensitive area and captured charge due to increased junction size.

Shifting more toward digitally-oriented techniques, logical masking prevents SEEs from propagating to critical nodes via combinatorial logic. A popular implementation of this is triple modular redundancy (TMR), where a block is triplicated and the final output is the majority of the triplicated blocks' results. Analogously, analog redundancy can average out error, though this can be at the expense of significantly increased power, noise, and area. For chains of digital logic where the area and power budget allows it, three-by-three voting can be used to ensure the input to any given triplicated section is always correct. For 3-to-1 voting in Figure 1.14a, the voting only hardens the logic against SETs within the Logic blocks. However, an SET which affects the output of any of the voter blocks will make the input of all three following Logic blocks incorrect, feeding the error into the downstream logic. For 3-to-3 voting like that in Figure 1.14b, an SET on the output of one of the voters will affect the directly connected Logic block, but the error will not propagate because only one of the three Logic blocks in a given stage was affected.




Figure 1.14: (a) 3-to-1 majority voting and (b) 3-to-3 voting. 3-to-3 voting ensures the input of all Logic blocks B and C are correct in the event of one SET (assuming good layout practice).

While not used in this dissertation, tools for triple modular redundancy generation [40] can quickly convert HDL modules into their TMR form. For placement and routing, [58] describes an interleaving method which can assist in reducing the routing overhead resulting from the additional cells.

Another form of spatial redundancy is dual interlocked cells (DICE) memory cells [13] (Figure 1.15).




Figure 1.15: The DICE memory cell, adapted from [13].

DICE cells are a specific latch topology which stores two copies of Q and \bar{Q} , such that a bit flip requires more than one output to be perturbed by a radiation-induced transient. Consider the cell in Figure 1.15 storing $Q = 0$. This means $Q\text{-}\bar{Q}\text{-}Q'\text{-}\bar{Q}' = 0101$. Suppose a single event transient perturbs Q such that it is temporarily brought high. This effectively

turns off P2, blocking propagation of the error to \overline{Q} . It also turns on N0, pulling \overline{Q} low, which turns off N3 and prevents the error from reaching Q'. As such, \overline{Q} and Q' are still correct, and so the positive feedback of the latch eventually corrects the initial error on Q.

1.2.3 Simulation and Testing

Pre-silicon simulation and post-silicon testing of radiation effects and events are necessary to design and confirm the radiation hardness of a device. The assertion that something is radiation hard must be accompanied by the specific, quantified radiation conditions under which it must operate, including distinctions between behavior in the presence of cumulative and immediate radiation effects. A generalized example operational requirement for total ionizing dose might look something like: “All components shall perform their required functions with no system-level degradation after exposure to two times the mission total ionizing dose (TID).” For single event effects, the qualification of SEE immunity might look something like: “For the purposes of part qualification, immunity to single event effects shall be demonstrated if all of at least two samples of the part do not exhibit the effect when each is exposed to a fluence of 10^7 ions with an equivalent LET (where valid) of at least $75\text{MeV cm}^2/\text{mg}$.” Some processes and devices such as wide bandgap FETs are snappily touted in marketing material as inherently radiation hardened. This often refers only to TID hardness, and a datasheet will have both the failure criteria and associated maximum dose to qualify the assertion of radiation hardness.

This subsection will address the validation techniques used for pre- and post-silicon validation of radiation hardness. As with all hardware, validation with silicon is the gold standard.

1.2.3.1 Pre-Silicon

TID Without the luxury of nicely pre-qualified devices where we are given specific limits on device performance up to a certain total ionizing dose, we resort to alternative methods of simulating radiation effects. To account for total ionizing dose, we treat it as an adjustment to “the usual” process corners (e.g. typical-typical, fast-slow, fast-fast, etc.). We unfortunately can’t provide the specifics of the model changes used in this dissertation, since nondisclosure agreements prevent us from giving any information on the underlying models of the PDK (including the types of model used). [29, 31] are only two examples of popular models which are used in industry today.

Threshold voltage shift can be treated as a modification of the traditional fast-slow corner with fast NMOS devices and slow PMOS devices, with the shift in threshold voltage increased to roughly match empirical values found in the literature, i.e. [21]. The vast majority of MOSFET models involve some concept of a threshold voltage with a constant base value, making the change simple for most models. Accounting for radiation-exacerbated mismatch is an even more straightforward matter of increasing the number of standard deviations considered versus a non-irradiated circuit. One consequence of the Central Limit

Theorem is that ensembles of random processes and variables (like those used in Monte Carlo experiments) converge to Gaussian distributions, so simply increasing the multiplier for standard deviations is statistically valid.

Adjusting the flicker noise corner can be more involved. Common back-of-the-envelope practice for determining the flicker noise power spectral density is expressed in Equation 1.7.

$$\text{PSD}_{1/f} = K_f \frac{I_D}{L_{\text{eff}}^2 C_{\text{ox}}} \frac{1}{f} \quad (1.7)$$

However, even not-quite-modern models often do not have a single constant, nicely abstracted parameter K_f , as it often changes with bias condition. [63] used the K_f abstraction when characterizing flicker noise pre- and post-irradiation and observed a dependence on drain current beyond the relationship described in Equation 1.7; after 10Mrad of irradiation, the NMOS devices in particular saw a factor of $4\times$ change in K_f with increasing drain current. Accurately emulating this relationship in a model can be complex depending on the model, so for our purposes we assumed a constant factor increase in the flicker noise power spectrum by a conservative value, estimated from empirical results in the literature [7].

An precise representation of increased off current like the profile seen in [35] can be highly nontrivial. Nevertheless, a relatively simple solution is to insert a current source in parallel with a typical MOSFET as in Figure 1.16, with the direction of the current flow determined by the polarity of the drain-source voltage.




Figure 1.16: Leaky MOSFET used to emulate radiation-exacerbated leakage. The sign and to an extent the magnitude of I_{leak} depends on the voltage measured by the voltmeter, clipped at a fixed value. The leakage also scales with device aspect ratio.

SEE SET emulation requires knowledge of the types and energies of ionizing radiation the system will be exposed to. For this, we look to the ions used in the 88-inch cyclotron at Lawrence Berkeley National Laboratory [41] for SEE testing.

To determine the charge collected from an SET, we use the linear energy transfer upon entry and the Bragg range. We calculate the energy E_{deposit} deposited within the charge collection depth z_{collect} with Equation 1.8

$$E_{\text{deposit}} = \begin{cases} E_{\text{original}} & \text{if } z_{\text{Bragg}} \leq z_{\text{collect}} \\ z_{\text{collect}} \times (\text{LET})_{\text{entry}} \times \rho_{\text{Si}} & \text{otherwise} \end{cases} \quad (1.8)$$

E_{original} is the particle's initial energy, z_{Bragg} is its Bragg range, and ρ_{Si} is the volumetric density of crystalline silicon. This is a rough estimate of energy loss, since linear energy transfer is not constant as the particle travels through the substrate. It is still nonetheless useful to gauge order of magnitude; SRIM/TRIM [88] could be used to obtain more precise values.

Silicon requires 3.6eV to generate an EHP, so we find that the worst case ion (bismuth in the 4.5MeV cocktail with an energy of roughly 904.16MeV) produces SETs with charge on the order of picocoulombs per micron of collection depth (Figure 1.17).

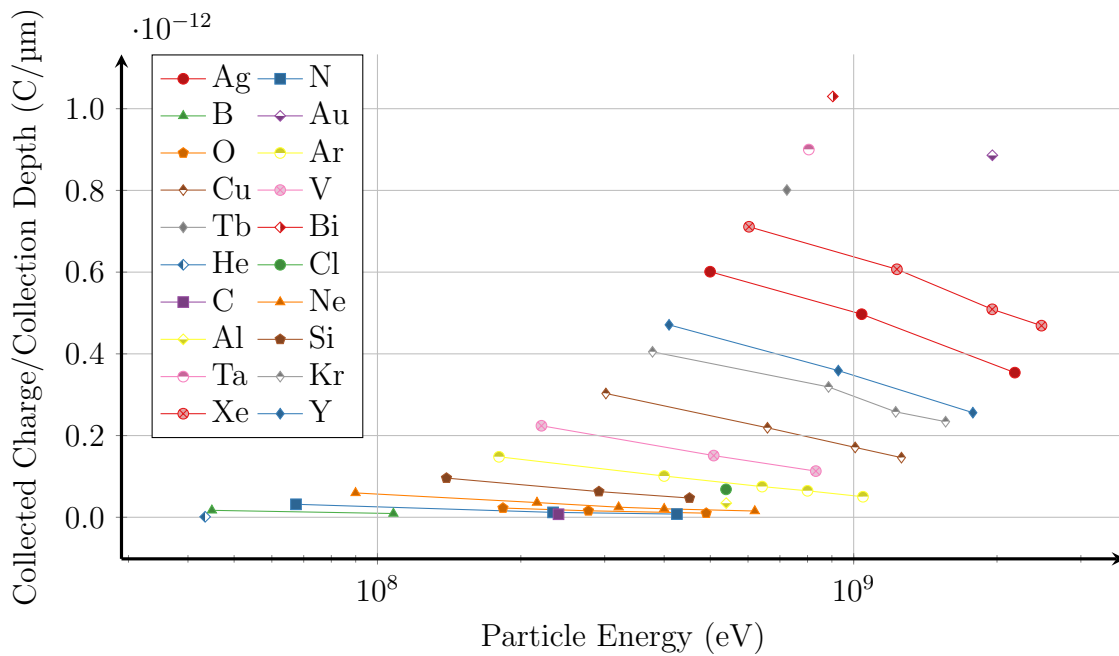


Figure 1.17: Estimated charge collected per micron of collection depth for an SET produced by the ions available in the cocktails at Lawrence Berkeley National Laboratory [41].

Using Equation 1.6 and empirical values for the drift and diffusion coefficients τ_r and τ_f [7], we can emulate an SET by injecting this current into the different nodes of our circuits to gauge their effect on the hardware's behavior (Figure 1.18)




Figure 1.18: SET emulator circuit. $\tau_r = R_r C_r$ and $\tau_f = R_f C_f$ for Equation 1.6. The input voltage source on the left provides a step function.

For Virtuoso 6.1.7 and beyond, Cadence includes a deepprobe component in the analogLib library which can traverse the hierarchy to access nodes internal to specific instances. This enables current injection into arbitrary nodes within cells from the testbench level, meaning no modification of the underlying cells is necessary for SET simulation.

1.2.3.2 Post-Silicon

TID testing follows the procedure outlined in Figure 1.19.

Oxide charge annealing is the process by which charge trapped in the oxide (not at interfaces) can neutralize over time. The neutralization annealing curve is independent of dose rate, and occurs by either tunnel annealing, where electrons tunnel from silicon into oxide traps, or thermal annealing, where electrons are emitted from the oxide valence band into oxide traps [25]. As with any process with underlying statistics, a large number of sample devices is ideal. In the absence of a high volume of samples, however, it's more practical to assume parts from a single wafer diffusion lot will have similar TID performance. [55] provides the specifics of the ESA-ESTEC Cobalt-60 facility for TID testing. SSL conducts their radiation testing at the Defense Microelectronics Activity at McClellan Air Force Base for TID (also Cobalt-60).

SEE testing involves exposing the device under test to high energy ions. The accepted figure of merit for SEE testing is the SEE cross section $\sigma \left[\frac{1}{\text{area}} \right]$, defined as the number of errors per ion fluence. The cross section is a function of a number of parameters—LET, flux, particle range, temperature, and operating voltage, to name a few—and there are many types of SEEs (Table 1.3) of varying severity. In particular, defining the safe operating

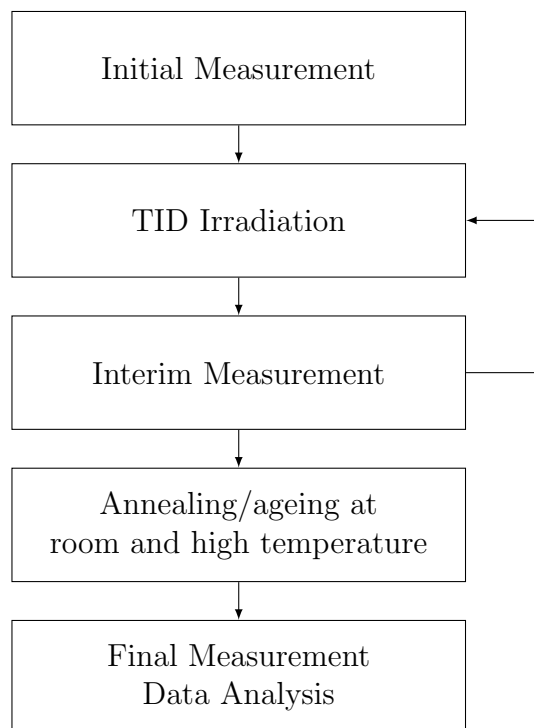


Figure 1.19: TID test flow. Adapted from [55].

area against hard errors, e.g. gate rupture, burnout, dielectric rupture, and others which can cause permanent device damage is critical to determine radiation hardness. For memory such as shift registers found in scan chains (Section 3.2.2.3), single event upset testing involves placing the registers in a known configuration, exposing the devices to the high energy ions for SEE testing, and then reading out the register values once more to ensure that they are correct. SSL conducts SEE testing at Lawrence Berkeley National Lab (Figure 1.17).

1.3 The Berkeley Analog Generator

The Berkeley Analog Generator (BAG) [14] is a Python-based framework which seeks to capture analog design methodology in a process- and specification-independent fashion to enable agile architecture exploration, fast iteration, and straightforward reuse. A number of others [71, 77, 78] have used BAG as a large component of their design flows and theses, fully integrating layout generation with complete measurement scripts to button-press generate design rule check (DRC) and layout-vs-schematic (LVS) clean blocks. Figure 1.20 shows an example BAG design flow, where a circuit designer tasked with a particular block can use prior generators and design scripts, or write their own if such scripts don't yet exist.

Setting up the infrastructure for schematic generation with BAG is doable in under

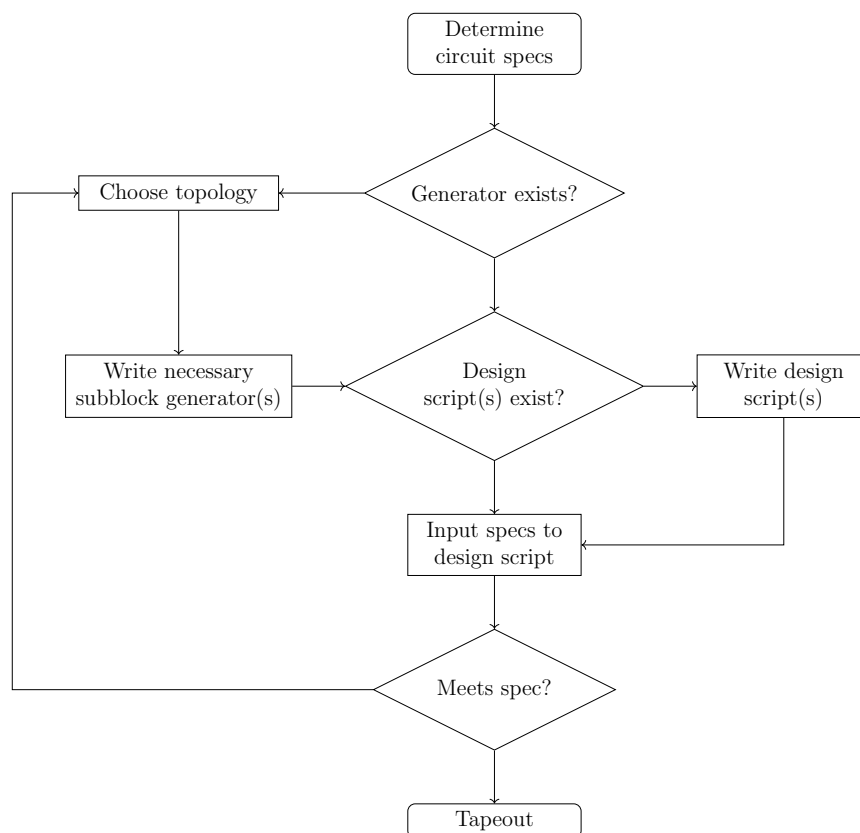


Figure 1.20: The BAG design flow for a block.

a day by one person (https://github.com/ucb-art/BAG2_cds_ff_mpt). Setting up the infrastructure for layout generation with XBase or LayGO is not. As such, BAG was used only at the schematic level for this dissertation. That said, the ease of reuse for prior scripts within BAG's framework, combined with the process-independent nature of design scripts if not generators, can save circuit designers substantial time; the chip in Chapter 3 began design using XFAB's XT018 6M process, then changed to TSMC 180nm 4 months prior to tapeout. This dissertation uses BAG2.0.

1.4 Acronyms

Acronym	Full Expansion
ASIC	application-specific integrated circuit
BAG	Berkeley Analog Generator
BSAC	Berkeley Sensor and Actuator Center
BWRC	Berkeley Wireless Research Center
CDR	clock and data recovery
CFD	constant fraction discriminator
DDD	displacement damage dose
DIBL	drain induced barrier lowering
DICE	dual interlocked cells
DMEA	Defense Microelectronics Activity
DRC	design rule check
EHP	electron hole pair
ESA	electrostatic analyzer
ITRS	international technology roadmap for semiconductors
LDO	low drop out (regulator)
LED	leading edge discriminator OR light-emitting diode
LVS	layout-versus-schematic
MCP	microchannel plate
MDS	minimum detectable signal
PDK	process development kit
PKA	primary knock-on atom
PWM	pulse width modulation
SEE/L/T/U	single event effect/latchup/transient/upset
SOI	silicon-on-insulator
SPAD	single photon avalanche diode
SPAN-I	Solar Probe Analyzer for Ions
SRIM	Stopping Range of Ions in Matter
SSL	Space Sciences Laboratory
STI	shallow trench isolation
SWaP	size, weight, and power
TDOA	time difference of arrival
TIA	transimpedance amplifier
TID	total ionizing dose
TMR	triple modular redundancy
TNID	total non-ionizing dose
TOF	time of flight
TOT	time over threshold

Chapter 2

Pulse Discrimination

Pulse discriminators are systems which only activate if the input signal meets some (potentially dynamic) threshold condition. Their function is a sort of event-to-digital trigger conversion, with output often taking the form of a digital pulse for use as a trigger in a larger system. This chapter discusses the various methods of pulse discrimination in the literature, their advantages and disadvantages in the context of SPAN-I's time-of-flight mass spectrometer, and ultimately arrives at an architecture suitable for our target use of space-based time-of-flight measurements. The sections are meant to be read in order to describe the thought process, but you can skip to Figure 2.9 for the end result and a summary of the benefits inferred by the architectural changes.

2.1 Topology Overview

At a low level, the pulse discriminator for SPAN-I must address two questions: first, is there a pulse; and second, when is the pulse? It is important to understand the nature of the inputs produced by SPAN-I. Chapter 1.1 discusses the operation of SPAN-I's TOF apparatus following its ESA. A previous iteration of SPAN-I [44] used a Z-stack MCP with nominal gain of $\approx 3 \times 10^7$. In an effort to reduce sensor dead time and increase particle throughput, SPAN-I changed to a chevron MCP with a reduced gain set to roughly 2×10^6 . Carbon foil yield changes with ion mass, with higher mass corresponding to a greater number of secondary electrons [48], and two different thicknesses of carbon foil are used for the START and STOP pulses. The result is input pulses which can be approximated with a double exponential similar to that of an SET in Equation 1.6, with the total charge Q ranging from 2Me^- to roughly 20Me^- , a rising time constant $< 1\text{ns}$ and a falling time constant on the order of $\approx 1\text{ns}$.

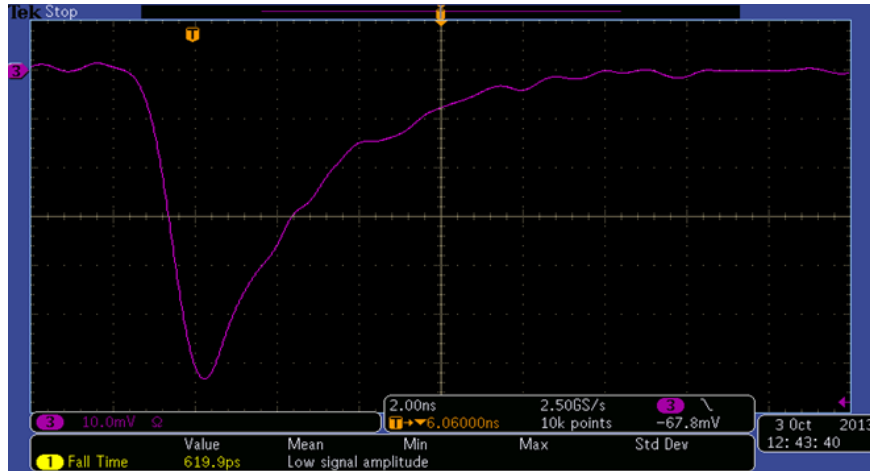


Figure 2.1: MCP output, converted to a voltage and averaged over a periodic event trigger.

Table 2.1 shows the target specification for both the previous version of SPAN-I which went on the Parker Solar Probe, as well as the target values for the hardware in Chapters 3 and 4.

Parameter	APL Chip [44]	Target
Mission(s)	Parker Solar Probe	Mars EscaPADE
Minimum Detectable Signal	40Me ⁻	2Me ⁻
Maximum Event Rate	<1Mevent/s	10Mevent/s
Signal Chain Integration	ASIC + Discrete	ASIC
Timing Walk w/o Shaping	< 100ps	600-800ps total
Jitter w/o Shaping	< 100ps _{rms}	
Timing Walk w/ Shaping	N/A	600-800ps total
Jitter w/ Shaping		
SEU Tolerance	Immune	Immune
TID Tolerance	100krads	100krads
Power	3-4mA _Q	3mA
ASIC Area	< 1mm × 1mm	< 2.5mm × 2.5mm
Process	TSMC 250nm CMOS	TSMC 180nm CMOS

Table 2.1: Measured parameters for SPAN-I’s previous chip designed by Johns Hopkins APL, and the target specification for the iteration of SPAN-I discussed in this dissertation. Area limits were set by available area on the BSAC shuttle. The 180nm process was chosen because BSAC does not need to pay for the area.

The simplest pulse discriminator which directly answers the question of a pulse’s existence is the leading edge detector (with the rather confusing acronym LED), where an input pulse is

compared against a static threshold. However, for systems which do not produce effectively digital or extremely uniform event pulses, this introduces timing walk, where the output timing shifts with pulse amplitude. Amplification is sometimes sufficient to mitigate this problem—this is what OOK communication links do—though it cannot remove it entirely. Even intentionally slewed amplifiers to fix the constant slope do not resolve this entirely due to finite bandwidth shifting the time at which the slewing begins [36].

Walk compensation methods abound, all of which require acquiring some additional information about the input. [26] simulated a front end with two LEDs at different thresholds, with the time between the two triggers used to estimate the input slope. The slope is related to the amplitude of the input, which then allows for walk correction by normalizing the input. This requires one time-to-digital conversion per pulse, and two for a TDOA measurement. [84] expands upon this with a more rigorous examination of the relationship between event energy (related to pulse amplitude) and walk, enabling more precise and accurate compensation than the linear function used in [26], though every pulse’s timing is independently determined. While this fitting method was effective for their chosen application of PET, it is ultimately a calibration scheme which needs adjustment for every sensor and application. [57] does not explicitly read out pulse slope or amplitude, but instead provides time over threshold (TOT) information, which can be used to determine pulse width—not applicable to SPAN-I—as well as provide an estimate of pulse amplitude. Once again, more than one time-to-digital conversion is needed for a single TOF measurement.

Constant fraction discrimination (CFD) is another popular [1, 72, 18, 44] method for pulse discrimination which triggers when the analog input pulse reaches a constant fraction f of its peak A (often shifted in time by a known constant value). The result is an output trigger whose timing does not depend on the amplitude of the analog pulse for (theoretically) zero timing walk. CFD operation can be generalized with the system in Figure 2.2




Figure 2.2: Block diagram of a common generalized CFD, implemented with two LTI operations $H_+(s)$ and $H_-(s)$ along with an ideal comparator.

In the time domain, the output is described by Equation 2.1

$$y = \text{sgn}[x * (h_+ - h_-)] \quad (2.1)$$

With this system, scaling x by a positive value C does not change the output y . That is, given two inputs x_1 and x_2 where $x_1 = Cx_2$, $C > 0$, $y_1(t) = y_2(t)$ —the timing of the edges

at the output do not depend on the value of C . This means that for a given pulse, the timing of the output edge is independent of pulse amplitude. This works for *any* two linear time invariant (LTI) operations with no need for any time- or analog-to-digital conversion to gain more information about the pulse, though it is common to include an LED in parallel to distinguish pulses from noise. It is worth noting that the majority of practical CFD implementations have inherent walk due to deliberate offset in the comparator to guarantee a low output for the zero-input condition. [36] addressed this by injecting a step equal to the offset magnitude into the pulse to cancel this upon a pulse's arrival.

It would be remiss not to acknowledge that all of these discrimination methods are easily implemented in the digital domain [22, 16]. For applications and SWaP budgets where such an ADC is feasible, immediate digitization scales well for more complicated signal processing algorithms. However, driving a single minimum size transistor in a 180nm CMOS process from a 3.3V supply at 1GHz consumes nearly $0.5\mu\text{W}$. An ADC with the requisite resolution and sampling rate would be impractically power-hungry for our purposes and made difficult by the fact that transistor f_{max} and f_T are in the tens of gigahertz [32].

2.2 Timing Walk

A standalone LED is inadequate to meet the walk requirements of Table 2.1 given the constraints of the 180nm process. Suppose we have a preamplifier with a transimpedance-bandwidth product of 3×10^{10} . Setting the threshold higher than 2LSB of a 9-bit DAC with a 3.3V full scale range like that of the APL ASIC consumes the entire combined walk and jitter budget even with ideal comparators (Figure 2.3).

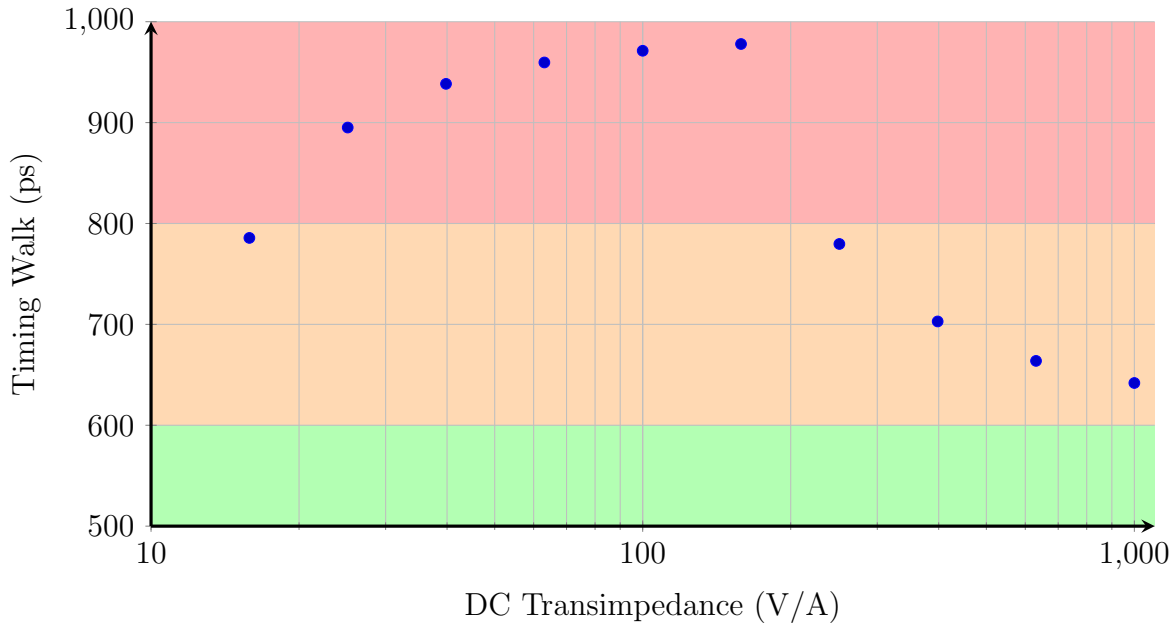


Figure 2.3: Plot of timing walk given 2-20Me⁻ pulses with $\tau_r = 500\text{ps}$, $\tau_f = 1\text{ns}$ to roughly match Figure 2.1, with a threshold set at $2 \times 3.3\text{V}/2^9 \approx 12.9\text{mV}$ for a transimpedance amplifier with a fixed gain-bandwidth product of 30GHz V/A. Increasing the transimpedance sees the timing walk asymptotically approach 611ps. This is the bleeding edge of what the process node can achieve under nominal operating conditions; decreasing the gain-bandwidth product to 10GHz V/A to account for variation in process, supply, and temperature makes it so even 1LSB of the 9-bit DAC is insufficient to meet walk requirements.

While the walk compensation methods described in Section 2.1 would no doubt improve this, they require additional time-to-digital conversions on top of algorithm implementations on a radiation hardened FPGA with already-constrained resources.

Considering the other popular means of pulse discrimination, CFDs are theoretically capable of achieving zero timing walk, and so we examine our options in the context of SPAN-I. Beyond the those listed in Table 2.1, there are several requirements:

- avoid triggering on noise
- < 2 output triggers per event
- monotonically increasing count rate vs. event rate
- tunable trigger fraction

The first of the list—distinguishing valid pulses from noise—is achievable with an LED, so we use one in parallel with a CFD (Figure 2.4).




Figure 2.4: High level concept of CFD usage in parallel with an LED. Here, the LED determines if there is a pulse, and the CFD provides the timing of the pulse.

2.3 Afterpulse Rejection

Afterpulsing is a non-ideal behavior where the response of an amplifying sensor element, e.g. MCP, SPAD, PMT, to a single event produces a secondary pulse after the initial response. In MCPs, this is due to the ionization of gas molecules which then drift back to the channel input to trigger an additional pulse [43].

To reject afterpulsing, we place a non-retriggerable one shot pulse generator—also called a monostable multivibrator—at the output of the system. One shot pulse generators are a class of circuit with one stable state. When the circuit is perturbed out of that stable state, it takes a fixed amount of time to return to the stable condition. A common use of these is to produce an output pulse of fixed duration t_{1shot} in response to an input trigger, e.g. a rising edge. So long as the duration of the one shot pulse is longer than the time it takes an afterpulse to appear and settle, the downstream logic will see only one output pulse associated with the event, ignoring the afterpulse. As an upper bound on t_{1shot} , we have the sum of the minimum time between events and any “hold” time required for any transient responses within the one shot to settle.

We chose a non-retriggerable circuit to prevent the output from “locking” high in the event of an unexpectedly high event rate. A non-retriggerable one shot does not respond to triggers which occur while the one shot is not in its stable state. By contrast, a retriggerable one shot will effectively restart the timer for t_{1shot} for every event which occurs while the one shot is unstable (Figure 2.5).

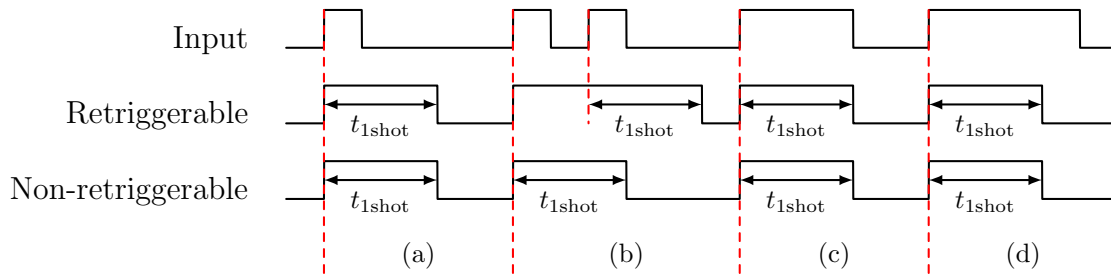


Figure 2.5: The response of a retriggerable one shot and a non-retriggerable one shot to (a) a single pulse with duration shorter than t_{1shot} (b) two pulses whose rising edges occur within t_{1shot} of each other (c) a pulse with duration equal to t_{1shot} (d) a pulse with duration longer than t_{1shot} . Each one shot produces HIGH pulses in response to rising edges.

Consider an extension of Figure 2.5b where the input is a signal where every rising edge, valid or otherwise, is less than t_{1shot} apart. A retriggerable one shot will register only a single long pulse at its output, which the downstream hardware can only interpret as a single event. In other words, as the event rate increases, the output trigger rate will initially increase, then fall once the timing between events is short enough to retrigger the one shot. This means that sudden fast bursts of input events can be almost entirely missed with no way of distinguishing it from single events. A nonretriggerable device, however, will simply reach some maximum output pulse rate in a manner akin to dead time on a sensor, which it will maintain even as the input event rate increases. Thus, a nonretriggerable one shot is necessary to maintain a monotonically increasing count rate versus event rate.

2.4 Constant Fraction Discrimination

This section will cover the specifics of various CFD architectures in the literature, their utility in radiation-ful environments, and their calibration requirements, finishing with a description of the architecture used in Chapters 3 and 4.

Equation 2.1 states that constant fraction discrimination can be achieved when H_+ and H_- are linear and time-invariant. While this condition is merely sufficient, not necessary, to make a CFD, it is nevertheless a useful starting point.




Figure 2.6: The core methods behind prior implementations of constant fraction discriminators.

One CFD method triggers exactly at the peak of a pulse when its derivative changes sign [72] (Figure 2.6a). This lacks flexibility—the system can only trigger at the first maximum which appears after the pulse is sufficiently large. Furthermore, spikes from differentiation (or any high pass element) risks railing the output and reintroducing walk in the absence of additional feedback. [18] uses the Nowlin Method and inserts a programmable attenuator onto the comparator’s inverting input (Figure 2.6b). While this adds a degree of tunability, it still does not resolve the problems with high pass elements. Replacing the high pass element with a low pass (Figure 2.6c) or all-pass element like a delay (Figure 2.6d) reduces the risk of railing. The delay-versus-attenuate CFD is well-established and widely used [1, 44, 28] technique which has both tunability and reduced risk of accidental nonlinearity.

One issue which all but the zero-derivative crossing methods encounter is a shape-dependent trigger fraction. Consider two pulses with different shapes—that is, they are not scaled copies of one another—passed through the frontend in Fig. 2.6d. Figure 2.7a shows that for each pulse, the fraction at which the system triggers is not only not the same as the attenuation factor f , but it varies depending on the shape of the pulse. While it is reasonable to expect the MCP to have consistent pulse shapes between paired START and STOP pulses—and so will not introduce walk for a double coincidence measurement—robustness against differences in pulse shape across ion flavor [48], MCP degradation and spatial variation [65], and carbon foil damage can degrade system performance.

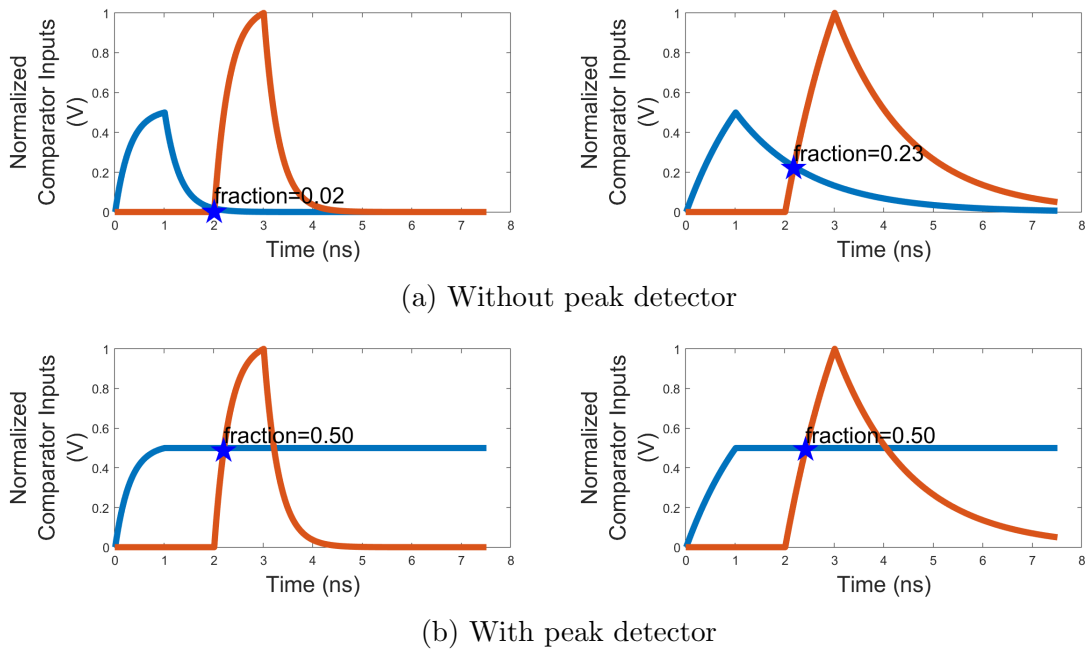


Figure 2.7: Example inputs to the CFD comparator, with (b) and without (a) the peak detector inserted in the shaping chain. The orange line is the delayed input pulse, the blue line is the input pulse, attenuated by a factor $f = 0.5$.

To maintain a constant fraction trigger irrespective of pulse shape, we insert a peak detector prior to the attenuator in the CFD (Figure 2.8).




Figure 2.8: The modified CFD with the peak detector added before the attenuator.

The peak detector implements a nonlinear function by holding the maximum value of its input—which can but won't be used for determining pulse amplitude—resulting in Figure 2.7b. As long as the delay t_d is greater than or equal to the rise time of the pulse (minus the time it takes the pulse to reach the fraction f), the CFD trigger fraction will always be f , and the trigger time will always be the sum of the time it takes the initial pulse to reach the fraction f of its max and t_d (Equation 2.2).

$$t_{\text{CFD}} = t_d + t_{\text{frac}} \tag{2.2}$$

We now have several constraints on the lower bound of t_d :

$$t_d \geq t_{\text{LED}} \tag{2.3a}$$

$$t_d \geq t_{\text{rise}} \tag{2.3b}$$

Equation 2.3a ensures that the pulse discrimination comes from the CFD, while Equation 2.3b guarantees a constant trigger fraction f regardless of pulse shape. If 2.3a is not satisfied, the system will behave as an LED with all its associated timing walk (Section 2.2). Otherwise if 2.3b is not satisfied, the system will behave as a conventional delay-versus-attenuate CFD (Figure 2.6d).

The upper bound of t_d depends on the logic which combines the LED and CFD branch outputs. The simplest solution would be to take the logical AND of the LED and CFD comparators. In this form, however, the upper bound on t_d is now quite tight—it cannot be long enough that the CFD triggers after the LED deasserts. Placing a latch on the output of the LED comparator will hold the LED output high until the CFD triggers. Alternatively, we can take advantage of the memory of the peak detector and connect the output of the peak detector to feed into both the CFD and LED branches as in Figure 2.9.

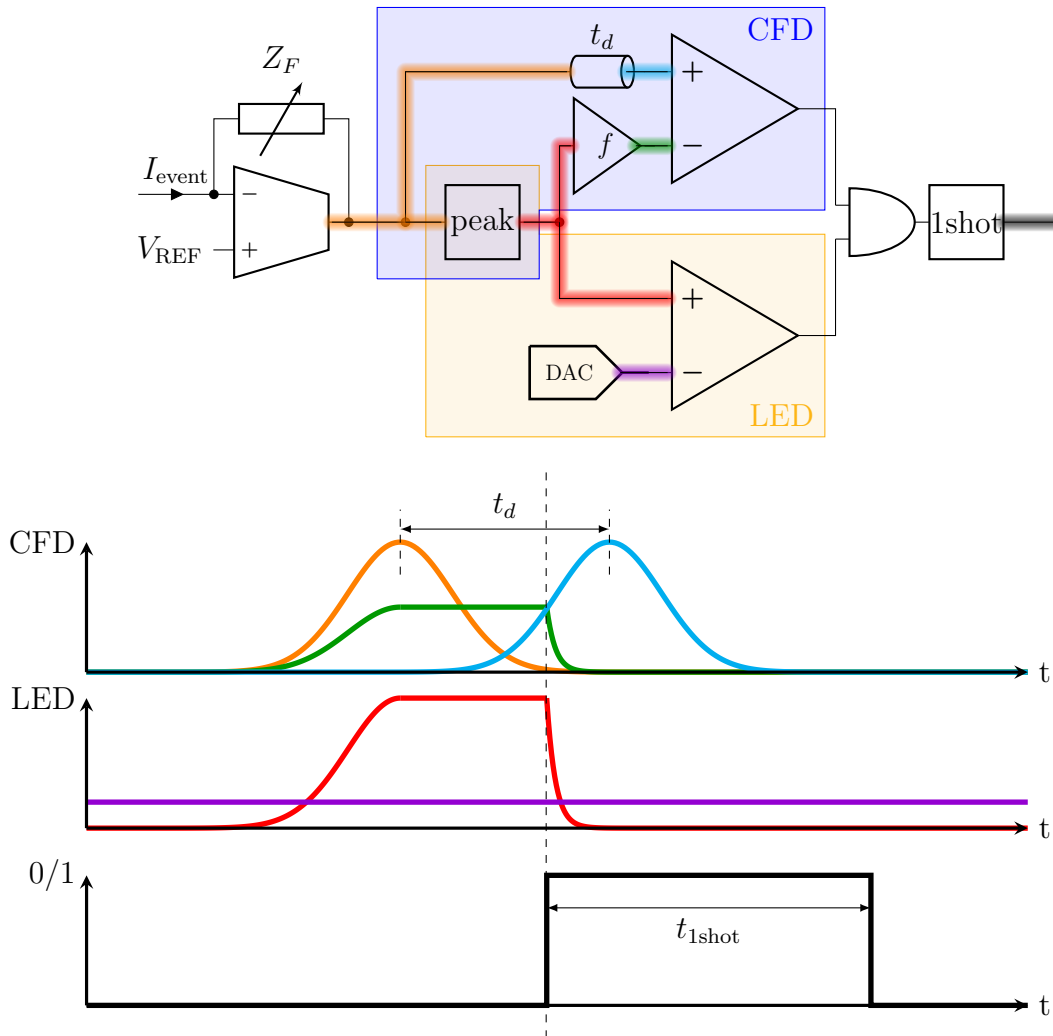


Figure 2.9: Block diagram of the front end and its operation with the CFD branch outlined in blue and the LED branch outlined in yellow. The one shot output is used to reset the peak detector

Now the upper bound on t_d is defined roughly by the time to the next pulse and the output one shot timing, and more precisely by Equation 2.4

$$t_d \leq t_{\text{LED,next}} - t_{\text{1shot}} \tag{2.4}$$

where t_{1shot} is the output pulse length of the one shot pulse generator described in Section 2.3. The peak detector is reset using one shot output, and with that the architecture satisfies the requirements listed in Section 2.2 and comes with a(n unused) built-in pulse amplitude hold for potential digitization for signal processing algorithms.

2.5 SEE Watchdog

Because the peak detector has memory, SEEs which occur within the peak detector can be held on its output. These result in output transients which have no correlation to the signal input to the peak detector. Sufficiently large transients can “stick” the CFD comparator low, permanently disabling the front end (Figure 2.10).

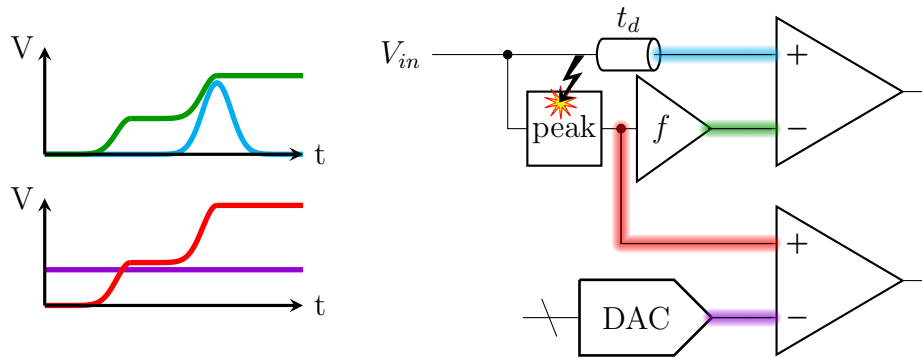


Figure 2.10: Diagram with one possible scenario for SEE-induced lockout. In general, lockout can occur if an SET on the peak detector raises the output of the attenuator to a level that real pulses can never reach. The CFD output remains low, and the system never resets the peak detector until the chip is reconfigured.

We combat this with a watchdog monitoring the digital outputs of the LED and CFD (Figure 2.11).




Figure 2.11: The peak detector’s SEE detection and correction watchdog circuit and operation in the event of an otherwise lock-inducing transient. (1) The peak detector experiences an invalid transient which causes its output to trigger the LED, starting the LED_1shot timer. (2) After t_{stuck} , if the CFD has not registered an event, `rst_stuck` asserts, (3) resetting the peak detector (along with the LED and CFD outputs).

Because the `rst_stuck` signal effectively introduces a dead time of t_{rst} , it is important to ensure that t_{rst} is not so long that watchdog resets overwhelm any actual signals. This depends on the single event rate within the sensitive area of the peak detector and will be addressed numerically with the real area of the device in Section 3.2.3.

Chapter 3

Chip V1

This chapter describes the internals, simulation results, and limited measurements taken from the first chip taped out for this project. It is not intended as a user manual. For detailed documentation of file locations, I/O, scan bits, etc., see Appendix E.

3.1 Chip Summary

The chip (Figure 3.1) was taped out on July 21, 2021 in TSMC’s 180nm process through the Berkeley Sensor and Actuator Center. The total chip area including pads and seal ring was $1.6\text{mm}\times 1.7\text{mm}$. The chip contains on-chip power management and reference generation, derived from a single 3.3V supply. The full front end is the CFD shown in Figure 2.9 and is referred to as the “main” or “full” chain. A pared-down front end with no on-chip pulse shaping is marked “shortened,” and is referred to as the “small” or “no-shape” chain. This small chain is similar to the APL chip used in [44] and is meant to be mutually exclusive with the main chain. The scan chain for configuring the chip is marked “SPI and Config” and consumes a significant portion of the chip area. Lastly, we included bandgap reference circuit and peak detector as standalone test structures, both contained in a dedicated power domain separate from that of the main and small signal chains.




Figure 3.1: Chip die photo with main structures annotated. Photo courtesy of Hani Gomez.

Table 3.1 provides a snapshot of the chip’s overall performance in relation to the target application. The remnant timing walk can be attributed to the limited bandwidth of the comparators, and will be discussed in greater detail in Section 3.2.2.2. Steps were taken in the second iteration of the chip to rectify this, and will be discussed in Chapter 4.

Parameter	APL Chip [44]	Target	Chip V1
Minimum Detectable Signal	40Me ⁻	2Me ⁻	8-10Me ⁻
Maximum Event Rate	<1Mevent/s	10Mevent/s	Simulated ≥ 10Mevent/s
Signal Chain Integration	ASIC + Discrete	ASIC	
Timing Walk w/o Shaping	< 100ps	600-800ps total	682ps
Jitter w/o Shaping	< 100ps _{rms}		177ps
Timing Walk w/ Shaping	N/A	600-800ps total	601ps
Jitter w/ Shaping			743ps
SEU Tolerance	Immune	Immune	Simulated
TID Tolerance	100krads	100krads	Simulated
Power	3-4mA _Q	3mA	2.9mA
ASIC Area	< 1mm × 1mm	≤ 2.5mm × 2.5mm	1.6mm × 1.7mm
Process	TSMC 250nm CMOS	TSMC 180nm CMOS	

Table 3.1: Chip V1 versus the target specifications.

All circuits were required to maintain consistent performance across the typical space-qualified temperature range of -55°C to 125°C with SEU immunity as well as TID hardness up to 100krads, and were simulated accordingly. Unfortunately, we could not conduct radiation hardness testing before the chips were returned to TSMC. The chips were returned to TSMC in January of 2023, with the BSAC administrative office handling customs and export control.

3.2 Design and Measurements

This section will describe the design considerations of the chip, as well as hardware and software setups used to test Chip V1 and show the data extracted from it. All tests in this section used a Teensy 3.6 development board for interfacing with the chip and board, and a Keysight E3631A DC power supply to provide the 3.3V high voltage supply from which all other voltages are derived.

3.2.1 Power and Reference Generation

With the exception of the electrically isolated test structures, the entire chip was designed to operate from a single 3.3V supply like that provided by SPAN-I's power management [44]. The scan chain, full/main front end, shortened/no-shape/small front end, and test structures operate on separate power domains, with the first three internally regulated down to a core

1.8V from the single 3.3V supply. The decision for separated 1.8V domains was several-fold: First, the small and full signal chains are mutually exclusive and were not designed to be used simultaneously—separating their supplies enabled separate current measurement and more isolated operation. Thus, we needed to provide a means of enabling/disabling the on-chip regulators. Second, the scan chain must always be on to be able to configure the chip. And finally, the choice of 1.8V internal operating voltage was because early characterization of the process’s 5V devices showed us that a signal chain with a 3.3V supply would be infeasible.

All internally regulated nodes were padded out for measurement and the potential for external override. Every reference voltage and current was derived from a single bandgap circuit [54] with a simulated DC supply rejection of 17.8dB with the 3.3V supply. The bandgap circuit was designed using the standard procedure of canceling temperature coefficients, which was then codified for use with BAG. As a means of sanity checking, the bandgap was duplicated as a standalone test structure with its output voltage padded out. For characterization, the device was placed in a TestEquity Model 107 temperature chamber which swept from 0°C to 80°C over the course of several hours while the bandgap voltage was measured with a Teensy 3.6 with 16B/13ENOB analogRead resolution. The reduced temperature range versus the true space qualification range was a limitation of the hardware on hand. The temperature was read from at least one of three possible sources: the temperature chamber via the backside RS-232 port, the Teensy internal temperature readout, or a TMP102 digital temperature sensor.




Figure 3.2: Bandgap voltage versus temperature with the envelope of the standard deviation. Spikes at low temperatures were from condensation within the chamber.

Figure 3.3 shows the reference routing network with the various LDOs, specifically how

the reference current from the bandgap circuit was mirrored and distributed across the chip, with voltages then derived using resistive DACs.




Figure 3.3: The reference routing network for power distribution. There are three LDOs on the chip, all implemented in a similar fashion.

The voltages generated by the resistive DACs for the LDOs were designed to span 1.75V to 2.05V with $10\mu\text{A}$, set via the scan chain. We recognize that the upper range extends past the stated acceptable operating voltage of the process core devices; this was done to guarantee performance across the full range of process and temperature corners.

The LDOs were designed with the assistance of the Berkeley Analog Generator, with the basic DC operating script forming the basis of the more rigorous design seen in [56]. For this chip we were forced to use a PMOS series device rather than NMOS to supply adequate current within the available silicon area.

Unlike prior work with the Single Chip Mote (Appendix A.1), this chip was allowed external decoupling capacitance. With a board-level 0603 nominally 10nF ceramic capacitor for each supply pad, we measured $< 5\text{mV}_{\text{pk-pk}}$ supply bounce for VDDSMALL and VDDMAIN in the presence of an incoming event pulse, and the same for VDDAON while scan was being programmed. This was measured using a Tektronix DPO70000DX series oscilloscope. We did not perform a full characterization of the DAC tuning range for the supply voltages to the same extent as the DACs in the signal chain, though we did confirm a voltage range of roughly 1.76V to 2.13V for the minimum and maximum settings.




Figure 3.4: The low dropout regulator used within the chip. The enable device (green) is not included in the always-on regulator, and the reference voltage V_{REF} is produced using an on-chip bandgap reference circuit.

3.2.2 No-Shape/Small Signal Chain

The no-shape or small signal chain (Figure 3.5) is a rework of the APL ASIC used in [44]. It is a subset of the full signal chain in Figure 2.9, with no integrated analog shaping beyond the voltage DAC used in the leading edge detector; all components in the small signal chain also appear in the main/full signal chain (Section 3.2.3).




Figure 3.5: Block diagram of the small/no-shape signal chain.

The small signal chain was operated as a conventional delay-versus-attenuate CFD described in Equation 2.1. Figure 3.6 shows the setup used to stimulate and read out timing

information from the small signal chain.




Figure 3.6: Measurement setup and procedure for gathering timing statistics for the small signal chain. (a) Configure the DG535, TDC, and chip scan chain. (b) Trigger the DG535. The DG535’s triggered output is then used as the START event to the TDC. A subsequent pulse nominally 1 μ s wide with an amplitude anywhere from 50mV to 600mV is routed down three paths: an attenuator (Kay Elemetrics 839); a coaxial cable roughly 60cm longer than that of the attenuator for an additional \approx 2ns delay; and directly to the positive input of the LED comparator. Each pulse amplitude test is performed 500 times with at least 100ns between pulses. (c) The chip output is latched and level shifted from the 1.8V core voltage to 3.3V to be fed into the TDC as a STOP event.

Each pulse amplitude test was repeated 500 times with at least 100ns between pulses.

To account for timing walk and jitter introduced by components other than the chip, calibration measurements were taken from boards without the chip populated. For jitter, we used a 1.8V output of the DG535 to drive the latch and level shifter which the device under test would otherwise be connected to. For timing walk, we used a Tektronix DPO70000DX series oscilloscope to approximate the relative shift in pulse peak times for different voltage amplitudes; somewhat ironically, we could not rely on a digital trigger or threshold to determine the walk of the DG535 and off-chip shaping components. Figure 3.7 shows the jitter and time difference of arrival (TDOA) of the output pulse from the chip after the contributions from the board were accounted for (Equations 3.1 and 3.2).

$$\sigma_{\text{measured}}^2 = \sigma_{\text{chip}}^2 + \sigma_{\text{shaping}}^2 + \sigma_{\text{DG535}}^2 + \sigma_{\text{board}}^2 \quad (3.1)$$

$$\Delta t_{\text{measured}} = \Delta t_{\text{chip}} + \Delta t_{\text{DG535}} + \Delta t_{\text{shaping}} \quad (3.2)$$

The simulated mean current consumption of the small signal chain with an event rate of 10Mevent/s was 1.76mA.

3.2.2.1 LED DAC

The DAC for the leading edge detector is intended as a threshold to distinguish incoming pulses from noise. Its code is set via the scan chain. It uses the resistive ladder shown in




Figure 3.7: (a) Jitter and (b) time difference of arrival of the measured pulses. Worst case jitter was measured at $176.9\text{ps}_{\text{rms}}$, and timing walk at 682ps .

Figure 3.8, selected for simplicity, essentially guaranteed monotonicity, and—at the expense of large area—low power consumption. The chosen mux architecture requires no additional encoding for binary selection. The DAC operates in constant voltage mode, where V_{DD} is the signal chain’s supply. The ladders contain $2^9 = 512$ elements, of which the mux connects to 256 for an 8-bit DAC. The 256 elements were chosen to begin at the seventy first element of the resistive ladder so a code of 128 corresponds to an output voltage of 700mV assuming a 1.8V supply.

The DAC transfer function was obtained by programming the scan chain to set the supply voltage and DAC code, then repeatedly measuring the output of the DAC with the Teensy’s `analogRead` (13ENOB over 3.3V for an ADC LSB of $\approx 403\mu\text{V}$). The DAC sees a capacitive load and is a DC signal, so the only real considerations for resistor value of $3.5\text{k}\Omega$ were area consumption and load leakage, which includes any external decoupling capacitance. The DAC’s simulated quiescent current consumption was $1\mu\text{A}$; it was kept intentionally low to budget more power for other components. Like the internally regulated supply voltages, the DAC output was padded out and connected to nominally 10nF of 0603 ceramic decoupling capacitance for the measurement. Equations 3.3 and 3.4 were used to calculate the differential and integrated nonlinearity of the DAC while ignoring gain error and offset.

$$\text{DNL}[k] = \frac{\text{step}[k] - \text{step}_{\text{avg}}}{\text{step}_{\text{avg}}} \quad (3.3)$$





Figure 3.8: (a) The resistive ladder DAC was chosen for its simplicity and guaranteed monotonicity. (b) The analog mux with N bits was constructed with $2^N - 1$ two-to-one muxes to enable direct feed of binary selection bits with no additional encoding.

$$\text{INL}[k] = \sum_{i=1} \text{DNL}[k] = \frac{V_{\text{DAC}}[k] - V_{\text{DAC,uniform}}[k]}{\text{step}_{\text{avg}}} \quad (3.4)$$


3.2.2.2 Comparators

At their core, all comparators were chains of fully differential stages (Figure 3.10) with low gain and high bandwidth, with conversion to a single-ended output reserved for the end of the chain.


A PMOS topology was chosen to mitigate residual timing walk; comparator speed will generally decrease with increased pulse amplitude because of the upshift in the comparator's input common mode and the reasonably static tail current source. Effort was made to place differential signals close to one another in the layout to increase the likelihood of an SEE affecting only the signal common mode. We opted to forgo common mode feedback in the name of power savings, relying on the tail current and resistors to set the bias point of each




(a)



(b)



(c)



(d)

Figure 3.9: (a) Voltage DAC transfer function with the supply voltage set to its lowest value. Each data point is the average of 100 measurements. The gain is 3.44mV/LSB for a full scale range of 877.22mV . (b) The DAC's RMS noise $\leq 1.17\text{mV}$, or 2.9LSB . Measured with nominally 10nF of decoupling capacitance on the output of the DAC. (c) DNL min/max $-0.17/0.23\text{LSB}$. (d) INL min/max $-0.19/0.31\text{LSB}$.

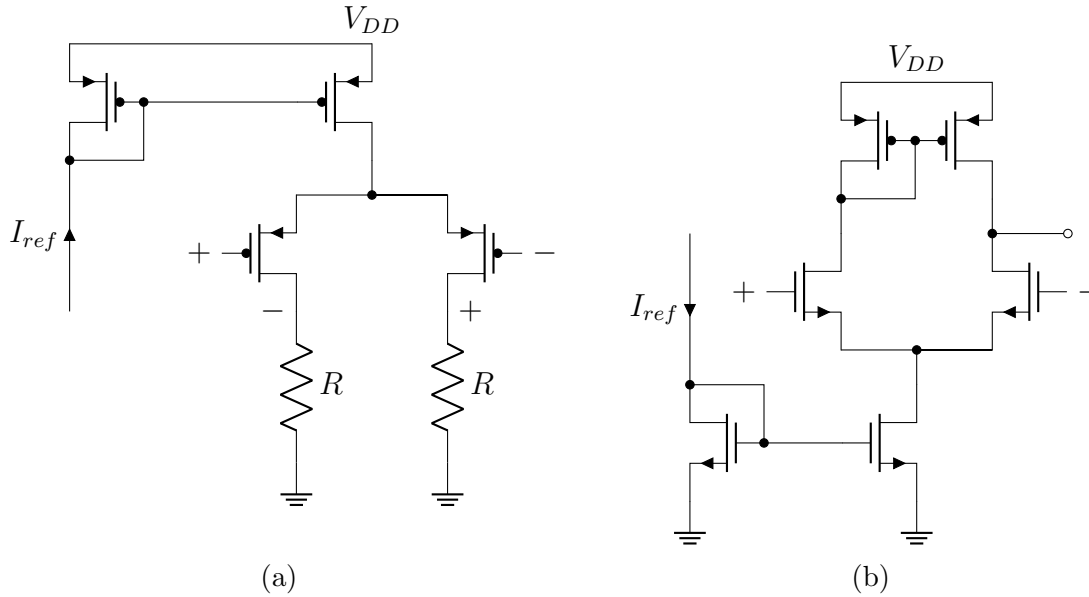


Figure 3.10: The core of every comparator consisted of several cascaded fully differential stages (a) and one final stage for a single-ended conversion (b).

stage. Supposing that each stage is identical and behaves in the small signal as a single pole system with a roughly constant unity gain frequency of $\omega_u = A_0\omega_0$ where $A_0 = g_m R$ is the DC gain and $\omega_0 = \frac{1}{RC}$, Equations 3.5 and 3.6 give the optimal number of stages and their individual gains to minimize delay in response to a step.

$$N_{\text{opt}} \approx \ln(A_v) \quad (3.5)$$

$$A_{0,\text{opt}} \approx e \quad (3.6)$$

From this, at least 6 low gain stages would be necessary for a single LSB of the LED voltage DAC to reach the supply voltage of 1.8V.

Offset in the CFD comparator introduces timing walk by adding a voltage error which does not scale with the input amplitude. To reduce the offset, we added autozeroing with a nulling amplifier (Figure 3.11), where the offset is sampled ϕ_1 when the chip is reconfigured, and the scan chain LOAD signal is toggled. Equation 3.7 is an expression for the output of the overall comparator after autozeroing, where A is the gain of an amplifier, B is the gain of the amplifier from its nulling port, and V_{OS} is the input-referred offset of the amplifier.

$$V_{\text{out}} = V_{\text{in}} (A_{\text{main}} + B_{\text{main}}A_{\text{null}}) + V_{\text{OS,main}}A_{\text{main}} + V_{\text{OS,null}} \left(\frac{B_{\text{main}}A_{\text{null}}}{1 + B_{\text{null}}} \right) \quad (3.7)$$

The nulling pin was added to both amplifiers in a rather brute force way, with the black differential pair in Figure 3.12 incorporated into only the first fully differential stage of each

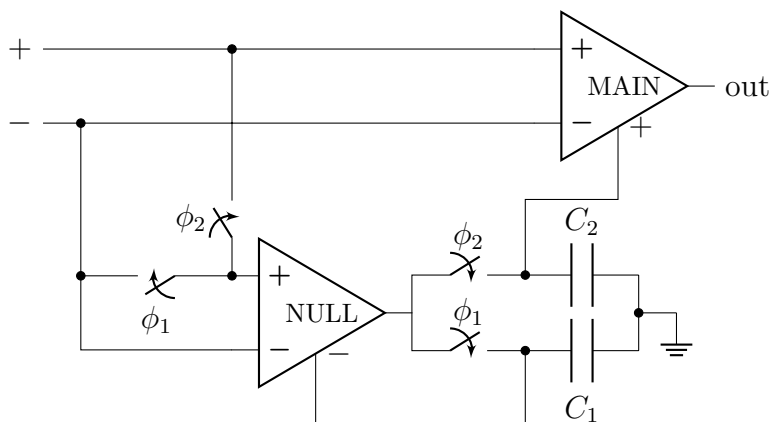


Figure 3.11: High level diagram of the autozeroed comparator. Because there is no clock, the sampling phase ϕ_1 occurs when the chip is reconfigured, i.e. the scan chain is LOADED into the rest of the chip. In our use, the nulling amplifier and the nulling pins are differential; we show the single ended variant here for clarity.

amplifier. There were no mismatch models in this PDK, so input referred offset was estimated at $\approx 5\text{mV} \times \mu\text{m}$ for a single transistor. Applying this in the worst case configuration to all the fully differential stages, the offset control was then designed such that the autozeroing was capable of correcting at least thrice the estimated error, referred to the output of the first stage.

The nulling amplifier was composed of 4 fully differential stages, while the main amplifier for the comparator contained 6 including the last stage with the single-ended output.

3.2.2.3 Scan Chain

The scan chain in Figure 3.13 is a common means for circuit designers to place their chips into known configurations and states for testing.

For hierarchy, we refer to a scan cell as the two flip flops associated with the same scan bit, along with all associated buffers. For SEU immunity, we employed several tactics. At the lowest device level, we used DICE latches (Section 1.2.2, Figure 1.15) to form the flip flops, laid out such that mutually redundant nodes are spaced at least $15\mu\text{m}$ apart. We used triple modular redundancy at the scan cell level with 3-to-3 voters between each cell (Figure 3.14), but not for individual logic gates within the cells due to area constraints.

The area overhead of the custom DICE latches and triple redundancy was significant, with the scan chain and its subsequent drivers and level shifters occupying nearly half of the designed chip area.




Figure 3.12: The core fully differential amplifier in grey, with its offset control in black and boxed on the left. The choice to burn additional current by adding R_{OS} was to more consistently define the gain of the offset control across corners; biasing resistors largely guaranteed the current source for I_{OS} behaved as such.

3.2.2.4 One-Shot Pulse Generators

Separate one shot pulse generators are used for afterpulse rejection (Section 2.3) and for clock-free timing in the watchdog (Section 2.5). Figure 3.15 shows the topology used for both, where a positive input pulse produces a positive output pulse. A NAND-based topology was chosen over NOR to have the discharge path of the high pass filter lead to ground. The NAND topology was also smaller than the flip flop-based topology, with every option ultimately requiring an RC of some kind for a delay Figure 3.16 shows an example of normal operation in the presence of a (short) input after a long period of no input.

In general, high pass filters like one shots are undesirable in the presence of SEEs (Section 1.2.1.2). To mitigate this problem, we used triple modular redundancy with 3-to-3 voters on the outputs of each triplicated section. Figure 1.14b outlines the sections which were triplicated together in one cell. Here, the inverter chains as a whole rather than their composing inverters were triplicated for expediency and area conservation, since the additional voters would have significantly increased area consumption and routing complexity.

We added a reset switch for a more consistent pulse width across different event rates for the one shot. The output of the RC high pass filter intentionally has nontrivial settling time,

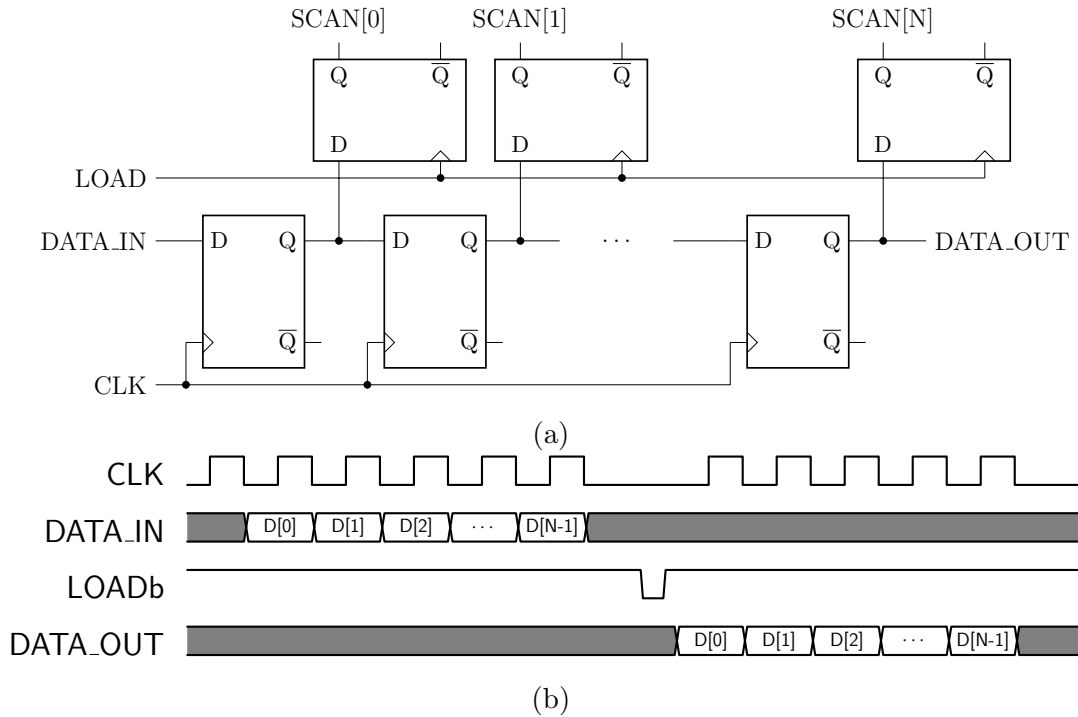


Figure 3.13: A standard scan chain circuit (a) and its operation (b). Buffers and the like have been removed for clarity.

since the RC time constant and the transition voltage of the subsequent inverter define the width of the output pulse. This however has the unintended consequence of making the RC node slow to settle, even after the output pulse has ended (visible in Figure 3.16). Figure 3.17 shows how this can lead to hysteresis in output pulse widths, where the duration of the output pulse depends on the time between input events. Variation in output pulse length directly affects afterpulse rejection/dead time (Section 2.3) and the timing trigger of the SEE watchdog (Section 2.5), so to ensure more consistent pulse widths across a wider range of event rates, we inserted a reset switch on the RC node driven by the logic in Equation 3.8.

$$\phi_{rst} = \overline{\text{in} + \text{out} + \text{out}_{\text{NAND}}} \quad (3.8)$$

If the input, output, and the output of the NAND (i.e. the input to the high pass filter) are all low, the reset switch is activated and the RC node is reset to 0V. Figure 3.18 shows the spread of output pulse widths of a one shot with and without the reset switch with $R = 20\text{k}\Omega$ and $C = 1\text{pF}$.

Without the reset, 29.4% of the output pulses varied by more than 20% from the median value, while only 4.4% of the output pulses varied that much with the reset switch. The difference in median pulse width with and without the reset switch is attributable to the




Figure 3.14: A triple modular redundant scan cell as it was used in the chip. As a defense against timing violations, the clock was routed in reverse order relative to the input data signal.




Figure 3.15: The one shot pulse generator, with triplicated components outlined. Each outlined section was followed by a 3-to-3 majority vote on its output(s). The reset switch is necessary to ensure consistent output pulse widths when input events are closely spaced in time relative to the RC time constant.

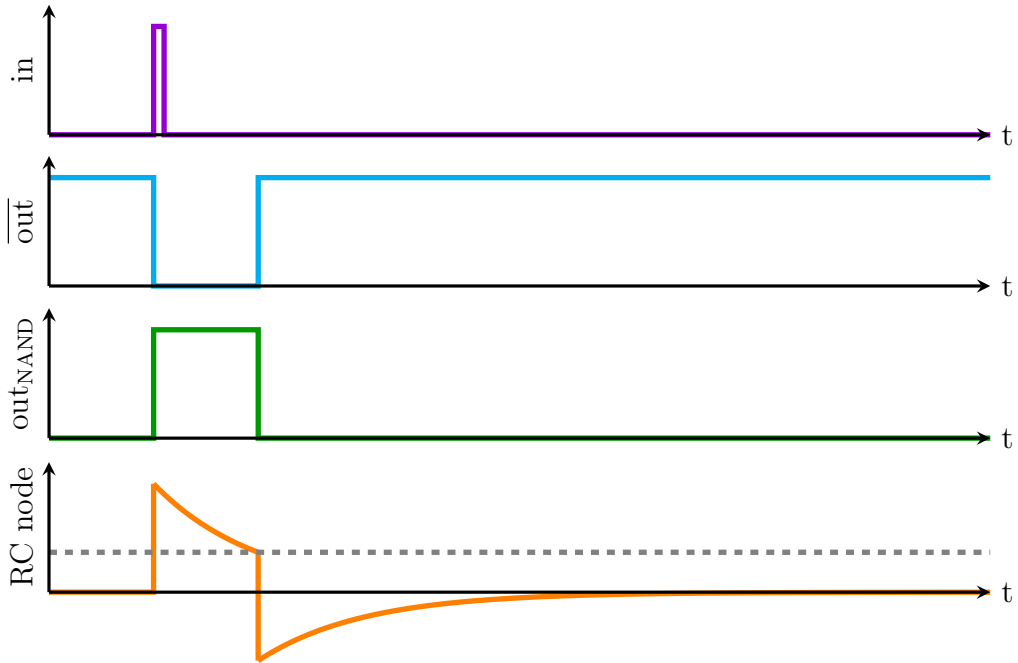


Figure 3.16: Normal one shot operation without the reset ϕ_{rst} in the presence of a single short pulse from the fully settled stable state. The grey line is the inverter switching point. Note that the the RC node takes finite time to settle to $\approx 0V$ even after the output pulse has terminated.




Figure 3.17: Inconsistent output pulse widths due to the prolonged settling time of the RC node.

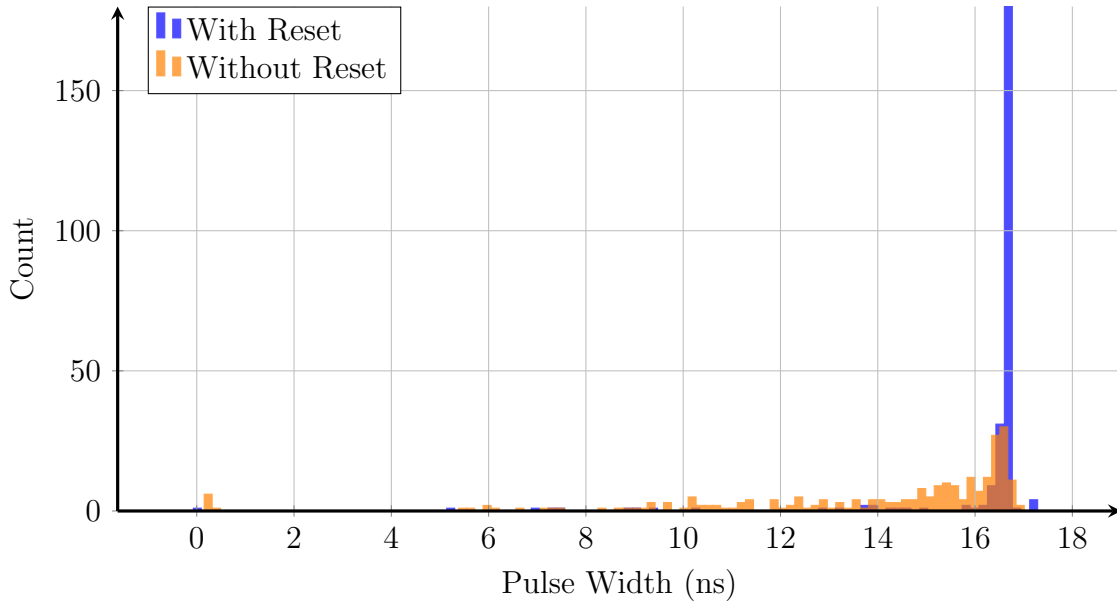


Figure 3.18: Simulated output pulse widths of a one shot pulse generator with (blue) and without (orange) the reset switch for $R = 20\text{k}\Omega$ and $C = 1\text{pF}$. Input pulses were 1ns wide with pulse spacing 1-100ns apart, randomly sampled. The median pulse width without the reset switch was 15.4ns and 16.8ns with the reset.

additional capacitance introduced on the RC node by the switch; the switch was necessarily large to have an on resistance $\approx 100\times$ smaller than R .

3.2.3 Full/Main Signal Chain

The full/main signal chain is a superset of the small/no-shape signal chain of Section 3.2.2. This subsection will only describe blocks which do not appear in the small signal chain. Figure 3.19 shows the setup used to stimulate and read out timing information from the main signal chain.

A similar calibration scheme for jitter and timing walk as for the small chain (Equations 3.1 and 3.2), with the appropriate components removed. Figure 3.20 the jitter and TDOA of the pulses after the contributions from non-chip components were accounted for.

3.2.3.1 Preamplifier

The preamplifier for this chip was a charge amplifier with a transimpedance amplifier (TIA) with an RC feedback network (Figure 3.21) and a gain approximately equal to the impedance of the feedback network (Equation 3.9). It was designed with the help of Mia Mirkovic.




Figure 3.19: Measurement setup for gathering timing statistics for the small signal chain. (a) Configure the DG535, TDC, and chip scan chain. (b) Trigger the DG535. The DG535's triggered output is then used as the START event to the TDC. A subsequent pulse nominally 2ns wide with an amplitude ranging from 0.1V to 1V is connected with a 50Ω termination to the PCB and AC coupled with a 2pF capacitor for current pulses of 1.2mA to 12mA into the preamplifier. Each pulse amplitude test is performed 500 times with at least 100ns between pulses. (c) The chip output is latched and level shifted from the 1.8V core voltage to 3.3V to be fed into the TDC as a STOP event.

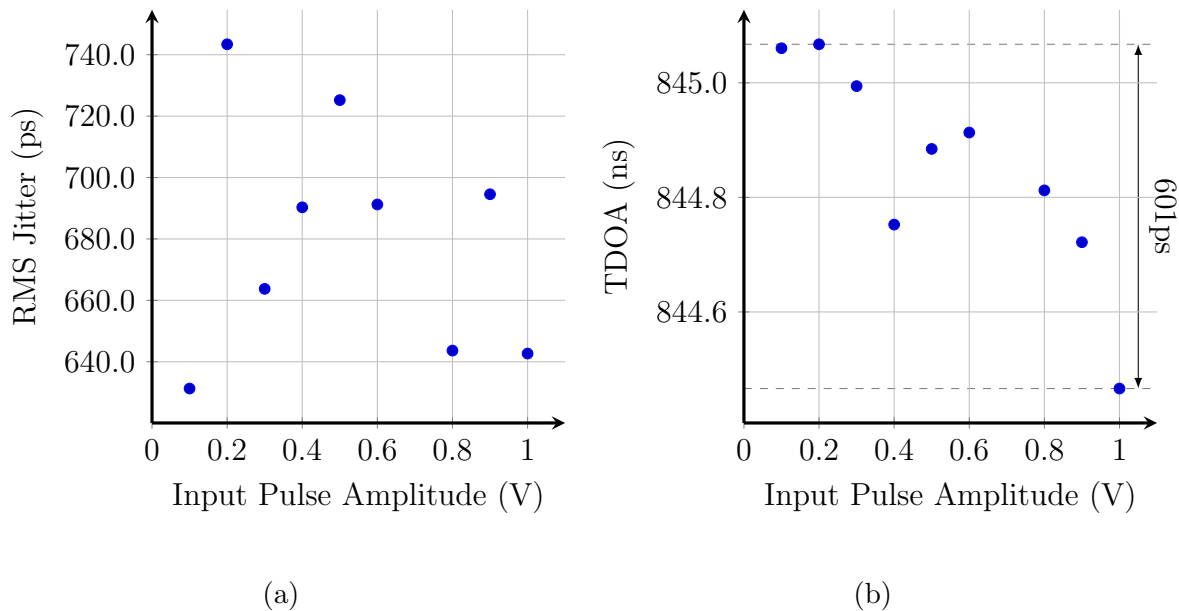


Figure 3.20: (a) Jitter and (b) time difference of arrival of the measured pulses for the main signal chain. Worst case jitter was measured at 743ps_{rms}, and timing walk at 601ps.

$$Z_F = \frac{R_F}{1 + sR_FC_F} \quad (3.9)$$




Figure 3.21: The preamplifier. The referencing used for biasing is generated by a resistor ladder DAC identical to the one used in the LED (Section 3.2.2.1).

We chose $C_F = 5\text{pF}$ and $R_F \in \{1, 2, 3, 4\}\text{k}\Omega$ so the time constant R_FC_F would be at least a factor of $5\times$ smaller than the minimum time between events and the preamplifier's step response would remain sufficiently linear. Because the input is broadband, we use compression of the output peak as a proxy measure of linearity. Figure 3.22 shows the peak of the preamplifier's output as the amplitude of the input pulse increases.

3.2.3.2 Peak Detector

The peak detector in Figure 3.23 is a standard diode-in-feedback topology with an added output buffer for output drive strength and multiple feedback for speed. The peak detector appears in the main signal chain but was also copied and padded out as a standalone test structure for sanity checking measurements. Additional ESD protection was placed at the storage node on the positive plate of C_{mid} .

The additional switch in grey at the input of the peak detector resets the preamplifier output was only included in V2 of the chip (Chapter 4); it was not necessary for V1 of the chip since the time between pulses was sufficiently long relative to the time constant of the preamplifier. Figure 3.24 shows how the preamplifier's RC low pass filter can prolong the falling edge of the input pulse. If the timing of the output one shot pulse generator is short relative to the tail of the preamplifier output, this can cause unintentional faux afterpulsing.

Static error measurements were taken by using the Teensy 3.6's analogWrite (16B) to slowly ramp the voltage to the desired value at a rate of approximately 3.3V/ms . The voltage is then read out via the Teensy 3.6's analogRead (13ENOB), the peak detector is reset, and the cycle continues. Figure 3.25 shows the output of the peak detector compared to the value measured from the Teensy. Each data point is the amalgamation of 200 samples. While at




Figure 3.22: (a) The simulated and ideal peak of the preamplifier's output as the amplitude of the input pulse increases. The ideal values are calculated by linearly fitting the simulated data. The charge corresponds to $2Me^-$ to $20Me^-$. (b) The compression of the peak as the size of the pulse increases.




Figure 3.23: The peak detector. The switch (grey) was not included in this version of the chip.

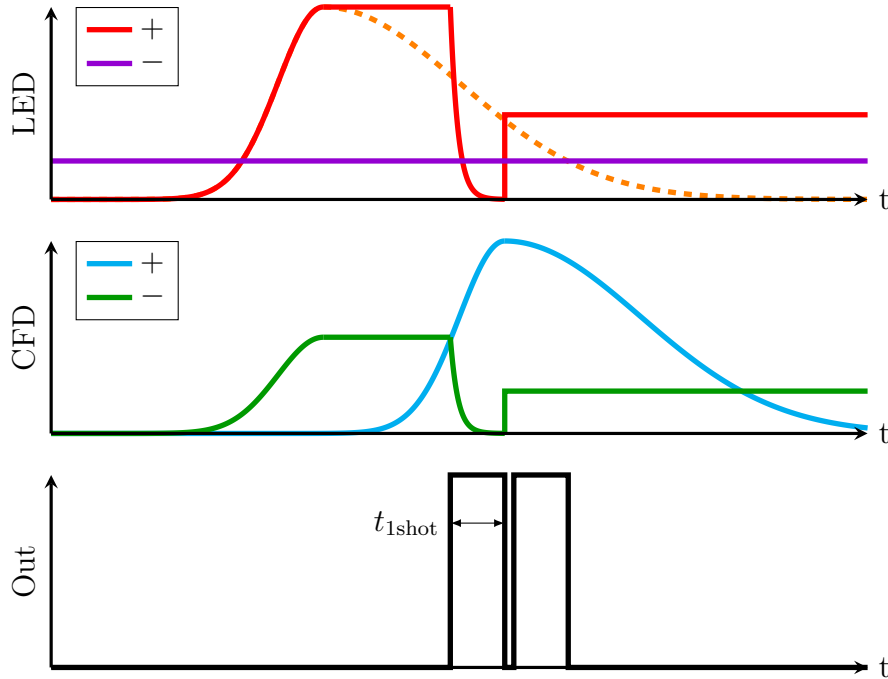



Figure 3.24: The positive and negative inputs to the LED and CFD comparators, without a reset on the output of the preamplifier. If $t_{1\text{shot}}$ is short relative to the settling time of the preamplifier, the peak detector output (red) will rise past the LED threshold and trigger the output once more, causing the appearance of afterpulsing.

face value the slope of the measured data appears smooth, the percent error in Figure 3.25b and the statistics of the measurement in Figure 3.25c have distinct spikes which occur with no discernible relation to the input voltage.


The spikes are due to transients on the output of the Teensy supplying the voltage. Examining the two largest spikes at 500mV and 569mV inputs at, we see in Figure 3.26 that the large variance is due to a single spike in voltage reading at the start of the repeated measurements. The Teensy as the source of these errors was confirmed by probing the output of the Teensy when it was disconnected from the chip under test.

3.2.3.3 Delay


Previous versions of SPAN-I used a discrete commercial component for its delay, with trial and error used between different delay values and individual components to achieve the desired walk and sensitivity. The main signal chain integrates the delay and its tuning as a Bessel filter. Bessel filters have maximally flat group delay in their passband, making them a convenient alternative to conventional transmission lines. For this chip, we implemented the second order Bessel filter in Equation 3.10 with the Sallen-Key topology shown in 3.27



(a)



(b)



(c)

Figure 3.25: (a) Static error of the peak detector. (b) The percent error of the measured output voltage relative to its target value. (c) The standard deviation of the measured output voltage. The error and statistics show noticeable spikes.

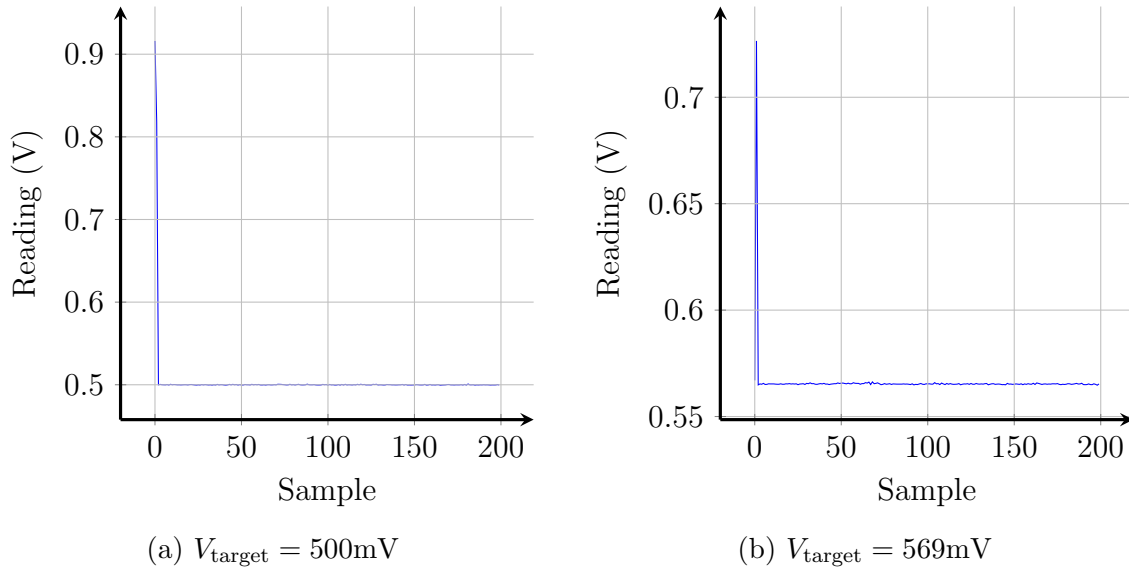


Figure 3.26: Tracking the spikes in voltage for (a) the largest and (b) second largest spikes in measurement variance.

and described by Equation 3.11. The topology was chosen for its simplicity and the use of only one amplifier relative to other, e.g. Tow-Thomas biquad topologies.

$$H_{\text{Bessel2}}(s) = \frac{3\omega_0^2}{s^2 + 3s\omega_0 + 3\omega_0^2} \quad (3.10)$$




Figure 3.27: Sallen-Key topology associated with the transfer function in Equation 3.11.

$$H_{\text{SK}}(s) = \frac{1}{s^2 + s \frac{1}{C_1} \frac{R_1 + R_2}{R_1 R_2} + \frac{1}{R_1 R_2 C_1 C_2}} \quad (3.11)$$

To match the filter component values to the desired transfer function, we constrained the solution space by beginning with “reasonable” starting values for C_1 and C_2 , leaving R_1 and R_2 to satisfy the conditions in Equation 3.12.

$$R_1 R_2 = \frac{1}{4\omega_0^2 C_1 C_2} \quad (3.12a)$$

$$R_1 + R_2 = \frac{1}{C_2 \omega_0} \quad (3.12b)$$

Across process and temperature corners, the simulated delay of each signal ranged from 3.66ns to 10.53ns across the full tuning range, with the tuning range at nominal operating conditions spanning 3.98ns to 8.29ns. The minimum value was chosen to account for the delay (1.7ns to 3.0ns across corners) introduced by the diode turn-on time in the peak detector.

3.2.3.4 Attenuator

The attenuator is an 8-element resistive divider with a mux tapping out all nodes except the reference at the bottom of the ladder for an attenuation factor $f \in \{\frac{1}{8}, \frac{2}{8}, \dots, \frac{8}{8}\}$ (Figure 3.28). The V_{REF} at the bottom of the ladder is a buffered copy of the output of the preamplifier biasing reference for DC cancellation, padded out to enable external override. Each individual resistor was chosen to be $2\text{k}\Omega$ for a collective resistance of $16\text{k}\Omega$. Resistive




Figure 3.28: The resistive divider attenuator.

loading of the peak detector was of particular concern, more so than low pass filtering of the

incoming signal; because the peak detector holds its final value, low pass filtering will settle on the correct final value. Thus, an adequately long delay will ensure a trigger at a constant fraction, and the linearity of the (admittedly nonideal) low pass filter will still enable zero theoretical timing walk between identically shaped pulses of different amplitude.

Chapter 4

Chip V2

This chapter describes the internals and measurements taken from the second chip taped out for this project. It is not intended as a user manual. For detailed documentation of file locations, I/O, scan bits, etc., see Appendix F.

4.1 Chip Summary

The chip (Figure 4.1) was taped out on November 16, 2022 in TSMC’s 180nm process through the Berkeley Sensor and Actuator Center. The total chip area including pads and seal ring was 2.5mm×2.5mm. The fundamental operation of Chip V2 is similar to that of Chip V1 in Chapter 3: Like V1, V2 contains on-chip power management and reference generation, derived from a single 3.3V supply; the full/main front end is described by Figure 2.9; the shortened/small/no-shape front end has the same high level architecture as Figure 3.5; there is a substantial scan chain which consumes a large percentage of the total silicon area; and there are standalone test structures on a power domain separated from the rest of the chip. However, low-level modifications were made at the individual block level to address verification misses from Chip V1. These will be discussed in greater detail in Section 4.2. Further debugging signals not originally made visible in Chip V1—particularly the outputs of the various comparators prior to the arming logic—were also padded out, and on-chip level shifters were added for the primary output to interface with off-chip components operating at 3.3V rather than the core 1.8V.

Table 4.1 provides a snapshot of the chip’s simulated performance in relation to the target application. Prior to this dissertation’s submission deadline, the amount of data taken from the chip was limited to extremely basic test structures, e.g. the bandgap reference voltage, which were identical to those in Chip V1. This is because of issues encountered during PCB assembly with packaged chips. More specifically, the materials used to package the chips could not withstand the temperatures used for surface mount soldering, resulting in electrical disconnects between the chip and the package.




Figure 4.1: Chip die photo with main structures annotated. Photo courtesy of Alexander Alvara.

4.2 Design and Measurements

This section will primarily describe the design considerations of the chip as they differ from those taken in Chapter 3, addressing both shortcomings in the design of the previous chip and discussing pre-silicon simulation results.

4.2.1 No-Shape/Small Signal Chain

The architecture of the no-shape/small signal chain is the same as that of Section 3.2.2. However, there were changes made at the block level, so this section will focus on those changes and their effect on the overall signal chain. Figure 4.2 shows the TDOA of the chip with an emulated benchtop setup with a realistic microchannel plate response. The timing

Parameter	Target	Chip V1	Chip V2
Minimum Detectable Signal	2Me^-	$8\text{-}10\text{Me}^-$	$2\text{-}3\text{Me}^-$
Maximum Event Rate	10Mevent/s	Simulated $\geq 10\text{Mevent/s}$	
Signal Chain Integration	ASIC		
Timing Walk w/o Shaping	600-800ps total	682ps	$< 300\text{ps}$
Jitter w/o Shaping		177ps	$< 100\text{ps}$
Timing Walk w/ Shaping	600-800ps total	601ps	$< 400\text{ps}$
Jitter w/ Shaping		743ps	$< 200\text{ps}$
SEU Tolerance	Immune	Simulated	
TID Tolerance	100krads	Simulated	
Power	$3\text{-}4\text{mA}$	2.9mA	3.3mA
ASIC Area	$\leq 2.5\text{mm} \times 2.5\text{mm}$	$1.6\text{mm} \times 1.7\text{mm}$	$2.5\text{mm} \times 2.5\text{mm}$
Process	TSMC 180nm CMOS		

Table 4.1: Chip V2 simulated performance versus the target specifications.

walk across an order of magnitude change in charge input was 206ps. The increase in time of arrival as the size of the pulse increases is due to a decrease in bandwidth of the comparators with a rising common mode input.

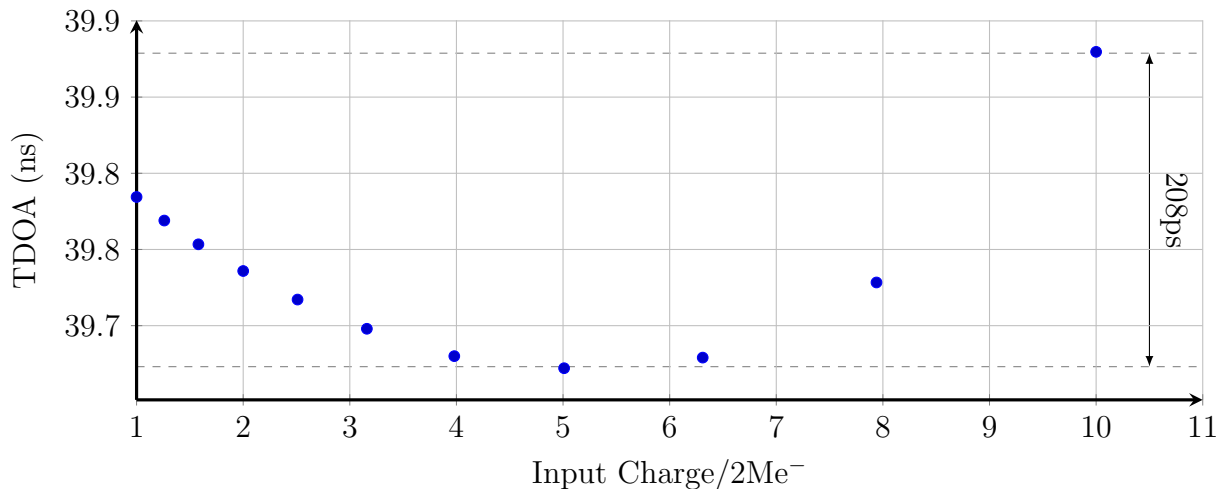


Figure 4.2: Time difference of arrival of the simulated pulses with an emulated benchtop setup and microchannel plate.

The concept of the comparators' design—many low-gain, high-bandwidth stages—remained the same between the first and second versions of the chip. However, after tapeout we observed a strong gain and bandwidth dependence on the comparators' input common mode,

with bandwidths varying up to two orders of magnitude across a 200mV change in input common mode. In the absence of some common mode tuning or feedback, process drift between stages can lead to large shifts in input common mode, eviscerating the effective gain and bandwidth of the comparator sub-stage. While some degree of indirect control is possible by setting the preamplifier DC bias, it proved unreliable when searching for a “sweet spot” for desired performance across a large quantity of devices in varying operating conditions, with the shifting common mode sometimes severe enough to stick the output of the final single ended stage high or low. A redesign of the comparator stage via design script with an expanded input common mode range yielded stages with a fairly consistent gain of $\approx 2V/V$ across temperature and input common mode (Figure 4.3).




Figure 4.3

To allow this to be calibrated out with more accuracy and precision, we incorporated tuning into each comparator stage capable of both pull-up and pull-down on each output of the stage, shown in Figure 4.4. This enables both fixed output offset correction, as well as output common mode tuning. Each stage is given 4 bits of pull-up tuning, 4 bits of pull-down tuning, and independent control of pull-up and pull-down for both the positive and negative outputs for a total of 12 tuning bits per stage. Tuning was designed to allow tuning of up to $\pm 250\text{mV}$ on both outputs in $\approx 7.8\text{mV}$ increments.

4.2.2 Full/Main Signal Chain

The architecture of the full/main signal chain is the same as that of Section 3.2.3, with the same changes made to the comparators in Section 4.2.1. Within the analog pulse shaping, the preamplifier and delay filter underwent the most significant changes, and the peak detector was modified to include the input reset switch in Figure 3.23. The worst case jitter simulated across corners and temperature was $192\text{ps}_{\text{rms}}$, with timing walk of 340ps across an order of magnitude change in pulse amplitude.

For the preamplifier, the insertion of the reset switch on the output of the preamplifier/input of the peak detector made it so the time constant $R_F C_F$ no longer needed to be




Figure 4.4: Trimmed comparator stage.




Figure 4.5: Timing walk versus input common mode over pulses ranging from 2Me^- to 20Me^- .

significantly shorter than the minimum time between input pulses. With that in mind, we decreased the size of the feedback capacitance from 5pF to 2pF and upscaled the reset resistance by a factor of $5\times$ so the preamp transient behavior appears more as a charge amplifier with a resistor- and switch-based reset. The decision to keep the resistor rather than replace it entirely with a reset switch retained a consistent DC feedback path.

For this chip, the delay was implemented as a fourth order Bessel filter, described by Equation 4.1.

$$H_{\text{Bessel4}}(s) = \frac{105\omega_0^4}{s^4 + 10\omega_0 s^3 + 45\omega_0^2 s^2 + 105\omega_0^3 s + 105\omega_0^4} \quad (4.1)$$

The increased order was to expand the bandwidth of the filter for the same group delay; for a group delay of 3.5ns, the frequency at which the group delay degrades by 50% increases by more than a factor of $2\times$ from the second to fourth order filter (Figure 4.6).

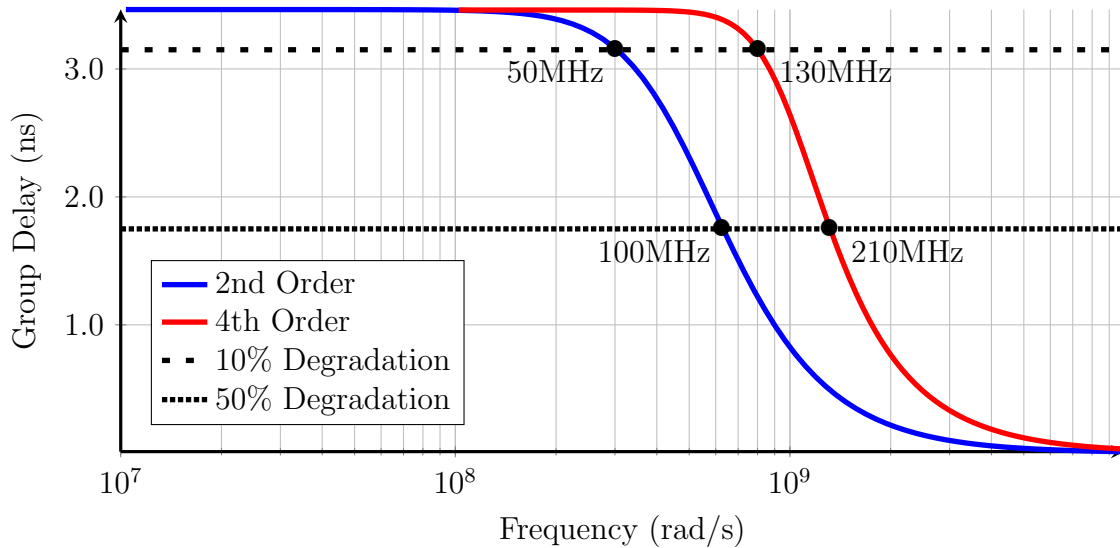


Figure 4.6: Group delays vs. frequency of a second order and fourth order Bessel filter, with the point of 10% and 50% change marked with a black horizontal line.

As with the second order filter, we used the Sallen-Key topology twice (Figure 4.7) for the transfer function in Equation 4.2

$$H_{\text{SK},2\times}(s) = \left(\frac{\frac{1}{R_{1A}R_{2A}C_{1A}C_{2A}}}{s^2 + s\frac{1}{C_{1A}}\frac{R_{1A}+R_{2A}}{R_{1A}R_{2A}} + \frac{1}{R_{1A}R_{2A}C_{1A}C_{2A}}} \right) \left(\frac{\frac{1}{R_{1B}R_{2B}C_{1B}C_{2B}}}{s^2 + s\frac{1}{C_{1B}}\frac{R_{1B}+R_{2B}}{R_{1B}R_{2B}} + \frac{1}{R_{1B}R_{2B}C_{1B}C_{2B}}} \right) \quad (4.2)$$

Looking to to Equations 4.1 and 4.2, we some algebraic manipulation we are left with the conditions in Equation 4.3.




Figure 4.7: Two Sallen-Key filters to give the fourth order transfer function in Equation 4.2.

$$\frac{1}{\omega_0} = C_{2B}(R_{1B} + R_{2B}) + C_{2A}(R_{1A} + R_{2A}) \quad (4.3a)$$

$$\frac{3}{7\omega_0^2} = R_{1B}R_{2B}C_{1B}C_{2B} + R_{1A}R_{2A}C_{1A}C_{2A} + C_{2A}(R_{1A} + R_{2A})C_{2B}(R_{1B} + R_{2B}) \quad (4.3b)$$

$$\frac{2}{21\omega_0^3} = C_{2A}C_{1B}C_{2B}R_{1B}R_{2B}(R_{1A} + R_{2A}) + C_{1A}C_{2A}C_{2B}R_{1A}R_{2A}(R_{1B} + R_{2B}) \quad (4.3c)$$

$$\frac{1}{105\omega_0^4} = R_{1A}R_{2A}R_{1B}R_{2B}C_{1A}C_{2A}C_{1B}C_{2B} \quad (4.3d)$$

Unlike the original method of coefficient matching for the second order Bessel filter, choosing capacitor values is insufficient to yield a(n obvious) unique solution for the resistor values. Rather than futz around endlessly with algebra to find an analytical expression for the hypersurface of solutions, we set a constant value for R and begin with numerical values produced by Nuhertz FilterSolutions (now owned by Ansys) for a fourth order Bessel filter. From there, scaling all the capacitors (or resistors) together by a constant factor N preserves the important characteristics of the transfer function—i.e. it will still be a Bessel filter—while scaling ω_0 by $\frac{1}{N}$. The choice to make all the resistors identical had the added benefit

Code	Delay (ns)	Attenuation Factor
00	10.1	0.78
01	13.5	0.71
10	15.7	0.74
11	18.4	0.67

Table 4.2: Delay filter passband group delay and attenuation.

of greater layout reuse, since the tuning of the transfer function is done with resistive DACs. Table 4.2 shows the passband group delay and attenuation versus code setting, measured

from AC and transient simulations. In spite of the nontrivial attenuation due to gain error and some amount of bandwidth limiting, the filter response is still ultimately linear, and the front end will still trigger at a constant (albeit slightly different) fraction.

Chapter 5

Conclusions and Future Work

With ever-compressing size, weight, and power constraints across a range of operating conditions, the design of sensor front ends has no one-size-fits-all solution. This thesis addressed high precision pulse discrimination for timing-based systems in a space-based operating environment, applied to the first attempt at integrating the pulse shaping front end for the time-of-flight mass spectrometer in SPAN-Ion. Its contribution is not only hardware to meet the demands of that specific application, but a code base of process-independent generators and design scripts for the Berkeley Analog Generator which will hopefully see use in future sensor front ends.

Chapter 1 examined the system requirements of a pulse discriminator for SPAN-Ion's time-of-flight mass spectrometer. In it, we addressed SPAN-Ion's operating environment and provided a summary of various techniques used to detect, correct, and mitigate radiation effects in integrated circuits. We also introduced the Berkeley Analog Generator as it was used in this work.

Chapter 2 included an architectural analysis for a front end suited for high precision pulse discrimination, within the limitations of the fabrication process on hand. The proposed constant fraction discriminator triggers relative to a well-defined shape-independent fraction of the input pulse's peak and produces theoretically zero timing walk. Considerations for afterpulse rejection, as well as a scheme for SEE detection/correction/rejection in both analog pulse shaping hardware and digital logic are presented as matters of practicality.

Chapter 3 presented the first attempt at fully integrated analog pulse shaping for SPAN-Ion and its legacy of time-of-flight mass spectrometers. The fully integrated front end fell slightly short of the sensitivity and jitter requirements to distinguish between higher mass/charge ions, but otherwise is capable of differentiating particles with $m/q \leq 20 \frac{\text{amu}}{e^-}$ with total front end current consumption in the low singles of mA. In addition to the fully integrated front end, the chip included a pared-down variant of the front end which lacked integrated shaping, an SPI interface, and integrated power management designed to operate from a single supply provided by the high voltage board of SPAN-Ion. The chapter also presented a substantial discussion of the design considerations for each component of the front end, power management, and SPI, with data measured from various test points throughout

the device.

Chapter 4 extended on the work done for the previous chip, addressing sources of remnant timing walk and incorporating low-level block adjustments for improved minimum detectable signal. Simulations indicate a timing walk and jitter capable of distinguishing between ions of m/q up to at least 40 with the time-of-flight setup in SPAN-Ion with a minimum detectable pulse of 2Me^- .

Even if Chip V2 works perfectly as intended, there is still room for expansion both for SPAN-Ion and its future descendants, as well as the more nebulous goal of building an open source generator and design script foundation for circuit design.

5.1 Button-Press Design with BAG

The generators and scripts written for this thesis were limited to the scope of this specific project and the infrastructure available for the chosen fabrication process. While a complete collection of generators and design scripts for every possible circuit to go into a sensor front end is unrealistic for an application-targeted thesis, there is room for better utilization of the BAG design framework. In particular, this work used schematic generation and design scripts applicable only at the schematic level. This is because layout generation for this process was not set up, and there was insufficient time to do so—for context, setting up XBase layout generation in one process took the majority of a research group the better part of a year. That said, the generators and scripts written for this work are generic enough for use in other timing-based applications, though expanding the code base to include layout generation would contribute significantly to fully automating the design of other similar front ends.

A typical top-down design flow might look something like Figure 5.1, with significant iteration occurring between and within each stage of the design.




Figure 5.1: A typical design flow.

In its full form, BAG is capable of encompassing the entirety of this design flow within the framework for true “button-press design.” This is achieved by creating, using, and

testing process-portable generators, all within the BAG framework. Much like the schematic generation used in this dissertation, BAG is capable of generating test benches and layout in a similar fashion, inserting the device under test into a test bench, and invoking simulators to execute the tests associated with a test bench. With all of this, it is possible to execute the entirety of the design flow within BAG, which in turn makes closed loop, automated iteration a possibility.

A closed design loop within BAG consists of several key components: The first is the generators, both for layout and schematic. Both layout and schematic generators take in physical parameters such as device length and width and produce a corresponding cell with those physical specifications. In the case of layout generation, LVS, DRC, and PEX may be run after creation, allowing for modeling as true to physical hardware as possible during the design loop, without the significant time overhead typically associated with layout.

Another necessary element to close the design loop is the design scripts and design managers. The infrastructure described in [71, 78] is specific to BAG3.0 and beyond; this work used BAG2.0. While the design manager does not exist in the same sense in BAG2 as it does in BAG3, it is still possible to create design scripts which operate much in the same way between the two versions. At a high level, the design scripts are the programmatic implementation of a designer's design procedure. These can vary in scope and complexity, and can similarly be done in a hierarchical fashion. For example, designing power management like the LDO in Figure 3.4 at a minimum involves sizing the series device as well as the amplifier to meet a target specification. The design script for the LDO can call upon a separate script to design the amplifier after sizing the series device, preserving good hierarchy practice.

A large proportion of the system modeling in Chapter 2 was done using MATLAB. Porting that modeling to the BAG framework would allow for more agile iteration in the design process, particularly when defining the target specifications of individual blocks. Limited as this work was in its use of BAG, the scripts and generators still significantly reduced the time required to port from XT018 to TSMC180. We hope the generators and design scripts written for this work will similarly save others time and energy in the pursuit of engineering solutions to interesting problems.

Bibliography

- [1] D. Abbaneo et al. “Design of a constant fraction discriminator for the VFAT3 front-end ASIC of the CMS GEM detector”. In: *Journal of Instrumentation* 11.01 (Jan. 2016), p. C01023. DOI: 10.1088/1748-0221/11/01/C01023. URL: <https://dx.doi.org/10.1088/1748-0221/11/01/C01023>.
- [2] European Space Agency. “TNID Total Non Ionizing Dose or DD Displacement Damage”. In: 2017. URL: https://indico.cern.ch/event/635099/contributions/2570676/attachments/1456363/2249942/4_Radiation_Effects_and_RHA_ESA_Internal_Course_4_10_May_2017_DD_CP.pdf.
- [3] J. R. Ahlbin et al. “The Effect of Layout Topology on Single-Event Transient Pulse Quenching in a 65 nm Bulk CMOS Process”. In: *IEEE Transactions on Nuclear Science* 57.6 (2010), pp. 3380–3385. DOI: 10.1109/TNS.2010.2085449.
- [4] M. Allenspach et al. “Evaluation of SEGR threshold in power MOSFETs”. In: *IEEE Transactions on Nuclear Science* 41.6 (1994), pp. 2160–2166. DOI: 10.1109/23.340557.
- [5] G. Anelli et al. “Radiation tolerant VLSI circuits in standard deep submicron CMOS technologies for the LHC experiments: practical design aspects”. In: *IEEE Transactions on Nuclear Science* 46.6 (1999), pp. 1690–1696. DOI: 10.1109/23.819140.
- [6] H. J. Barnaby. “Total-Ionizing-Dose Effects in Modern CMOS Technologies”. In: *IEEE Transactions on Nuclear Science* 53.6 (2006), pp. 3103–3121. DOI: 10.1109/TNS.2006.885952.
- [7] R. Baumann. “Radiation-induced soft errors in advanced semiconductor technologies”. In: *Device and Materials Reliability, IEEE Transactions on* 5 (Oct. 2005), pp. 305–316. DOI: 10.1109/TDMR.2005.853449.
- [8] H.N. Becker, T.F. Miyahira, and A.H. Johnston. “Latent damage in CMOS devices from single-event latchup”. In: *IEEE Transactions on Nuclear Science* 49.6 (2002), pp. 3009–3015. DOI: 10.1109/TNS.2002.805332.
- [9] J.M. Benedetto et al. “Variation of digital SET pulse widths and the implications for single event hardening of advanced CMOS processes”. In: *IEEE Transactions on Nuclear Science* 52.6 (2005), pp. 2114–2119. DOI: 10.1109/TNS.2005.860679.

- [10] Dolores A. Black et al. “Modeling of Single Event Transients With Dual Double-Exponential Current Sources: Implications for Logic Cell Characterization”. In: *IEEE Transactions on Nuclear Science* 62.4 (2015), pp. 1540–1549. DOI: 10.1109/TNS.2015.2449073.
- [11] David Burnett. “Crystal-free wireless communication with relaxation oscillators and its applications”. PhD thesis. EECS Department, University of California, Berkeley, Apr. 2019. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-5.html>.
- [12] David C. Burnett et al. “Two-Chip Wireless H₂S Gas Sensor System Requiring Zero Additional Electronic Components”. In: *2019 20th International Conference on Solid-State Sensors, Actuators and Microsystems & Eurosensors XXXIII (TRANSDUCERS & EUROSENSORS XXXIII)* (2019), pp. 1222–1225.
- [13] T. Calin, M. Nicolaidis, and R. Velazco. “Upset hardened memory design for submicron CMOS technology”. In: *IEEE Transactions on Nuclear Science* 43.6 (1996), pp. 2874–2878. DOI: 10.1109/23.556880.
- [14] Eric Chang et al. “BAG2: A process-portable framework for generator-based AMS circuit design”. In: *2018 IEEE Custom Integrated Circuits Conference (CICC)*. 2018, pp. 1–8. DOI: 10.1109/CICC.2018.8357061.
- [15] Athanasios Chatzidimitriou et al. “Multi-Bit Upsets Vulnerability Analysis of Modern Microprocessors”. In: *2019 IEEE International Symposium on Workload Characterization (IISWC)*. 2019, pp. 119–130. DOI: 10.1109/IISWC47752.2019.9042036.
- [16] H. Choi et al. “Study of timing performance parameters for a SiPM-based digital positron annihilation lifetime spectrometer”. In: *Journal of Instrumentation* 17.12 (Dec. 2022), p. C12007. DOI: 10.1088/1748-0221/17/12/C12007. URL: <https://dx.doi.org/10.1088/1748-0221/17/12/C12007>.
- [17] C. Dufour et al. “Heavy ion induced single hard errors on submicronic memories (for space application)”. In: *IEEE Transactions on Nuclear Science* 39.6 (1992), pp. 1693–1697. DOI: 10.1109/23.211355.
- [18] G. L. Engel et al. “Multi-Channel Integrated Circuits For Use In Research With Radioactive Ion Beams”. In: *AIP Conference Proceedings* 1336.1 (2011), pp. 608–613. DOI: 10.1063/1.3586175. eprint: <https://aip.scitation.org/doi/pdf/10.1063/1.3586175>. URL: <https://aip.scitation.org/doi/abs/10.1063/1.3586175>.
- [19] E.W. Enlow et al. “Response of advanced bipolar processes to ionizing radiation”. In: *IEEE Transactions on Nuclear Science* 38.6 (1991), pp. 1342–1351. DOI: 10.1109/23.124115.

- [20] Lars Evensen et al. “Guard ring design for high voltage operation of silicon detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 337.1 (1993), pp. 44–52. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/0168-9002\(93\)91136-B](https://doi.org/10.1016/0168-9002(93)91136-B). URL: <https://www.sciencedirect.com/science/article/pii/016890029391136B>.
- [21] F. Faccio and G. Cervelli. “Radiation-induced edge effects in deep submicron CMOS transistors”. In: *IEEE Transactions on Nuclear Science* 52.6 (2005), pp. 2413–2420. DOI: 10.1109/TNS.2005.860698.
- [22] M. Fang, N. Bartholomew, and A. Di Fulvio. “Positron annihilation lifetime spectroscopy using fast scintillators and digital electronics”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 943 (2019), p. 162507. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2019.162507>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900219310460>.
- [23] M.W. Fishburn. “Fundamentals of CMOS Single-Photon Avalanche Diodes”. PhD thesis. Delft University of Technology, 2012.
- [24] D.M. Fleetwood. “Fast and slow border traps in MOS devices”. In: *IEEE Transactions on Nuclear Science* 43.3 (1996), pp. 779–786. DOI: 10.1109/23.510713.
- [25] D.M. Fleetwood, P.S. Winokur, and J.R. Schwank. “Using laboratory X-ray and cobalt-60 irradiations to predict CMOS device response in strategic and space environments”. In: *IEEE Transactions on Nuclear Science* 35.6 (1988), pp. 1497–1505. DOI: 10.1109/23.25487.
- [26] Geng Fu et al. “Improved walk-correction method for timing measurements in PET detector”. In: *2014 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. 2014, pp. 1–3. DOI: 10.1109/NSSMIC.2014.7431004.
- [27] Matthew Gadlage et al. “Temperature Dependence of Digital Single-Event Transients in Bulk and Fully-Depleted SOI Technologies”. In: *Nuclear Science, IEEE Transactions on* 56 (Oct. 2010), pp. 3115–3121. DOI: 10.1109/TNS.2009.2034150.
- [28] M. Ghioni et al. “True constant fraction trigger circuit for picosecond photon-timing with ultrafast microchannel plate photomultipliers”. In: *Review of Scientific Instruments* 68.5 (May 1997), pp. 2228–2237. ISSN: 0034-6748. DOI: 10.1063/1.1148074. eprint: https://pubs.aip.org/aip/rsi/article-pdf/68/5/2228/8811928/2228__1__online.pdf. URL: <https://doi.org/10.1063/1.1148074>.
- [29] G. Gildenblat et al. “PSP: An Advanced Surface-Potential-Based MOSFET Model for Circuit Simulation”. In: *IEEE Transactions on Electron Devices* 53.9 (2006), pp. 1979–1993. DOI: 10.1109/TED.2005.881006.

- [30] Balkaran S. Gill, Chris Papachristou, and Francis G. Wolff. “A New Asymmetric SRAM Cell to Reduce Soft Errors and Leakage Power in FPGA”. In: *2007 Design, Automation and Test in Europe Conference & Exhibition*. 2007, pp. 1–6. DOI: 10.1109/DATE.2007.364504.
- [31] BSIM Group. 2021. URL: <http://bsim.berkeley.edu/>.
- [32] Kyung Han et al. “RF characteristics of 0.18- μm CMOS transistors”. In: *Journal-Korean Physical Society* 40 (Jan. 2002), pp. 45–48.
- [33] J.H. Hohl and G.H. Johnson. “Features of the triggering mechanism for single event burnout of power MOSFETs”. In: *IEEE Transactions on Nuclear Science* 36.6 (1989), pp. 2260–2266. DOI: 10.1109/23.45433.
- [34] F.N. Hooge. “1/f noise is no surface effect”. In: *Physics Letters A* 29.3 (1969), pp. 139–140. ISSN: 0375-9601. DOI: [https://doi.org/10.1016/0375-9601\(69\)90076-0](https://doi.org/10.1016/0375-9601(69)90076-0). URL: <https://www.sciencedirect.com/science/article/pii/0375960169900760>.
- [35] Zhiyuan Hu et al. “Comprehensive Study on the Total Dose Effects in a 180-nm CMOS Technology”. In: *IEEE Transactions on Nuclear Science* 58.3 (2011), pp. 1347–1354. DOI: 10.1109/TNS.2011.2132145.
- [36] Georg Kirchner et al. “Time-walk-compensated SPAD: multiple-photon versus single-photon operation”. In: *Laser Radar Ranging and Atmospheric Lidar Techniques*. Ed. by Ulrich Schreiber and Christian Werner. Vol. 3218. International Society for Optics and Photonics. SPIE, 1997, pp. 106–112. DOI: 10.1117/12.295659. URL: <https://doi.org/10.1117/12.295659>.
- [37] R. Koga et al. “Single event functional interrupt (SEFI) sensitivity in microcircuits”. In: *RADECS 97. Fourth European Conference on Radiation and its Effects on Components and Systems (Cat. No.97TH8294)*. 1997, pp. 311–318. DOI: 10.1109/RADECS.1997.698915.
- [38] S.L. Kosier et al. “Bounding the total-dose response of modern bipolar transistors”. In: *IEEE Transactions on Nuclear Science* 41.6 (1994), pp. 1864–1870. DOI: 10.1109/23.340518.
- [39] S.L. Kosier et al. “Physically based comparison of hot-carrier-induced and ionizing-radiation-induced degradation in BJTs”. In: *IEEE Transactions on Electron Devices* 42.3 (1995), pp. 436–444. DOI: 10.1109/16.368041.
- [40] S. Kulis. “Single Event Effects mitigation with TMRG tool”. In: *Journal of Instrumentation* 12.01 (2017), p. C01082. URL: <http://stacks.iop.org/1748-0221/12/i=01/a=C01082>.
- [41] Lawrence Berkeley National Laboratory. *Cocktails and Ions*. <https://cyclotron.lbl.gov/base-rad-effects/heavy-ions/cocktails-and-ions>. Aug. 2020.

- [42] Ronald C. Laco. “Improving Integrated Circuit Performance Through the Application of Hardness-by-Design Methodology”. In: *IEEE Transactions on Nuclear Science* 55.4 (2008), pp. 1903–1925. DOI: 10.1109/TNS.2008.2000480.
- [43] Joseph Ladislav Wiza. “Microchannel plate detectors”. In: *Nuclear Instruments and Methods* 162.1 (1979), pp. 587–601. ISSN: 0029-554X. DOI: [https://doi.org/10.1016/0029-554X\(79\)90734-1](https://doi.org/10.1016/0029-554X(79)90734-1). URL: <https://www.sciencedirect.com/science/article/pii/0029554X79907341>.
- [44] Roberto Livi et al. “The Solar Probe ANalyzer—Ions on the Parker Solar Probe”. In: *The Astrophysical Journal* 938.2 (Oct. 2022), p. 138. DOI: 10.3847/1538-4357/ac93f5. URL: <https://dx.doi.org/10.3847/1538-4357/ac93f5>.
- [45] Filip Maksimovic. “Monolithic Wireless Transceiver Design”. PhD thesis. EECS Department, University of California, Berkeley, May 2020. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2020/EECS-2020-33.html>.
- [46] Lloyd W. Massengill et al. “Technology scaling and soft error reliability”. In: *2012 IEEE International Reliability Physics Symposium (IRPS)*. 2012, pp. 3C.1.1–3C.1.7. DOI: 10.1109/IRPS.2012.6241810.
- [47] Timothy C. May and Murray H. Woods. “A New Physical Mechanism for Soft Errors in Dynamic Memories”. In: *16th International Reliability Physics Symposium*. 1978, pp. 33–40. DOI: 10.1109/IRPS.1978.362815.
- [48] J. P. McFadden et al. “MAVEN SupraThermal and Thermal Ion Composition (STATIC) Instrument”. In: *Space Science Reviews* 195.1 (Dec. 2015), pp. 199–256. ISSN: 1572-9672. DOI: 10.1007/s11214-015-0175-6. URL: <https://doi.org/10.1007/s11214-015-0175-6>.
- [49] Flynn B. McLean and Timothy R. Oldham. “Basic Mechanisms of Radiation Effects in Electronic Materials and Devices”. In: 1987.
- [50] Alan L. McWhorter. “1/f noise and related surface effects in germanium”. In: 1955.
- [51] T. L. Meisenheimer and Daniel M. Fleetwood. “Effect of radiation-induced charge on 1/f noise in MOS devices”. In: *IEEE Transactions on Nuclear Science* 37 (1990), pp. 1696–1702.
- [52] Sahar Mesri. “Design and User Guide for the Single Chip Mote Digital System”. MA thesis. EECS Department, University of California, Berkeley, May 2016. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-71.html>.
- [53] G. C. Messenger. “Collection of Charge on Junction Nodes from Ion Tracks”. In: *IEEE Transactions on Nuclear Science* 29.6 (1982), pp. 2024–2031. DOI: 10.1109/TNS.1982.4336490.

- [54] P.K.T. Mok and Ka Nang Leung. “Design considerations of recent advanced low-voltage low-temperature-coefficient CMOS bandgap voltage reference”. In: *Proceedings of the IEEE 2004 Custom Integrated Circuits Conference (IEEE Cat. No.04CH37571)*. 2004, pp. 635–642. DOI: 10.1109/CICC.2004.1358907.
- [55] Michele Muschitiello, Ali Zadeh, and Alessandra Constantino. *Irradiation characterization of EEE components for space application*. Oct. 2017.
- [56] Jackson Paddock. “Redesigning Power Systems on a Single Chip Micro Mote with Berkeley Analog Generator Low Dropout Series Regulator Generation”. MA thesis. EECS Department, University of California, Berkeley, May 2021. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-124.html>.
- [57] T Poikela et al. “Timepix3: a 65K channel hybrid pixel readout chip with simultaneous ToA/ToT and sparse readout”. In: *Journal of Instrumentation* 9.05 (May 2014), p. C05013. DOI: 10.1088/1748-0221/9/05/C05013. URL: <https://dx.doi.org/10.1088/1748-0221/9/05/C05013>.
- [58] Jeffrey Prinzie, Karel Appels, and Szymon Kulis. “Optimal Physical Implementation of Radiation Tolerant High-Speed Digital Integrated Circuits in Deep-Submicron Technologies”. In: *Electronics* 8.4 (2019). ISSN: 2079-9292. DOI: 10.3390/electronics8040432. URL: <https://www.mdpi.com/2079-9292/8/4/432>.
- [59] H. Quinn et al. “Radiation-induced multi-bit upsets in SRAM-based FPGAs”. In: *IEEE Transactions on Nuclear Science* 52.6 (2005), pp. 2455–2461. DOI: 10.1109/TNS.2005.860742.
- [60] D. Radaelli et al. “Investigation of multi-bit upsets in a 150 nm technology SRAM device”. In: *IEEE Transactions on Nuclear Science* 52.6 (2005), pp. 2433–2437. DOI: 10.1109/TNS.2005.860675.
- [61] L. Ratti et al. “Dark Count Rate Degradation in CMOS SPADs Exposed to X-Rays and Neutrons”. In: *IEEE Transactions on Nuclear Science* 66.2 (2019), pp. 567–574. DOI: 10.1109/TNS.2019.2893233.
- [62] Behzad Razavi. “The StrongARM Latch [A Circuit for All Seasons]”. In: *IEEE Solid-State Circuits Magazine* 7.2 (2015), pp. 12–17. DOI: 10.1109/MSSC.2015.2418155.
- [63] V. Re et al. “Total ionizing dose effects on the noise performances of a 0.13 μm CMOS technology”. In: *IEEE Transactions on Nuclear Science* 53.3 (2006), pp. 1599–1606. DOI: 10.1109/TNS.2006.871802.
- [64] J. A. Richardson, L. A. Grant, and R. K. Henderson. “Low Dark Count Single-Photon Avalanche Diode Structure Compatible With Standard Nanometer Scale CMOS Technology”. In: *IEEE Photonics Technology Letters* 21.14 (2009), pp. 1020–1022. DOI: 10.1109/LPT.2009.2022059.
- [65] G A Rochau et al. “Measurement and modeling of pulsed microchannel plate operation (invited)”. en. In: *Rev Sci Instrum* 79.10 (Oct. 2008), 10E902.

- [66] D.M. Schmidt et al. “Comparison of ionizing-radiation-induced gain degradation in lateral, substrate, and vertical PNP BJTs”. In: *IEEE Transactions on Nuclear Science* 42.6 (1995), pp. 1541–1549. DOI: 10.1109/23.488748.
- [67] D.M. Schmidt et al. “Modeling ionizing radiation induced gain degradation of the lateral PNP bipolar junction transistor”. In: *IEEE Transactions on Nuclear Science* 43.6 (1996), pp. 3032–3039. DOI: 10.1109/23.556902.
- [68] R.D. Schrimpf et al. “Hardness-assurance issues for lateral PNP bipolar junction transistors”. In: *IEEE Transactions on Nuclear Science* 42.6 (1995), pp. 1641–1649. DOI: 10.1109/23.488761.
- [69] J.R. Schwank. “Advantages and limitations of silicon-on-insulator technology in radiation environments”. In: *Microelectronic Engineering* 36.1 (1997). Proceedings of the biennial conference on Insulating Films on Semiconductors, pp. 335–342. ISSN: 0167-9317. DOI: [https://doi.org/10.1016/S0167-9317\(97\)00076-2](https://doi.org/10.1016/S0167-9317(97)00076-2). URL: <https://www.sciencedirect.com/science/article/pii/S0167931797000762>.
- [70] J.R. Schwank et al. “Radiation effects in SOI technologies”. In: *IEEE Transactions on Nuclear Science* 50.3 (2003), pp. 522–538. DOI: 10.1109/TNS.2003.812930.
- [71] Krishna Settaluri. “Photonic Links – From Theory to Automated Design”. PhD thesis. EECS Department, University of California, Berkeley, Apr. 2019. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-8.html>.
- [72] M.L. Simpson et al. “An integrated CMOS time interval measurement system with subnanosecond resolution for the WA-98 calorimeter”. In: *IEEE Journal of Solid-State Circuits* 32.2 (1997), pp. 198–205. DOI: 10.1109/4.551911.
- [73] Helmuth Spieler. “Introduction to radiation-resistant semiconductor devices and circuits”. In: *AIP Conference Proceedings* 390.1 (Jan. 1997), pp. 23–49. ISSN: 0094-243X. DOI: 10.1063/1.52282. eprint: https://pubs.aip.org/aip/acp/article-pdf/390/1/23/11412480/23_1_online.pdf. URL: <https://doi.org/10.1063/1.52282>.
- [74] G.M. Swift, D.J. Padgett, and A.H. Johnston. “A new class of single event hard errors [DRAM cells]”. In: *IEEE Transactions on Nuclear Science* 41.6 (1994), pp. 2043–2048. DOI: 10.1109/23.340540.
- [75] Jeffrey L. Titus. “An Updated Perspective of Single Event Gate Rupture and Single Event Burnout in Power MOSFETs”. In: *IEEE Transactions on Nuclear Science* 60.3 (2013), pp. 1912–1928. DOI: 10.1109/TNS.2013.2252194.
- [76] Jens Verbeeck, Paul Leroux, and Michiel Steyaert. “Radiation effects upon the mismatch of identically laid out transistor pairs”. In: *2011 IEEE ICMTS International Conference on Microelectronic Test Structures*. 2011, pp. 194–197. DOI: 10.1109/ICMTS.2011.5976845.

- [77] Zhongkai Wang and Elad Alon. “Analog Generators for SerDes Clock Generation and Distribution”. PhD thesis. EECS Department, University of California, Berkeley, May 2023. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2023/EECS-2023-36.html>.
- [78] Nicholas Werblun. “Closing the Analog Design Loop with the Berkeley Analog Generator”. MA thesis. EECS Department, University of California, Berkeley, May 2019. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-23.html>.
- [79] C.F. Wheatley, J.L. Titus, and D.I. Burton. “Single-event gate rupture in vertical power MOSFETs; an original empirical expression”. In: *IEEE Transactions on Nuclear Science* 41.6 (1994), pp. 2152–2159. DOI: 10.1109/23.340556.
- [80] Brad Wheeler et al. “A Low-Power Optical Receiver for Contact-free Programming and 3D Localization of Autonomous Microsystems”. In: *2019 IEEE 10th Annual Ubiquitous Computing, Electronics and Mobile Communication Conference (UEMCON)*. 2019, pp. 0371–0376. DOI: 10.1109/UEMCON47517.2019.8992964.
- [81] Bradley Wheeler. “Low Power, Crystal-Free Design for Monolithic Receivers”. PhD thesis. EECS Department, University of California, Berkeley, May 2019. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-36.html>.
- [82] Susan Wood et al. “Simulation of Radiation Damage in Solids”. In: *IEEE Transactions on Nuclear Science* 28.6 (1981), pp. 4107–4112. DOI: 10.1109/TNS.1981.4335684.
- [83] A. Wu et al. “Radiation-induced gain degradation in lateral PNP BJTs with lightly and heavily doped emitters”. In: *IEEE Transactions on Nuclear Science* 44.6 (1997), pp. 1914–1921. DOI: 10.1109/23.658962.
- [84] Siwei Xie et al. “Methods to Compensate the Time Walk Errors in Timing Measurements for PET Detectors”. In: *IEEE Transactions on Radiation and Plasma Medical Sciences* 4.5 (2020), pp. 555–562. DOI: 10.1109/TRPMS.2020.2981388.
- [85] Alan Yates. *Lighthouse; the Steam VR Tracking System*. 2016.
- [86] G.U. Youk et al. “Radiation-enhanced short channel effects due to multi-dimensional influence from charge at trench isolation oxides”. In: *IEEE Transactions on Nuclear Science* 46.6 (1999), pp. 1830–1835. DOI: 10.1109/23.819161.
- [87] Luya Zhang and Ali Niknejad. “A Multi-Parametric Microwave-Optical Biomolecular Sensor”. MA thesis. EECS Department, University of California, Berkeley, Dec. 2019. URL: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-159.html>.
- [88] James F. Ziegler, M.D. Ziegler, and J.P. Biersack. “SRIM – The stopping and range of ions in matter (2010)”. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 268.11 (2010). 19th International Conference on Ion Beam Analysis, pp. 1818–1823. ISSN: 0168-583X. DOI: <https://doi.org/10.1016/j.nimb.2010.02.091>. URL: <https://www.sciencedirect.com/science/article/pii/S0168583X10001862>.

Appendix A

SCM3x

The Single Chip Micro-Mote (SC μ M) Version 3 [11, 45, 52, 81] is a $2.5 \times 3 \times 0.3\text{mm}^3$ device designed and fabricated in TSMC65LP CMOS to be a fully integrated 802.15.4 transceiver with a Cortex M0 and a diagnostic ADC to monitor the chip’s internal temperature as well as interface with external sensors. This appendix includes information on the regulator used for digital power management and the sensor ADC subsystem which was designed with David Burnett. These subsystems first appeared in SC μ M3 and were reused in SC μ M3B and SC μ M3C.

A.1 Digital LDO

The low dropout regulator designed for digital (Figure A.1) was used twice within SCM3—once to power the Cortex M0 microprocessor, and a second time to power the hodgepodge of digital miscellany that is collectively referred to as auxiliary digital (auxdig).

The PMOS device was chosen due to a lack of headroom—the output of the LDO nominally sits at 1.0V, compared to the battery voltage down to 1.2V. The giant capacitor and resistor were an extremely janky attempt to slow down the amplifier for stability purposes, before Osama let it be known that there would be substantial decoupling capacitance on the output of the LDO. The area of the LDO ends up being vastly increased by the resistors and capacitors, which ultimately aren’t necessary; [56] did substantial work using the Berkeley Analog Generator to design LDOs with improved performance which didn’t require the additional passives. That said, the LDO is stable with and without decoupling capacitance on the output. The amplifier was an NMOS folded cascode, chosen for its high DC gain. The reference voltage is generated by a bandgap-produced current source used in conjunction with a resistive DAC [45].

Table A.1 provides a summary of the LDO’s typical performance under nominal operating conditions; the PSRR bandwidth doesn’t end up mattering much since the load is digital. RMS worst-case noise across corners—process, temperature, load, and battery voltage—was less than $1.05\text{mV}_{\text{rms}}$. For SCM3, issues with timing violations were observed in silicon when




Figure A.1: Annotated schematic of the digital LDO.

the output of the regulator dropped below 700mV.

Parameter	Simulated Value
Reference	1V
Battery	1.5V
Static Load	1mA
Static Error	-2.756mV
Static Current Consumption	4.763 μ A
Phase Margin	85.62 $^{\circ}$
PSRR	43.87dB
PSRR Bandwidth	461.9Hz
Load Regulation	96.50 μ
Amp Capacitor	3.60pF
Load Capacitor	452pF
Total Gate Area	756.6 μ m ²

Table A.1: Simulated performance of the digital LDO.

A.1.1 Cadence Locations

Library	Path	Relevant Cells
VDD	/tools/projects/lydiale/scum_v3/VDD	all
SCM_v3b_analog	/tools/projects/fil/scm_v3/SCM_v3b_analog	vdd_digital_ldo_v11, vdd_auxdigital_ldo_v8

Table A.2: Location of digital LDO libraries in Cadence Virtuoso

A.1.2 Known Idiosyncrasies

Initial measurements on the first SCM3 saw a dramatically bouncing supply ($> 200\text{mV}$). The chip in question was superglued to a PCB and wire bonded using the Swarm lab West Bond bonder with a board where surface moun components were soldered by hand. This error was not reproducible on additional chips and has never been observed since; for the purposes of an academic chip, we chalk it up to assembly issues.

Some SCM3C chips are unable to cold boot. First pass attempts at a solution observed that briefly connecting the digital output to a slightly higher voltage ($\approx 1.2\text{V}$) was sometimes sufficient to enable the chip to boot. This was referred to as “VDDD tap,” named for the “tap” of bringing a wire down on the board to briefly connect the output of the digital LDO (VDDD) to the external voltage. In November 2022, Jacob Louie, working under David Burnett, discovered a firmware workaround to the VDDD tap issue. In essence, when the processor would write a value to a register and immediately attempt to read from it, the read would always return 0. When booting, this means that the instruction pointer register would read out as 0, causing the program counter to jump there and crash the chip. The solution—which has been a consistent fix—is to inject delay between the write and read of a register in the form of no-ops in the assembly (printfs in the C code). While not directly relating to the digital LDO, the problem and its solution are notable.

A.2 ADC + PGA Sensor Interface

The sensor ADC subsystem is similar to the diagnostic ADCs found in industry chips; it is meant to monitor battery health, report the chip’s internal temperature, and enable SCuMto interface with external sensors (not at the same time). The subsystem, shown in Figure A.2 consists of a low dropout regulator (LDO), a programmable gain amplifier (PGA), and a 10-bit analog-to-digital converter (ADC) designed by David Burnett. Figure A.4 shows the internals of each component.

[45] describes the design of the PTAT reference. Within the PGA, a unit capacitance of roughly 20fF was used to mitigate mismatch effects while still meeting the speed requirements of the system. The reference for the PGA was generated using a PTAT with a different




Figure A.2: Block diagram of the sensor ADC subsystem.


temperature coefficient than the PTAT intended for temperature measurement. Given the extremely low power consumption of the PTAT, kickback from the switching action of the PGA introduced significant perturbation (peak 10%) on the PTAT voltage during conversion. Figure A.3 shows the data from a temperature sweep from 0°C to 80°C in the same temperature chamber used in Figure 3.2. [12] demonstrated the use of the the full front end with an H₂S gas sensor, with wireless information transmission at 2.4GHz. Figure A.5 shows the differential and integrated nonlinearity of the ADC with the PGA bypassed. Large spikes in DNL around apparent at bit flips within the DAC, which is somewhat expected given the binary DAC structure; the ADC is also nonmonotonic.

A.2.1 Code Base and Cadence Locations


The code base for SC μ M3C has changed significantly since the ADC was last characterized. That said, Titan Yuan was able to perform a first-pass ADC characterization with minor updates to the code found at https://github.com/PisterLab/scum-test-code/tree/adc/scm_v3c/sensor_adc. Several working examples using this code base can be found at https://github.com/PisterLab/scum-test-code/blob/adc/scm_v3c/run_me.py. The SC μ M3C User Manual provides a detailed explanation of the steps taken in the SC μ Mfirmware; a copy of the manual can be found by contacting current members of the group. Table A.3 shows the location of the sensor ADC both as it appeared in the initial design, and any copies which appeared in the SC μ M3x tapeouts.

Library	Path	Cells
sensoradc	/tools/projects/db/cadence3/sensoradc	all
SCM_v3b_analog	/tools/projects/fl/scm_v3/SCM_v3b_analog	TEMPCORE_TOPa_v5

Table A.3: Location of sensor ADC libraries in Cadence Virtuoso



(a)



(b)

Figure A.3: (a) ADC code readout with a PGA gain setting of $2V/V$. The slope corresponds to approximately $1.2^{\circ}\text{C}/\text{LSB}$. (b) The number of measurements associated with each temperature measurement, taken with a TMP102 digital temperature sensor. Temperature precision was 0.01°C .

A.2.2 Known Idiosyncrasies

If using the on-chip finite state machine, the most significant bit (MSB) of the ADC will always read out as low (Figure A.6). This issue was reproduced across a variety of chips across several years and three generations of graduate students, and is a confirmed bug which spawns from the fact that the reset of the ADC is tied to the soft reset of the entire chip. There are two methods to circumvent this issue: Both require that the ADC be controlled via GPIO loopback, where the GPIOs that are connected to the ADC control signals are set to be both inputs and outputs. This allows the Cortex M0 to drive the GPO aspect of the GPIOs, while the GPI component feeds to the ADC finite state machine. The first method involves triggering a soft reset after reach ADC sample. The second method extends the ADC settling time via scan chain, allowing the MSB to settle.

When triggering an ADC conversion via on-chip FSM, the first reading after a power cycle

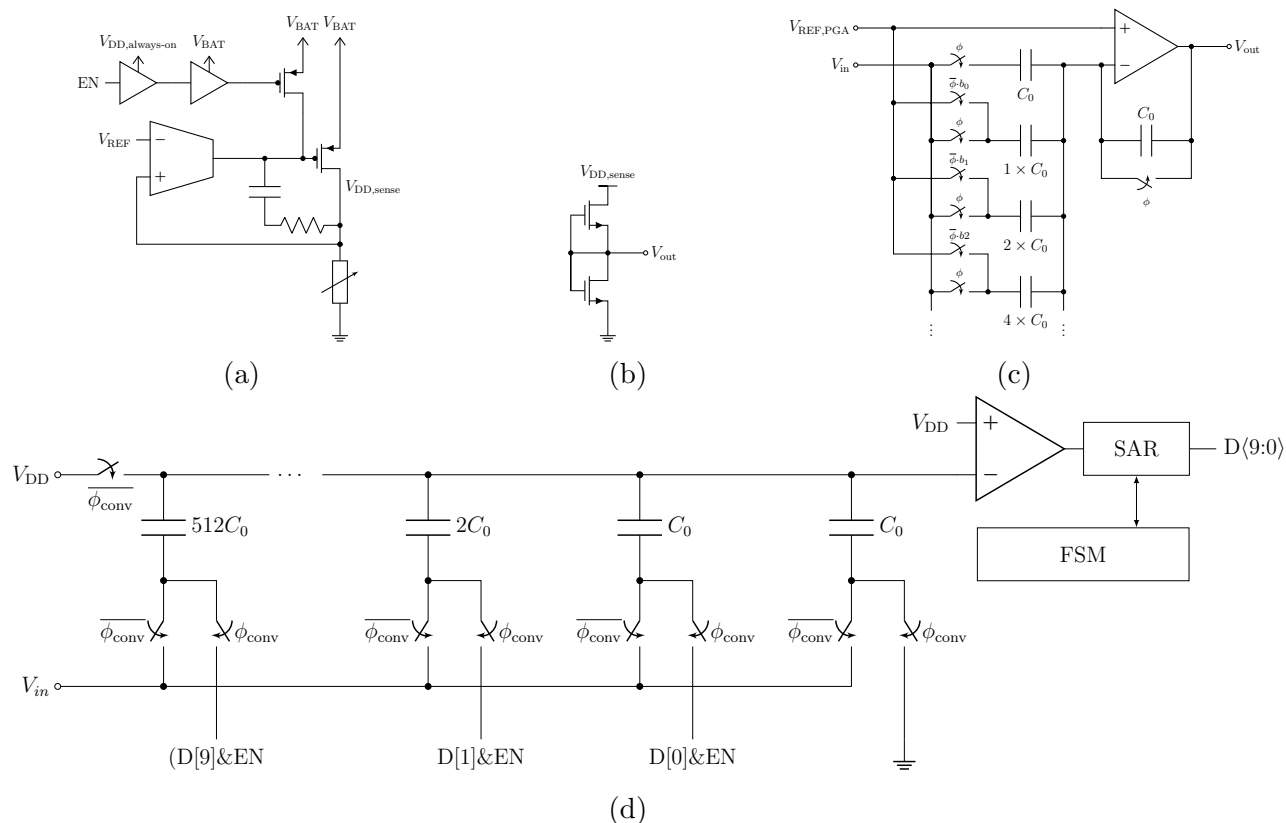


Figure A.4: Subblocks of the front end. (a) Regulator. (b) PTAT. Body connections are to ground. [45] (c) Programmable gain amplifier. (d) Successive approximation register analog to digital converter.

is always all ones save for the MSB (see the above), giving a reading of 511. Unfortunately, this seems to be due to an incorrect startup at the initial boot of the chip. More specifically, the on-chip FSM doesn't reset properly before taking the first reading.

The simulated noise of the PGA is less than 2LSB; the measured noise contribution of the PGA is more than 5LSB. Furthermore, the PGA at times does not cooperate with input voltages exceeding 700mV, with the output changing with little apparent correlation to the input voltage.

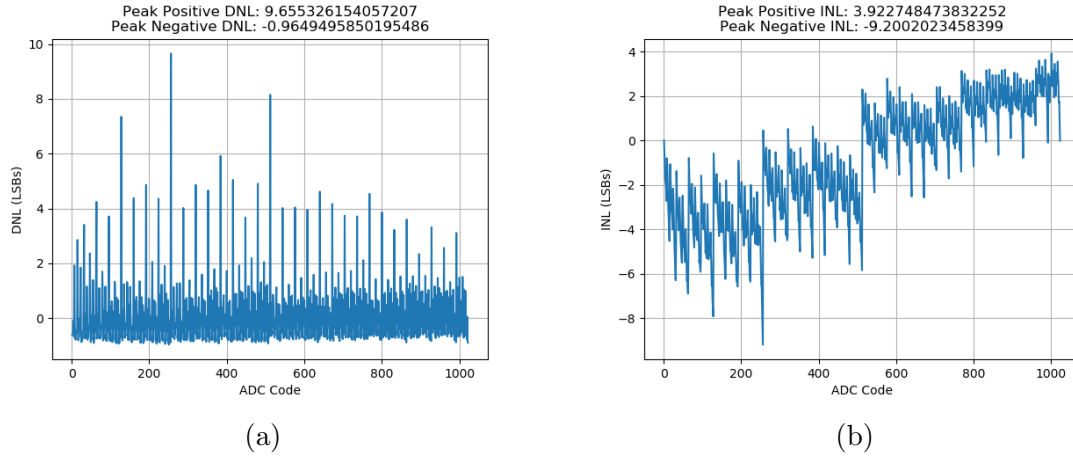


Figure A.5: Sensor ADC (a) DNL and (b) INL. There are clear spikes at bit flips which are consistent with binary DACs, with some nonmonotonicity with the DAC.

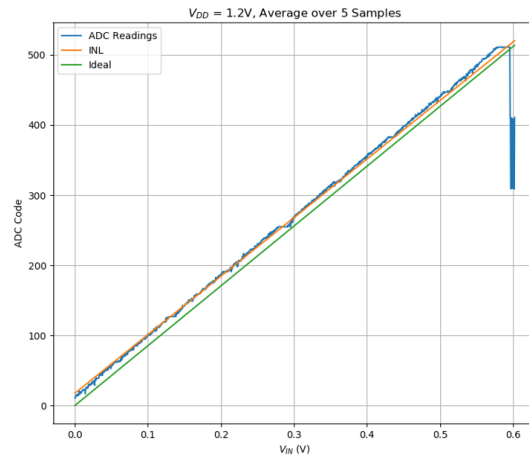


Figure A.6: The ADC output code versus input voltage. The nominal FSR is $V_{DD,sensor} = 1.2V$, and each data point is the result of 5 averaged samples. Unfortunately the raw data for this plot has been lost.

Appendix B

SCM3C Digital Flow Documentation

This appendix is documentation for the flow used for the Single Chip Mote Version 3C. It assumes you have access to TSMC65LP on BWRC infrastructure. Specific paths within the BWRC file system were correct as of 2021.

B.1 Code Base

The top-level repository for SC_μM3C digital lives in <https://bwrcrepo.eecs.berkeley.edu/SCuM/scum3b-digital>. The associated BWRCRepo group is <https://bwrcrepo.eecs.berkeley.edu/groups/SCuM/>. Contact Titan Yuan or Yu-Chi Lin for access or request it directly on the repository page. To clone the repository,

```
git clone git@bwrcrepo.eecs.berkeley.edu:SCuM/scum3b-digital.git
```

Using csh, source the cshrc with the following command:

```
cd lp-setup
csh    // For those using bash rather than csh
source .cshrc_tsmc65_rf
```

For those who have taken EECS251, CS250, or EE241B prior to 2019, the repository structure is similar to that used in those classes.

- lp-setup: Where all the Verilog simulations, synthesis, and place-and-route happen
 - dc-syn: Location from which to run synthesis scripts
 - icc-par: Location from which to run place-and-route scripts
 - pt-pwr: Location from which to run PrimeTime power scripts
 - pt-pwr-syn: Location from which to run post-place-and-route PrimeTime power scripts

- setup: Contains the .tcl files used in setting up synthesis and place-and-route
- src: Contains all Verilog source files
- vcs-sim-gl-par: Location from which to run post-place-and-route simulations
- vcs-sim-rtl: Location from which to run RTL simulations

B.2 Simulation

B.2.1 RTL

The directory `vcs-sim-rtl` contains the setup from which to run tests. For this particular project, we don't have any tests that aren't just in Verilog test benches, so no special handling (e.g. for C-based, Python-based scripts) is required. Long story short:

1. Modify the Makefile variable `vsracs` to include your Verilog modules and test bench files.
2. `make run`. This compiles your Verilog and runs any test benches you've included in the Verilog source files. This doubles as a sanity check for syntax, disconnected wires, incorrect wire widths, and a whole host of things that could cause problems during synthesis. This also generates a .vpd file for you to view your waveforms using DVE.
3. `make dve` to view waveforms. This will reference the .vpd file created during the run, and it's pretty handy if you don't want to rerun a simulation to get a single signal out.

Disclaimer: If you have hundreds to thousands of signals in a bus, some may not display correctly in the waveform viewer because it didn't save properly in the .vpd. Fear not, your simulation still ran correctly.

B.2.2 Post-P&R Verification

This section uses ModelSim Version 10.3b. For those looking for the Verilog used in the final tapeout, the result can be found in "backup38" in the ICC results.

B.2.2.1 Getting Started

- Copy the `output.v` and `.sdf` from `icc-par/yourRunDirectory/results` as needed to the `modelsim-gl-par` directory (or skip this if you don't want a local copy)
- Edit the `.mpf` file to point to these files (or add them in the GUI later)
- Launch with: `vsim <mpf file>`. A screen like Fig. B.1 should appear.




Figure B.1: ModelSim window after launch, showing Library and Project tabs on the left and Wave tab on the right.

On the right side, switch to the Library tab, right click 'work', hit delete, then hit OK. Right click empty space and select "New > Library" then hit OK with the default settings, as shown in Fig. B.2

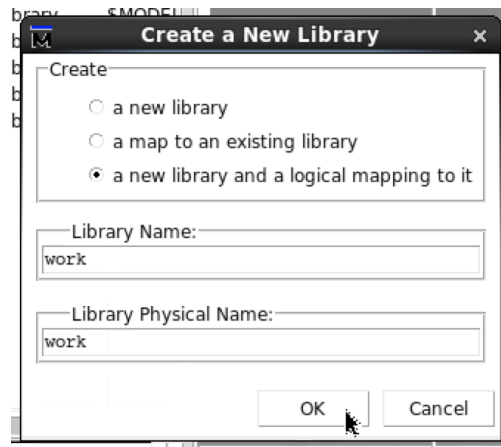


Figure B.2: Creating a new work directory

Switch back to the Project tab, right click empty space, and select “Compile > Compile All”. This step should take less than a minute to complete.

Notes about these files:

- Most .v files are outside of the Modelsim working dir. If any fail, check their paths.
- The .rom.v has its properties set to “Do Not Compile” so a green check mark won’t appear next to it.
- the SDF file path is specified in top_testbench.v on a line starting with “\$sdf_annotate(”

Inside top_testbench.v, there’s a `bootload_source_select` flag. A little bit after should have `emulate_optical()` and `emulate_3wb()` commands on separate lines. These must all match to test bootloading optical or wired 3WB programming. Either:

- set `bootload_source_select = 1'b0`, uncomment `emulate_optical()`, and comment out `emulate_3wb()`, OR
- set `bootload_source_select = 1'b1`, comment out `emulate_optical()`, and uncomment `emulate_3wb()`.

B.2.2.2 Load

- Switch back to the Library tab and open the “work” library. Now that all files have been compiled, all contained verilog functions are listed here.
- Scroll to `top_testbench`, right click + “Simulate”. This will take a minute or so to load the simulation but won’t start it yet. You’ll know it’s done loading with the message “SDF Backannotation Successfully Completed” in the transcript at the bottom.

- Switch to the new sim tab (left) to display the signal hierarchy; you can select signals from the new center column by right click + “Add Wave”

Save sets of plotted signals by bringing the wave tab into focus, then going “File > Save Format...”. Load sets of plotted signals by bringing the wave tab into focus then going ”File > Load > Macro File”. A pre-made `wave.do` file is in the ModelSim directory with a set of signals.

If no signals are visible, go back to the left-side Library tab, right-click `top_testbench`, and select “Simulate without optimization”. This will take a while but the full signal hierarchy should be visible once it’s finished.

B.2.2.3 Run Sim & Examine Results

Click “run” (Fig. B.3) or type “run” in the transcript box and hit enter, then wait for output as specified by the testbench. The optical bootload process takes half a day but the 3wb bootload process takes less than an hour. If there’s a lot of simulation output, the transcript window will only display the last few hundred lines but all output is logged to a file called “transcript” in the simulation directory.




Figure B.3: The “run” button circled in red

B.3 Scripts

B.3.1 Setup

`common_setup.tcl` This is, in general, the file to reference when changing variables. If there is a variable you can’t identify, chances are it’s set in here, `icc_setup`, or `constraints`.

Edit cases:

- Library and techfile references
- Power domain parameters
- Plan group parameters
- Min/max metal routing layers

constraints.tcl This does a few things, and these are all change cases:

- Defines and sets clock constraints
- Creates power domains as per what was set in `common_setup`
- Sets driving and loading cells for inputs/outputs and clock sources
- Max fanout, transition times, capacitance for inputs, outputs, and internal nets

dc_setup.tcl This suppresses a bunch of warnings and uses a bunch of the variables set in the Makefrag and dc scripts to change settings during DC. We didn't change this from what it was already set.

- Number of cores for computation
- Setting constraints file location
- Setting RTL source file location

dc_setup_filenames.tcl Don't change this. This determines where all the files go in the file structure during a design compiler run.

icc_setup.tcl This assigns a boatload of of ICC variables. A lot of variables here deal with optimization settings as well as floorplan creation settings, constraints files, etc.

If you want any source files referenced, this is where they're defined.

B.3.2 Synthesis

This flow uses the Synopsys Design Compiler. All scripts can be found in `/dc-syn/dc_scripts`. We strongly recommend taking a look at Section B.5.2 before proceeding.

Makefile This is what gets referenced when you type `make` in the `dc-syn` directory. A number of the variables set in `vars` are specific to the process and should be set appropriately. Assuming all of those variables have been set correctly, the only change that needs to be made during synthesis is `vsrcls`, which lists all the Verilog modules which need to be synthesized.

For reasons unknown, the variables set in this file can't be applied past the scope of just this file (i.e. we haven't been able to use it in the tcl files). That said, the variables *are* used within this file for referencing and should be maintained.

dc.tcl This runs the mapping from Verilog to gate-level Verilog. This references variables set in `common_setup`, `dc_setup`, and `dc_filename_setup`. Change this only when

- You want to include any kind of gating or logic optimization
- You want to spit out any more reports

find_regs.tcl We don't modify this. This finds all the registers in the design and writes a VCS command to set them to a known state at simulation time 0. Fair warning that this can mask functional problems if you use this in simulation but your real-life device doesn't initialize correctly. If you'd like to use it, the file contains a description of how to use it in simulation in the comments.

fm.tcl This is used to verify the synthesized netlist vs. your Verilog. We didn't modify this as it came out of the box. At the moment, the most recent version of Formality available on the BWRC servers is not compatible with the most recent version of DC (the latter being more up-to-date than the former). To get around this, we deprecated the version of DC we used to run Formality.

B.3.3 Place-and-Route

We strongly recommend taking a look at Section B.5.3 before proceeding.

init_design_odl.tcl Don't modify this file. This creates the floorplan and sets various floorplan constraints. In this case, we set the floorplan input to a user file. Files used beyond setup:

- floorplan.tcl
- pin_placement.tcl
- pin_placement_continued.tcl

create_plangroups_dp.tcl Again, don't modify this file. This creates plan groups for hierarchical synthesis and PNR. Files used beyond setup:

- plangroup_constraints.tcl
- macro_placement.tcl

create_odl_dp.tcl Don't change this. We actually modified this so it wouldn't create an on-demand netlist—the settings in ICC made placement it modifies made it unworkable and would make PNR fail. Files used beyond setup:

- pg_constraints.tcl

routeability_on_plangroups_dp.tcl Don't modify this. This performs custom power network synthesis and analysis, and then performs the logical connections with power as well. Files used beyond setup:

- power_network_synthesis.tcl
- power_network_analysis.tcl
- pg_constraints.tcl

pin_assignment_budgeting_dp.tcl Don't modify this. We're not entirely sure what this does.

place_opt_icc.tcl Don't modify this. This optimizes the placement of the cells within your design while abiding by your plan group settings.

clock_opt_cts_icc.tcl Don't modify this. This synthesizes the clock tree in a DC-style fashion. Your clock settings should be carried over from DC.

User beware that this software is notoriously unstable; you should ensure that all of your scripts complete correctly in their entirety before proceeding. For example, the software encountered memory issues with specific settings configurations, and the command for fixing hold time violations would never be reached.

clock_opt_psyn_icc.tcl Don't modify this. This performs further optimizations on the clock tree.

clock_opt_route_icc.tcl Don't modify this. Places and routes the physical clock tree with additional adjustments for physical placement.

route_icc.tcl Don't modify this. This performs initial signal routing between cells.

route_opt_icc.tcl Don't modify this. This performs optimizations on signal routing between cells and inserts buffers as necessary.

chip_finish_icc.tcl Don't modify this. This inserts filler cells (DCAP, FILL) specified in common_setup.tcl.

metal_fill_icc.tcl Don't modify this. This inserts metal fill depending on the DRC rules fed into the tool. We didn't actually end up using this.

outputs_icc.tcl This generates all of the outputs for use outside of this miserable godforsaken tool. At the moment, it generates the GDS, the various types of Verilog, and all the other output.X files available. Modify this only if you'd like to add or remove a particular output.

floorplan.tcl This script deals with

- Metal layer routing direction
- Min/max metal layers
- Defining do-not-place locations (if they aren't already defined in some LEF)
- Defining do-not-route locations (if those aren't already defined in some LEF)
- Floorplan creation (see documentation for `create_floorplan` in ICC shell)
- Setting macro location and keepout margins (macro placement)

pin_placement.tcl This contains pin physical constraints for pins in the analog toplevel. Note that this removes any unused/no-Conn'd pins. This is generable by using

`pin_scraper.ipynb`

- `pin_scraper.ipynb`

pin_placement_unused_kept.tcl This contains more physical constraints for pins which appear in the analog toplevel. The purpose of this was to keep some of the pins we removed in `pin_placement.tcl` (e.g. analog config pins, analog scan chain) as a "just in case".

plangroup_constraints.tcl This script

- Creates the plan groups based on the specifications in `common_setup`
- Adds no-place padding to the plan groups (which it didn't obey for mine, I had to manually put down placement blockages)
- Creates voltage areas associated with those plan groups

macro_placement.tcl This does nothing beyond stating that macros, after being placed, should not be moved. Actual macro placement occurs in `floorplan.tcl`.

pg_constraints.tcl I'm realizing now that "pg" and "plangroup" are very easy to confuse (oops). At any rate, this

- makes the logical connections between the power and ground nets as per the specifications of the net names and cell names.

No routing is performed in this (a fact which I learned quite late and after a great deal of confusion and frustration). If you want separate power domains and want power/ground to connect to something other than a default top level power net, you should specify those logical connections before making any logical connections to the top level power net.

power_network_synthesis.tcl This deals with the physical creation and routing of the power network.

- Creates power straps associated with each voltage region. This should be modified based upon need
- Creates power rings for each power domain using the locations specified in `common_setup`
- Performs power routing for both standard cells and macros
- Checks power routing connectivity and spits out information in a report

power_network_analysis.tcl This analyzes the rails associated with the power nets. We barely used this and this should be changed for any future runs.

- See ICC documentation on the command we used for greater flexibility.

pin_scraper.ipynb This was a script that took in a .txt file (the source text file seems to have been lost, but see Section B.4.1 on how to generate it) and spat out the pin names and their locations in the layout. It also ran a comparison against a Verilog module to see if there were any pins that should not be routed because they don't connect to anything, e.g. like a lot of the scan chain bits do. It then proceeded to assign target routing locations for all of the digital pins based on where they appear in analog, in line with the grid limitations and minimum spacing requirements of the process. Currently it naively assigns everything to a single metal layer.

B.4 Running the Flow

B.4.1 Analog Blocks

This section describes how to incorporate analog blocks into a digital top flow. Blocks which call for more specificity in routing than just avoiding shorts (e.g. don't route X metal

over this entire block) require the use of abstract views (Section B.4.1). Blocks which don't require that specificity can go ahead and use the Cadence Virtuoso utility for generating LEF files (Section B.4.1).

Pin Location Pulling

Note that this process does not give you metal layer information for the pins. To generate a list of (x,y) coordinates of pin label locations from a layout:


- Open Calibre PEX on the layout cell of interest.
- Under Outputs → Netlist → Format choose DSPF, then additionally check HSIM under the same dropdown.
- Choose R-only Extraction type to reduce the netlist size (DSPF does not work with No R/C).
- Run PEX. The PEX netlist is the file of interest for the next steps.
- Each line that contains a pin label coordinate starts with *|P.
- To parse into a file containing only pin names and locations, run:

```
grep -rnw netlist_filename -e "*|P" >> output_file
```

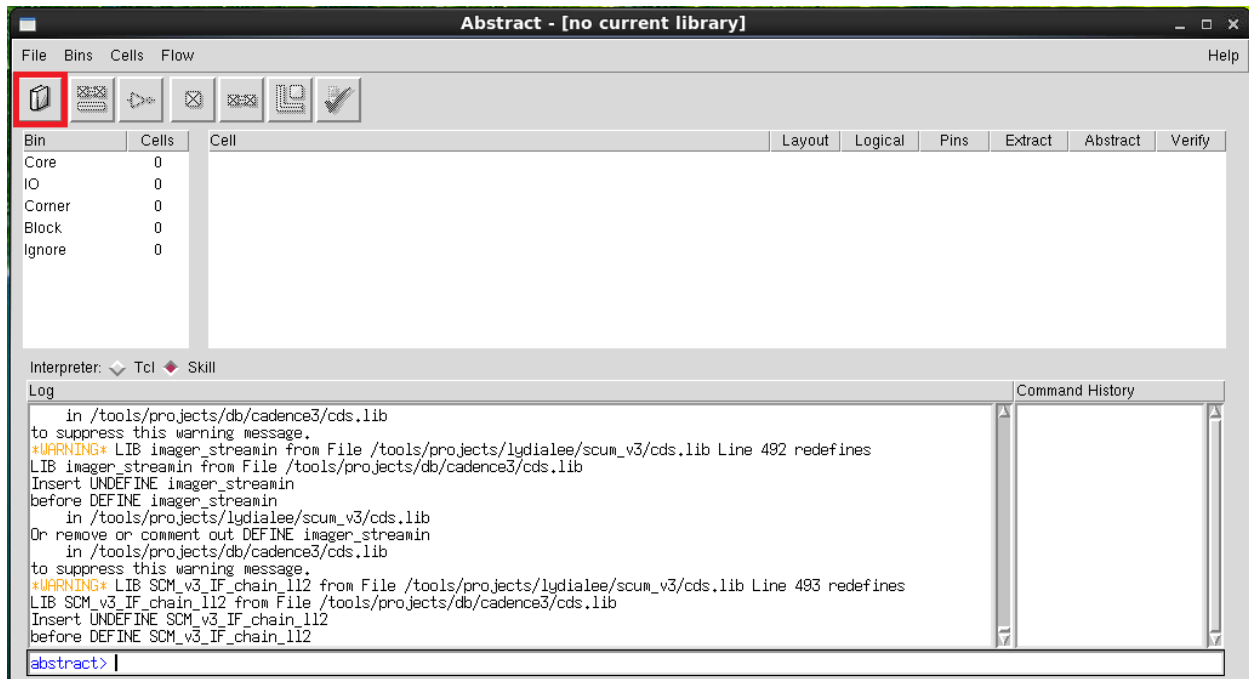
Abstract Generation

We never actually used this with SCM, but the instructions from Sidney Buchbinder have been recorded here for documentation purposes. Either with Flow → Abstract or the Abstract button This flow uses Abstract, a software which should come along with Cadence Virtuoso. Wherever you typically run Virtuoso from, type

```
abstract &
```




It's important that this is run in the same location as Virtuoso because it will use the cds.lib to determine which libraries and cells are visible. Click the Library button (marked in red) and select the library containing the analog block(s) you're interested in.



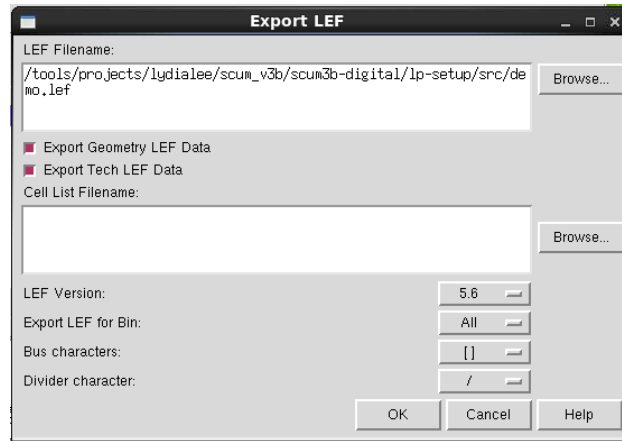
On the left you should see a list of Bins. When you load the library, Abstract will automatically place the cells in that library into those bins. Which bin a cell is in determines how it's handled when generating abstracts. You can make your own bins with your own settings, but we didn't need to.

We only use Block and Ignore. The cells you care about should go into Block and everything else should be in Ignore. To migrate cells, select them in the list, then at the top choose Cells → Move. The window this opens will take a second to load, but you'll end up with a list of bins to move your selection to. Choose the appropriate bin and hit OK.




Now Abstract options need to be set up. The Abstract options are specific to the technology and deal with how different routing layers are handled, e.g. no M1 routing over the entire block, M2 routing is permitted over the block (the place-and-route tools later should avoid unintentional shorts). Once set, these options can be saved in an abstract options file and reused.

There is a template abstract options file at `lp-setup/gen_stdcells/abstract.options`. Loading it goes File → Import → Options. The requirements of the specific block being abstracted will dictate the settings that need to be changed. This document has a gitturdun walkthrough (Section B.4.1) of how to set up Abstract Options below, but the block-specific sections have documentation for the Abstract Options settings for that analog block.



Abstract Options

Either with Flow → Abstract or the Abstract button



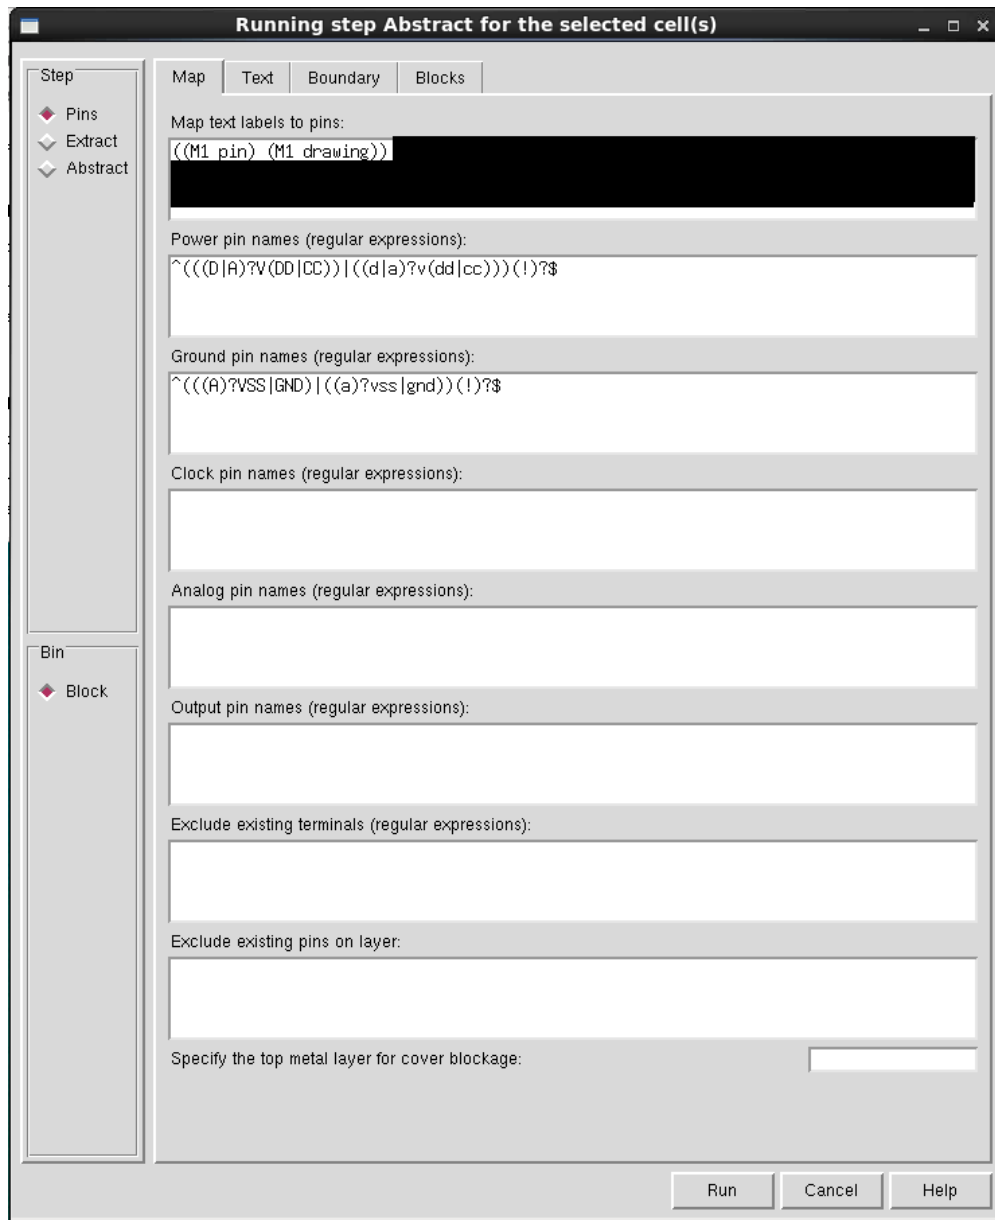
you'll open up the dialog to decide all of your settings. There are 3 steps:

1. Pins
2. Extract
3. Abstract (silly, I know)

which you'll set before running.

Pins This will take the labels of your pins or labels and use the text for it to establish pins. In all layouts, pins need some kind of label to be recognized for LVS. Which layer the labels and pins are on is process-specific. In our case, we use a label type layer, i.e.

LayerName label. If this layer is incorrect in your settings, Abstract will freak out and give you incorrect output (if it gives you output at all).



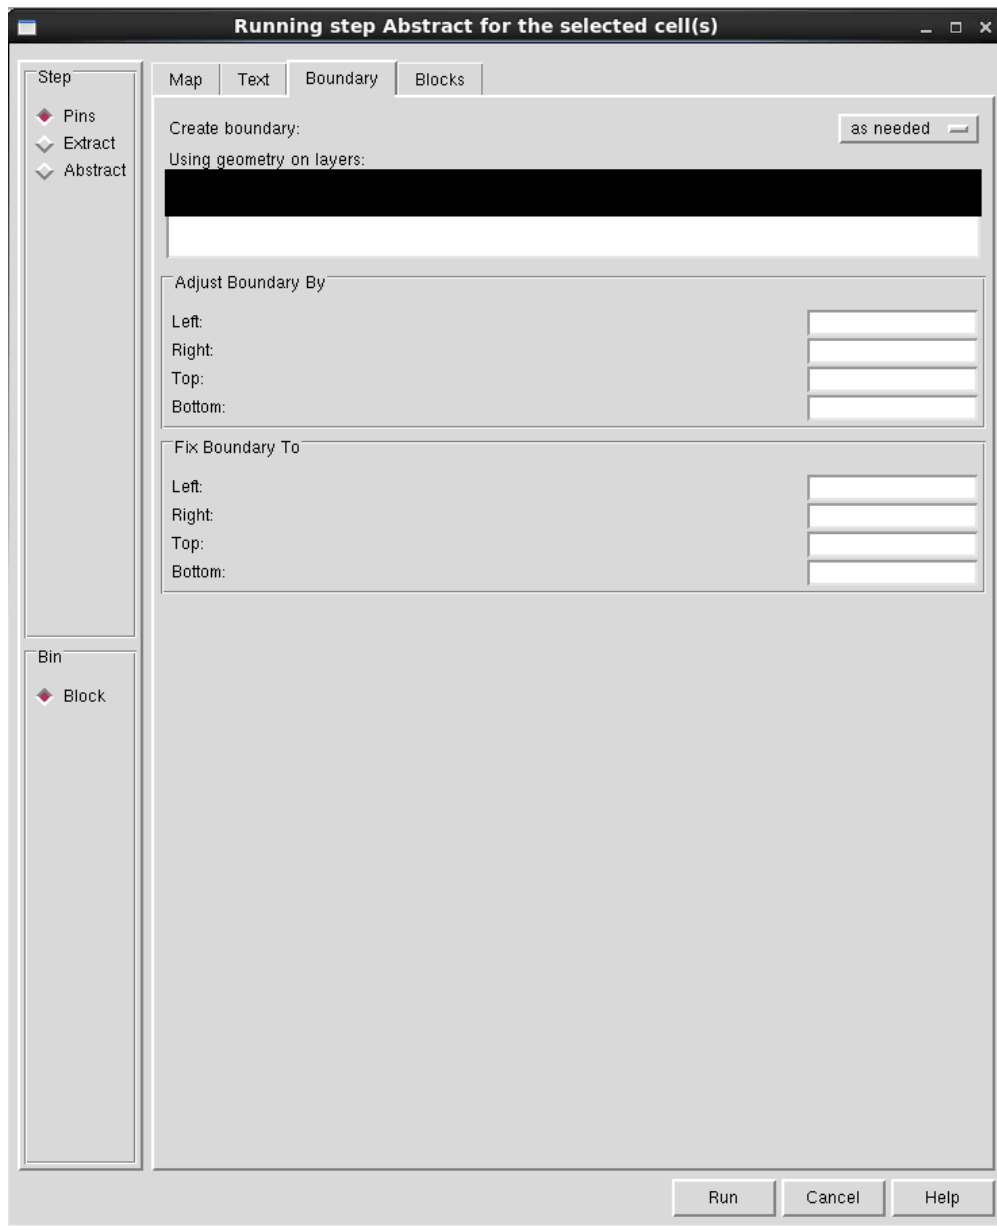
For the box 'Map text labels to pins', inserting `((layerName label) (layerName drawing))` will take any pins on the label layer and map them to shapes on the drawing layer. Our use didn't require anything beyond this, but the manual on this is surprisingly concise and complete if our use in the template isn't satisfying.

In place-and-route tools, you need to distinguish between power/ground pins and potentially clock pins. When you run Pins, Abstract will grab all the pins/labels/ports/equivalent

in the layout and use regular expressions to determine what type of pin it is. Everything that isn't power or ground is a signal pin.

The Text tab will let you manipulate the text of your layout labels if you have any strange formatting. For example, Cadence uses angular brackets for bus notation, but you may want square brackets for the LEF file to match up with Verilog syntax; this is the place to change that.

The Boundary tab lets one adjust the PR (place-and-route) boundary. Most if not all processes should have a layer for the PR boundary in layout.




The only setting we really care about is the drop-down at the top right for creating a boundary:

- off won't create the size of the block
- always takes the maximum extent of all of the listed layers of any purpose in the process and use that as the PR boundary
- as needed uses whatever pre-existing PR boundary there is and assumes the block fits within that boundary

We use "as needed". We recommend always drawing a PR boundary, having the bottom left corner be at (0,0), and making it a clean-ish size, e.g. to the nearest micron.

The Blocks tab, we don't know what it does.

Extract Extract will go through the geometry in the layout and figure out the coordinates of the polygons attached to each pin. There are separate tabs for power and signal pins. In both, you need to list all of your metal layers and potentially all of your via layers. For us, this automatically populates.



One should specify the geometry for all layers (examples in the template and described in the manual). Failing to do so will interpret every layer of that type (including pins, labels, etc.) as part of the geometry. In the case where a text label extends past the edge of the polygon associated with it, the router may end up not connecting properly because it thinks the text region is part of the polygon.

Be sure to take into account layers specific to non-wiring components, e.g. metal resistors and capacitors.

Connectivity is always strong because the connectivity of the pins in layout should always be strong (what weak is, we have no idea).

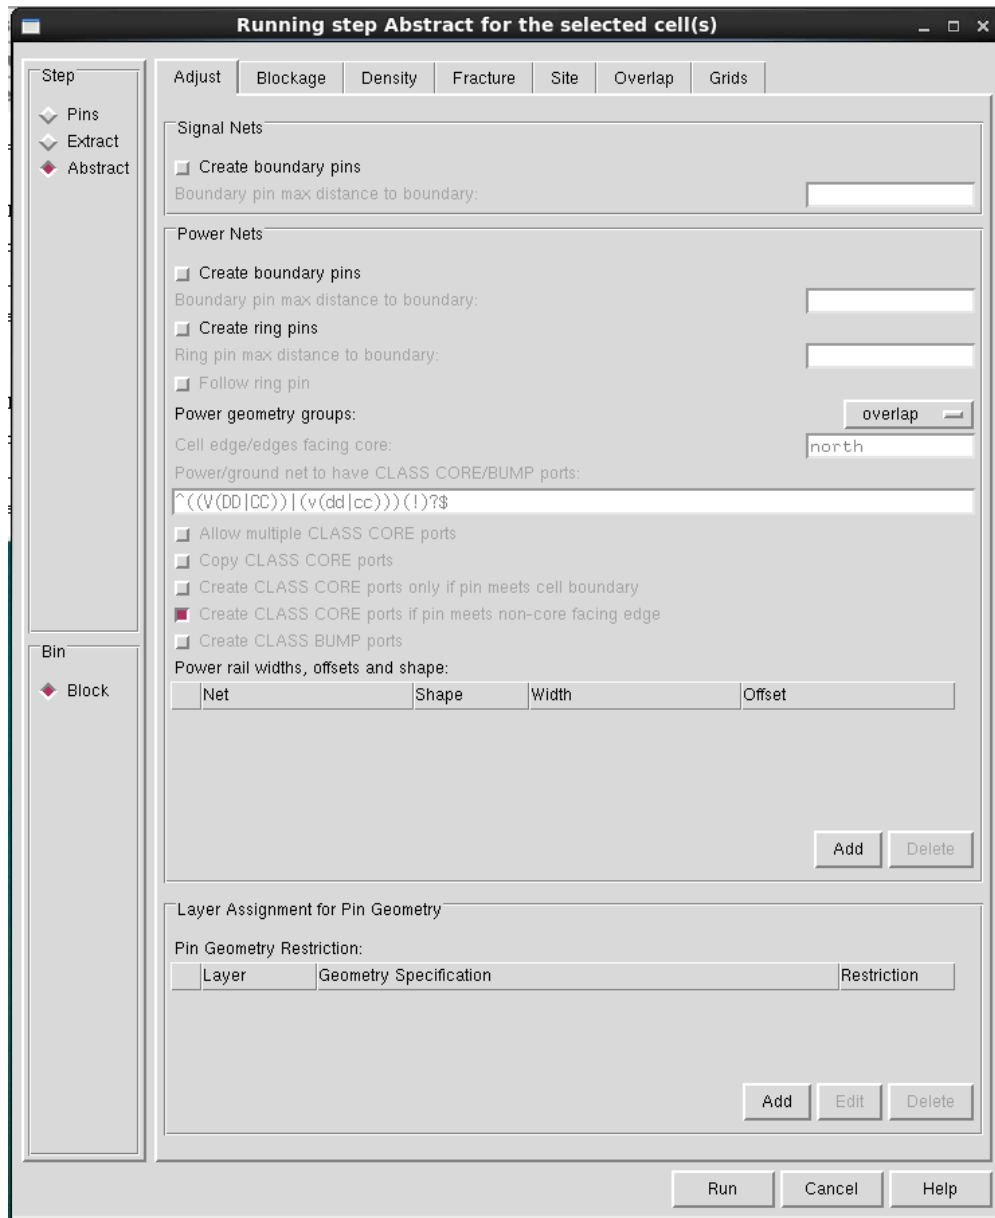
For creating pins, there's no harm in creating pins for everything except for time and file size. If one knows one will never want to route to a particular layer, one can turn off creating pins for that particular layer.

Max depth (0 - 32) defines how far down the hierarchy one wants to extract metals for. Again, this becomes a time and file size constraint. Because this process has a lot of metal routing layers, we don't need to go too many layers down.

In the Antenna tab, there's no reason not to choose all of the antenna options. These are intended to avoid antenna rule DRC violations and you don't want to miss any of them.

The General tab should be autopopulated and just describes which vias connect which metal layers.

Abstract



In the Adjust tab, creating boundary pins will only tell the PR tool about a small region at the edge of the block. Not creating boundary pins tells the PR tool that the connectivity can be done over the entire region of the block.

CORE/BUMP deals with I/O pins. Hopefully we can ignore this because TSMC handles it?

The Blockage tab is the one we care about. You should list metal and via layers of interest. This defines keepout regions for routing. The blockage types are

- Cover doesn't allow routing of the specified layer over the block (though edge routing is fine). This is useful for highly sensitive blocks, e.g. a radio

- Detailed provides more granularity for keepout; it prevents routing that leads to undesirable shorting, but doesn't prevent routing over the entire block
- Shrink we don't know what it does

The Pin Cutout check box makes it such that you can always route a via down to the metal. The single check box for cutting the window around the pins does the exact same thing, so of course Cadence had to have it in two places. Just in case, have both checked.

We don't know what the rest of the tabs in Abstract do.

To save your Abstract options for future use, it's File → Export → Options.

Hopefully by this point you have all the abstract files for your cell. Take a look at the abstract file for the cell and eyeball it to make sure nothing was horribly messed up. Unfortunately, we don't have a more efficient way of debugging this aside from manually inspecting the abstract output file.

Fair warning that this explanation is by no means complete—so far our only solution has been to RTFM (which can be found in `/tools/cadence/IC/<your version of IC>/doc/abstract/abstract.pdf`).

LEF Generation

We never actually used this with SCM because of issues with LEF incorporation. The instructions here from Sidney Buchbinder are purely for documentation purposes. If you've used an abstract view, in the Abstract home window, File → Export → LEF. Choose to export geometry and the tech data. When choosing the LEF version, 5.6 works for us, though other groups have had issues and needed to mark as 5.7 (and that somehow worked even though the LEF files were otherwise identical). We set the bus character as square braces and divider character as the slash.

The following struck-out paragraph is included for completeness since it is purported to work in other process nodes. However, we were not able to get this functioning in TSMC65LP on the BWRC infrastructure. ~~Alternatively, in the Virtuoso home window, File → export → LEF. Set the LEF file name appropriately in the lef directory. Select the appropriate library and output cell(s), and set the Output View(s) to layout or abstract, depending on whichever you're using. When choosing the LEF version, 5.6 works for us, though other groups have had issues and needed to mark as 5.7 (and that somehow worked even though the LEF files were otherwise identical). We currently check the box for No Technology, but we don't know what it does. We also don't know how or if this handles bus characters.~~

Ideally you have the LEF file at this point to be used with the rest of the digital flow.

LEF Incorporation

We never actually figured out how to get the LEF to Milkyway conversion functioning! The conversion would inexplicably drop layers in certain locations with no indication why.

You have a LEF file, but Synopsys tools use Milkyway files. Synopsys has yet another separate piece of software called Milkyway with which you can convert the LEF file to a MW file. The Makefile to do this is in the `lef` directory. You'll want to modify the Makefile

- Change the library name to something unique; if a directory of the same name already exists, the software will throw an error.
- Point to the appropriate toplevel.
- `make`

Nominally, `make` will generate a ton of files and generate the Milkyway file for the cell(s) you've specified.

B.4.2 Synthesis

We strongly recommend taking a look at Section B.5.2 before proceeding.

In the directory `dc-syn`,

- Modify the `Makefile` such that the `toplevel` and `toplevel` instance are those of what you're synthesizing
- `dc_setup.tcl`, modify the variable `RTL_FILES` to point to the Verilog of interest.
- `make`

B.4.3 Place-and-Route

We strongly recommend taking a look at Section B.5.3 before proceeding.

In the directory `icc-par`, you can run `make dp_hier` to run the first portion of the place-and-route (no optimization, strictly for examining device placement). `make ic` will finish the rest of the place-and-route where the previous command left off. Alternatively, you can run `make run` and have both run with no intermediate break.

B.5 Known Idiosyncrasies

B.5.1 General

- Cadence Virtuoso 6.1.7 experiences issues with streaming in/out GDS files for SCM3C which cause it to fail LVS. We've confirmed stream-in/out correctness in Virtuoso 6.1.8 and 6.1.5.

B.5.2 Synthesis

- Ensure that *all* of your source Verilog is synthesizable! This means you need to go through and ensure that there are no `initial begin` statements which are active come synthesis time. (Example: `correlator.v`) Rather than warning us, Design Compiler crapped out and tied every single output in the result to logical 0.
- Version compatibility between DC and Formality is a serious issue. In particular, the auto-generated files from newer versions of DC cannot be read by slightly older versions of Formality. This is an issue even for versions which are 1-2 iterations apart.

B.5.3 Place-and-Route

- Check your log files to make sure your scripts finished correctly. We discovered that certain setting configurations (e.g. specifying the location of the pin HCLK, specifying a specific no-route gap at the top of the aux digital plan group) would cause the tool to crash very consistently and in the same respective location in the scripts. We're 100% confident the issue wasn't us; the software would literally "encounter a fatal error", spit out a stack trace, and then fail.
- Don't expect any kind of response from Synopsys if you're from academia and asking a question on their help website. We aren't paying enough for them to answer in a timely let alone thorough fashion.
- The documentation in some of this software either doesn't exist for certain commands (so the command exists, but the manual has no entry), is spotty and rife with typos ("What does that line even mean?!" - Us, on many occasions), or is fantastically and very clearly written. Good luck.

B.6 Signoff

B.6.1 Generating a Floorplan GDS

Congratulations! You now have a placed-and-routed cell. The script `outputs_icc.tcl` should have produced a GDS file and an output Verilog file for your cell, which you'll now want to use to check DRC and LVS.

B.6.2 Running LVS

Running LVS (layout-versus-schematic) checking is significantly more annoying than DRC, mostly because you have to deal with text rather than just layer correctness at this point.

B.6.2.1 Verilog Module Generation

This ideally should have been handled in `outputs_icc.tcl`. Depending on if you have buses, you may want to remove the `-split_bus` option from your Verilog creation command in the tcl file.

If the option is turned on, it'll separate `signal[1:0]` into `signal[1]` and `signal[0]`, which may actually cause netlist confusion later on.

B.6.2.2 SPICE Netlist Generation

There are two parts to this section

1. Fix the Verilog so Virtuoso doesn't get confused
2. Generate the SPICE netlist from the fixed Verilog

These are wrapped up in the `icc-par/Makefile` for convenience.

Skim the Makefile for `lvs_fix`. You may need to modify the line which starts with `v2lvs` to include the correct technology SPICE files.

```
make lvs_fix
```

This creates and runs a Perl script to modify your Verilog so it's LVS-friendly, and it generates the SPICE netlist from that Verilog using Calibre's V2LVS software. As it is, that SPICE netlist will be in the same location as the original netlist, only with "fixed.2" at the end of the file name.

Because this chip uses SRAM (i.e. macros), we need to include the associated SPICE netlist files at the top of our V2LVS output. I cannot list the names here for NDA reasons, but you can find them in

```
/tools/projects/oukhan/SCuM3/digitalsynthesis/okSCMDigital/custom/
```

and they will all end with `.spi`

DO NOT COPY THESE FILES! These have fairly particular restrictions on their distribution.

B.6.2.3 Streamed-In GDS Fixes

In our particular technology, we want the text of pins to be on the pin-type layer (as opposed to drawing-type), and we want `<>` rather than `[]` for buses. Unfortunately, the GDS stream-in didn't do either for you, even if you checked the boxes which said they'd do that.

In Cadence Virtuoso, open the layout of the cell. Close all other cells. In the CIW window (basically the Virtuoso command prompt)

```
load("mygitrepo/lp-setup/setup/fix_icc_labels.il")
fix_icc_labels()
```

Depending on how your GDS was exported, you also might not have pins for power routing. Add them as necessary.

Save the layout.

Appendix C

Optical Receiver

C.1 Background

The optical bootloader on the Single Chip Mote Version 3C was designed as a means to wirelessly transfer program data to SRAM for low power relative to radio based bootloaders. The first iteration of the bootloader was designed by Brad Wheeler and Andy Ng [81] with an active power consumption of $1.52\mu\text{W}$ and a standing power consumption of 640nW . The clock and data recovery (CDR) scheme was pulse width modulation (PWM), with long pulses corresponding to a logical HIGH and short pulses corresponding to logic LOW. However, lab measurements revealed a limited range for programming and high transmit irradiance ($1.7\text{mW}/\text{mm}^2$) to maintain a bit error rate of 1 payload bit error per hundred 64kB programming cycles ($\approx 2 \times 10^{-8}$); a decrease in irradiance to $1.5\text{mW}/\text{mm}^2$ produced a bit error rate of $\approx 6 \times 10^{-4}$.

It was later discovered in [80] that the same front end could be used for outside-in lighthouse localization [85] of the Single Chip Mote with HTC Vive V1 lighthouse base stations. Lighthouse localization is a method for a sensor with minimal compute to determine its azimuth and elevation relative to a base station. The base station flashes a synchronization pulse, then sweeps a beam at a fixed angular velocity horizontally; given the fixed angular velocity and time between the synchronization pulse and when the swept beam strikes the sensor, the receiver can calculate its azimuth relative to the base station. The process is repeated for a sweep in the vertical direction for a calculation of elevation angle.

In practice, programming reliability is hit-or-miss at best, with programming ranges limited to low singles of centimeters; this gets worse when chips are sealed in UV epoxy. The use of a high-powered infrared transmitter with the potential to cause permanent eye damage is also less than desirable, and so some time was spent redesigning the optical receiver for improved sensitivity.

Parameter	Programming	Lighthouse V2
BER	2×10^{-8}	10^{-3}
Yield	$\geq 99\%(2.6\sigma)$	
Bandwidth	N/A	1.84MHz
Area	130 $\mu\text{m} \times 130\mu\text{m}$	
Power	minimize	

Table C.1: Optical receiver target specifications.

Reader beware: All work and values presented in this presentation are based on simulation, *not* lab measurements.

C.2 Chip Summary

The chip (Figure C.1) was designed in TSMC’s 65nm LP process through the Berkeley Wireless Research Center.

It includes an analog front end, digital for clock and data recovery (CDR), scan chain, a test photodiode array of 12 diodes routed in parallel (separated to meet density DRC), a pared-down signal chain consisting of only a preamplifier and comparator, and a test SPAD device with a quenching circuit that was largely unsimulated and added at the last second.

The analog front end of the chip was designed to fit into 16,900 μm^2 , with the total chip area coming out to 1mm \times 1mm. Table C.2 provides a snapshot of the chip’s simulated performance in relation to the target applications. All values described in this appendix are simulated.

Parameter	[81]	Programming	Lighthouse V2
Irradiation	1.5mW/mm ²	<20mW/mm ²	<10 μ W/mm ²
BER	6×10^{-4}	2×10^{-8}	10^{-3}
Yield	Unknown	$\geq 99\%(2.6\sigma)$	
Bandwidth	Unknown	N/A	1.84MHz
Area	130 $\mu\text{m} \times 130\mu\text{m}$		
Power	1.5 μ W	15 μ W	

Table C.2: Simulated optical receiver versus the tested version in [81].




Figure C.1: Chip layout screenshot. The analog front end fits in $130\mu\text{m} \times 130\mu\text{m}$.

C.2.1 I/O

Pin	Pad Location (Counter Clockwise)
SPAD_VOUT	IN0
SPAD_VDDL	IW0
SPAD_VDDH	IW1
MAIN_VDD	W0
MAIN_LPF_VBN	W1
ACLK	W2
IBG	W3
MAIN_VCMFB	W4
MAIN_PREAMPINN	W5
ASC_VDD	W6
VSS	S0
IPD	S1
CDR_DATA_IN	S2
CDR_CLK_IN	S3
CDR_RESET	S4
CDR_DATA_OUT	S5
CDR_CLK_OUT	S6
ASC_VDD	E6
ASC_PHIB	E5
ASC_PHI	E4
ASC_SCAN_IN	E3
ASC_LOAD	E2
ASC_SCAN_OUT	E1
ASC_RESETB	E0
VSS	N6
COMP_OUTN	N5
COMP_OUTP	N4
VDDIO	N3
PARED_VDD	N2
PARED_PREAMPINP	N1
PARED_PREAMPINN	N0

Table C.3: Optical receiver chip I/O and associated pad locations.

Pin Name	Description	Type	Domain
ACLK	Comparator clock. Note that this is not internally level shifted and so should be in the appropriate power domain.	Input	MAIN + PARED
ASC_LOAD	Raise to latch bits into scan chain. This should be performed after the desired bits have been clocked in.	Input	VDDIO
ASC_PHI	Input clock for clocking the analog scan chain. Should be the inverse of ASC.PHIB.	Input	VDDIO
ASC_PHIB	Inverted input clock for clocking the analog scan chain. Should be the inverse of ASC.PHI.	Input	VDDIO
ASC_RESETB	Reset the analog scan chain to all 0. Active low.	Input	VDDIO
ASC_SCAN_IN	Input data to clock into the analog scan chain. CHECK IF LSB OR MSB GOES IN FIRST	Input	VDDIO
ASC_SCAN_OUT	Output data from the scan chain. This should match ASC_SCAN.IN.	Output	VDDIO
ASC_VDD	Always-on supply. Nominal 0.8V.	Power	AON
CDR_CLK_IN	Clock input of the CDR. Signal chain requires always-on supply.	Input	VDDIO
CDR_CLK_OUT	Output clock from clock and data recovery. Nominally 5MHz. Runs off always-on supply.	Output	VDDIO
CDR_DATA_IN	Positive input to the clock and data recovery. Signal chain requires always-on supply.	Input	VDDIO
CDR_DATA_OUT	Output data from clock and data recovery. Should change analogous to output clock. Runs off always-on supply.	Output	VDDIO
CDR_RESET	Reset for CDR. Active high (rising edge). Requires always-on supply.	Input	VDDIO
COMP_OUTN	SR latched negative output of the comparator in the selected signal chain. Select buffer runs off of always-on supply. Selected signal chain runs off of either main or pared supply.	Output	VDDIO
COMP_OUTP	SR latched negative output of the comparator in the selected signal chain. Select buffer runs off of always-on supply. Selected signal chain runs off of either main or pared supply.	Output	VDDIO
IBG	Input current fed to resistor ladder. Nominal 1uA.	Input	MAIN
IPD	Connected to the output of the photodiode test structure. The board should include a linear TIA structure at the output. Bias voltage should not exceed core voltage (nominally 0.8V; rated up to 1.2V)	Output	N/A
MAIN_LPF_VBN	Gate bias voltage produced by constant-gm circuit; tunable via analog scan chain. Can be overridden with an externally applied voltage in the event of boot-up issues with the constant-gm.	I/O	MAIN
MAIN_PREAMPINN	Negative input terminal of the preamp in the main signal chain. Can be used to test the resistor DAC; the mux; and the low-pass filter.	I/O	MAIN
MAIN_VCMFB	Voltage for preamp common mode feedback. Nominal 0.5V.	Input	MAIN + PARED
MAIN_VDD	Supply for main signal chain. Nominal 0.8V.	Power	MAIN
PARED_PREAMPINN	Inverting input to the preamplifier in the pared-down signal chain. Requires pared-down supply to be on.	Input	PARED
PARED_PREAMPINP	Noninverting input to the preamplifier in the pared-down signal chain. Requires pared-down supply to be on.	Input	PARED
PARED_VDD	Supply for pared-down signal chain. Nominal 0.8V.	Power	PARED
SPAD_VDDH	Nominally >3.5V. NOTE THAT THIS WILL PROBABLY DAMAGE THE CHIP. Do NOT use the same chip for SPAD testing and standard signal chain testing.	Power	SPAD_HI
SPAD_VDDL	Nominally >3.5V.	Power	SPAD_LO
SPAD_VOUT	Output voltage of the SPAD test structure.	Output	SPAD_HI; SPAD_LO
VDDIO	I/O voltage to communicate with external devices. Nominally rated up to 2.5V; though operation up to 3.3V has been demonstrated (reliability not guaranteed at that point)	Power	IO
VSS	Ground. 0V.	Ground	All

Table C.4: Description of optical receiver I/O.

C.2.2 Scan

Signal Name	Description	Bits (MSB:LSB)
main_tune_tia_res	Thermometer coded. Adjusts the TIA gain by changing the value of the feedback resistor. 0.0000 corresponds to the highest resistance value; 1.1111 corresponds to the lowest resistance value. Nominal resistance (i.e. transimpedance) can be calculated as $[6 - (\text{No. of ones})] * 500\text{kohm}$. Nominal condition all 0.	4.3.2.1.0
main_enb_lpf	1=disables the low pass filter in the main signal chain connected to the output of the TIA.	5
main_tune_lpf_constgm	Binary coded. Adjusts the corner frequency of the low-pass filter. At a low level; adjusts the resistance of the constant-gm for generated bias voltage. Corner frequency increases with binary value. 0000 corresponds to roughly 245kHz bandwidth; 1111 corresponds to roughly 1.5MHz bandwidth. Nominal condition is all 0.	9.8.7.6
main_tune_bgrdac_coarse	Adjusts the value of the resistive DAC by a single large step of 32/512. 0 = higher resistance; 1 = ideally 0 resistance. Nominal condition is 0.	10
main_tune_bgrdac_fine	Binary coded. Adjusts the value of the resistive DAC by steps of 1/512. 0x00 corresponds to the highest resistance; 0xff corresponds to ideally 0 resistance. Nominal condition is all 0.	18.17.16.15.14.13.12.11
main_tune_preamp_constgm	Binary coded. Adjusts the resistance of the constant-gm for preamp biasing in the main signal chain. 0x0 = highest resistance; lowest gm. 0xf = lowest resistance; highest gm. Nominal condition is all 0.	22.21.20.19
main_en_comp	1 to enable the main signal chain comparator with the gated clock. 0 to disable. Nominal condition is 1.	23
main_sel_preamp_in	Chooses the source of the negative input of the preamp in the main signal chain. 0 sources from the low-pass filter; 1 sources from the resistive DAC. Nominal condition is 0.	24
pared_tune_preamp_constgm	Binary coded. Adjusts the resistance of the constant-gm for preamp biasing in the pared-down chain. 0x0 = highest resistance; lowest gm. 0xf = lowest resistance; highest gm. Nominal condition is all 0.	28.27.26.25
pared_en_comp	1 to enable the pared-down signal chain comparator with the gated clock. 0 to disable. Nominal condition is 1.	
sel_comp_out	Selects which signal chain to choose from. 1 = pared-down signal chain; 0 = main signal chain.	
tune_zero_count_min	Binary value. For PWM; the minimum number of clock cycles a pulse (going off rising edges of the CDR clock) must cover in order to be considered a 0. Do NOT use 0.	38.37.36.35.34.33.32.31
tune_one_count_min	Binary value. For PWM; the minimum number of clock cycles (going off rising edges of the CDR clock) a pulse must cover in order to be considered a 1. Do NOT use 0.	46.45.44.43.42.41.40.39

Table C.5: Optical receiver scan bits.

C.3 Analog Front End

The analog front end here is a fairly run-of-the-mill topology with a transimpedance amplifier for current-to-voltage conversion, a high pass filter implemented as the signal minus

the low pass filtered version of itself, a preamplifier for that subtraction in addition to more gain and isolation from comparator kickback, and finally a clocked comparator. Similar to the leading edge detector (confusingly given the acronym LED) of Section 2.4, there is the option of using a constant threshold to determine if the incoming signal is a logical high or low. Figure C.3 shows an example of the analog front end operation with the low pass filter in use.




Figure C.2: A block diagram of the analog front end of the optical receiver.




Figure C.3: General operation of the analog front end for a pulse width modulated input.

C.3.1 Photodiode

For the photodiode, we estimated a capacitance of roughly $1\text{fF}/\mu\text{m}^2$. The photodiode had an area of $25\mu\text{m} \times 100\mu\text{m}$, corresponding to a signal of 10nA for programming at a distance of 10cm with the same diode used in [81], 600pA for lighthouse localization at 5m , a DC

ambient current of 56nA , and shot noise of roughly $190\text{nA}_{\text{rms}}$. This range was chosen to account for 3σ variation in the TIA bias point.

C.3.2 Transimpedance Amplifier

The transimpedance amplifier (TIA) is an inverter with resistive feedback (Figure C.4), similar to that found in the original work by Brad and Andy, but without the additional biasing devices. It was chosen for its simplicity as well as the fact that there was already a BAG schematic design script written for it in the context of bit error rate accounting for gain, bandwidth, noise, and offset.




Figure C.4: The self-biased transimpedance amplifier.

For sensor front ends, maximizing gain early in the signal chain is desirable to minimize input-referred noise, within the bandwidth spec. The gain of the TIA is tunable via the scan chain for $500\text{k}\Omega \times 1, \dots, 6$. At the highest gain setting, we measured a gain of $2.95\text{M}\Omega$, a 3dB bandwidth of 1.6MHz , an input-referred noise of $200\text{pA}_{\text{rms}}$. Figure C.5 shows an eye diagram of the output of the TIA, loaded with the rest of the hardware, with a 1.84Mbps OOK SCM-programming input. The eye opening is 17.5mV .

C.3.3 Low Pass Filter

The low pass filter in Figure C.6 is implemented as a $G_m - C$ filter, where the pole frequency is defined by the transconductance G_m of the amplifier and the output capacitance C (assuming parasitics are negligible).

$$\omega_p = \frac{G_m}{C} \quad (\text{C.1})$$

For the constant- g_m circuit in Figure C.6b, the PMOS devices are matched and the right side NMOS device is large relative to the NMOS on the left in order to provide gate biasing for

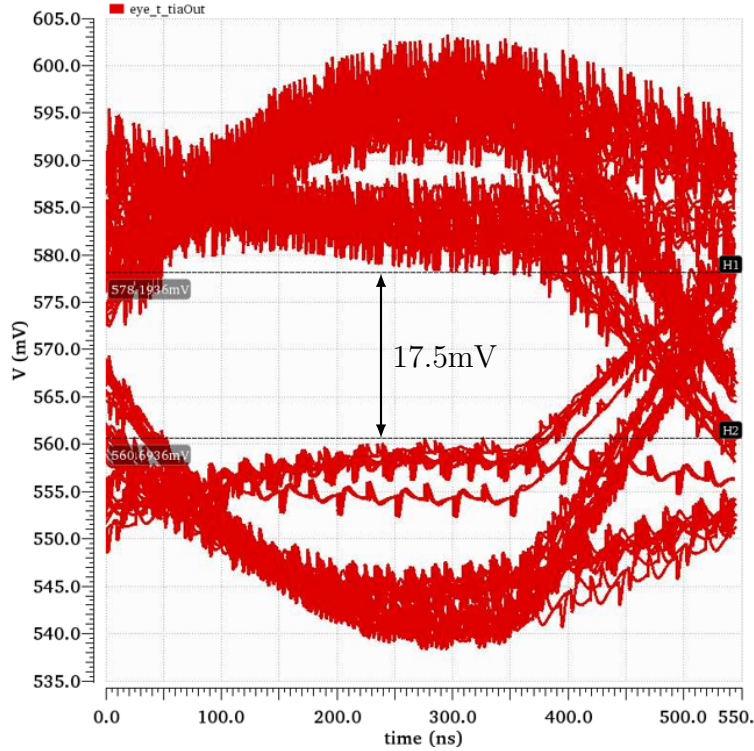


Figure C.5: An eye diagram of the output of the TIA when it's fed a 1.84Mbps OOK signal to program SCM with a 10nA signal. The eye opening is roughly 17.5mV.

an NMOS tail current. Assuming the bias current on both sides of the circuit are identical, the g_m of the left side NMOS device is approximated with Equation C.2.

$$g_m \approx \frac{2}{R_{\text{TUNE}}} \left(\frac{\sqrt{K}}{\sqrt{K} - 1} \right) \quad (\text{C.2})$$

The tail devices of each amplifier are intended to mirror the NMOS devices in the constant- g_m circuit, so we can tune the corner frequency by adjusting the G_m of the OTAs (Equation C.1).




Figure C.6: (a) The second order low pass filter with a corner frequency of $\omega_0 = \frac{G_m}{C}$. The G_m of each operational transconductance amplifier is controlled with the tunable constant- g_m circuit shown in (b), used for adjusting the tail current of each OTA.

The gain error introduced by the finite gain of the OTAs is 2.5%, with a 3dB cutoff frequency of 245kHz base, and -50.2dB attenuation as $\omega \rightarrow \infty$.

C.3.4 Preamplifier

The preamplifier is intended as an additional gain element as well as a shield between earlier-stage elements and kickback from the comparator. It is implemented as the fully differential amplifier in Figure C.7a, with an additional common mode feedback loop to ensure appropriate biasing (Figure C.7b). The tail bias was provided by a nearby constant- g_m circuit like that in Figure C.6b, also tunable via the scan chain to account for process variation.

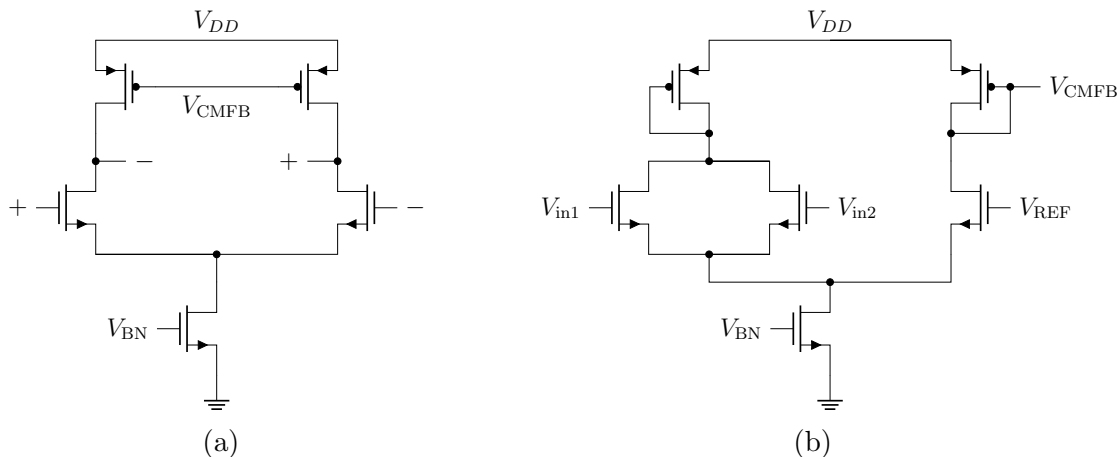


Figure C.7: (a) The main amplifier for the preamp. (b) The common mode feedback amplifier used in the preamp to maintain biasing.

Here, V_{REF} is the common mode bias, externally supplied (though I *meant* to get it from a bandgap reference eventually). The reference is nominally 500mV. The gain of the preamp is simulated at 25V/V, with a bandwidth of 3.2MHz. The input-referred noise contribution is 22pArms. Accounting for 2.6σ for 99% yield, the input-referred offset was simulated at 1.6nA.

Looking to the eye diagram in Figure C.8, we see significant asymmetry in the shape of the eye. The upper portion of the eye is 130mV, whereas the lower edge of the eye only reaches -15mV —a difference of more than $8\times$. This is because for OOK signaling, the DC level of the input shifts in the absence of some form of encoding, e.g. 4b5b; the same issue was encountered with the original optical receiver in [81].

As an added added precaution or get-out-of-jail free card, we included a DAC with an analog mux with the optional of an external override to substitute the low pass filter.

C.3.5 DAC

The DAC was *meant* to be implemented as a bandgap reference-generated $1\mu\text{A}$ current, fed through a resistive DAC to produce a voltage. While the resistive DAC is present, the reference is not. In this case, the voltage is read out from the top of the resistive ladder, which can be set via an 8-bit fine tuning DAC and a 1-bit coarse tuning DAC (Figure C.9). Each resistor element is roughly $1.17\text{k}\Omega$ for a maximum of 600mV and a minimum of 264mV—slightly off from the ideal 262mV due to finite switch resistance.




Figure C.8: An eye diagram of the output of the preamp when the TIA is fed a 1.84Mbps OOK signal to program SCM with a 10nA signal. The eye opening spans from 130mV to -15mV because of DC imbalance, suggesting 4b5b is necessary for any OOK.




Figure C.9: The resistive ladder intended for a $1\mu\text{A}$ bias current.

As it is, the current source was never implemented and the output of the DAC was simply padded out for external override.

C.3.6 Comparator

For programming, we assume a 5MHz sampling clock source from SCM. For lighthouse localization, we assume a 20MHz clock, sourced from SCM or elsewhere. It's implemented as a strongarm latch [62] followed by an SR latch (Figure C.10).




Figure C.10: (a) The strongarm latch and (b) subsequent SR latch used in the clocked comparator.

The input-referred noise was simulated at $1.1\text{pA}_{\text{rms}}$, with an input-referred offset of 130pA accounting for 3σ . The average power consumed, including the power used to drive it, was simulated as $1\mu\text{W}$ while running from a 20MHz sampling clock.

C.4 PWM Scheme

In pulse width modulation (PWM), pulses longer than a time $T_{\text{threshold}}$ constitute a HIGH bit, with nonzero pulses of shorter duration constitute a LOW bit. The original design relied on the timing of an asynchronous digital delay cell in an effort to conserve space and power. Rather than rely on essentially a one shot timer, we opted to take advantage of the fact that the output of the analog front end is already sampled by the digital clock (5MHz for programming and 20MHz for lighthouse localization). From here, we implement our PWM

scheme in hardware description language (HDL), enabling easier reuse and scaling with more advanced process nodes.

The CDR block can be found at https://bwrcrepo.eecs.berkeley.edu/SCuM/optical_bootloader/-/blob/master/verilog/cdr_pwm.v. The block is implemented as the finite state machine (FSM) using Verilog; Table C.6 is documentation of the digital interface.

At a high level, the CDR scheme determines the length of a pulse by the number of consecutive clock cycles the input remains high, n_{pulse} . From there, we look to Equation C.3 to determine what the data is (or if it even constitutes data).

$$\text{data_out} = \begin{cases} 1 & n_{\text{pulse}} \geq \text{one_count_min} \\ 0 & \text{zero_count_min} \leq n_{\text{pulse}} < \text{one_count_min} \\ \text{N/A} & n_{\text{pulse}} < \text{zero_count_min} \text{ (no clock)} \end{cases} \quad (\text{C.3})$$

Looking at the Verilog with about 5 years of additional experience and hindsight, several potential issues immediately pop out:

- The inputs which determine the thresholds for what constitutes HIGH versus LOW versus noise are not internally latched, and so they must be held stable during operation.
- There is no failsafe to ensure `zero_count_min < one_count_min`.

C.5 SPAD Test Structure

Single photon avalanche detectors (SPADs) are diodes operated in Geiger mode which—when perturbed by an incident photon—exit the metastable state at which they have been biased and provide an avalanche current to mark the event.

Sensitivity is quantified by wavelength-dependent photon detection probability (PDP)

$$\text{PDP}(\lambda) = P(\text{avalanche}|\text{event}) \cdot \text{QE} \quad (\text{C.4})$$

where the event is the absorption of a photocarrier and QE is the SPAD’s quantum efficiency. Accounting for device fabrication, the photon detection efficiency (PDE) is the photon detection probability, multiplied by the fill factor of the SPAD.

Uncorrelated device noise is characterized by the average rate of events in the absence of photons—known as the the dark count rate (DCR)—measured in counts per second (Hz). DCRs today are on the order of singles to hundreds of counts per second at room temperature, though exposure to high total doses of ionizing radiation increases both median DCR and its spread when measured across multiple SPADs [23, 61]. Temperature dependence follows from its two primary mechanisms, with band-to-band tunneling dominating at low temperatures and trap-induced noise dominating at higher temperatures. Lower temperatures correspond to lower DCR at an exponential rate; while processing dictates the specifics of the exponential relationship, [64] measured 20°C per order of magnitude of DCR.

Name	IO	Width	Description
MAX_CYCLES	parameter		The maximum number of cycles expected from a single pulse.
clk	input	1	Input clock. This should be the same as the clock used for the comparator. Data is sampled on the rising edge of the clock.
reset	input	1	Asynchronous active high reset.
data_in	input	1	Input PWM stream.
zero_count_min	input	$\text{ceillog}_2(\text{MAX_CYCLES})-1$	The minimum number of cycles required to recognize an incoming pulse as a valid bit. This is not internally latched and so must be held stable during operation.
one_count_min	input	$\text{ceillog}_2(\text{MAX_CYCLES})-1$	The threshold number of cycles to distinguish between an incoming 0 (short pulse) and 1 (long pulse). This is not internally latched and so must be held stable during operation.
clk_out	output reg	1	Output clock which ticks high whenever an incoming bit is determined to be a 1 or a 0.
data_out	output reg	1	Output data. Edge-aligned with clk_out.

Table C.6

Dead time is the interval after avalanche is triggered during which a SPAD can no longer detect incoming photons. The self-sustaining nature of the SPAD’s avalanche current necessitates a quenching circuit to stop the avalanche and reset the device to a detection-ready state. For megabit communication links, dead time approaching tens of nanoseconds—readily achievable with modern fabrication methods—is sufficient to avoid significant error due to dead time.

Afterpulsing occurs when carriers captured in traps during the initial photon-induced avalanche breakdown are released after the device has been quenched, triggering avalanche once more. Afterpulsing probabilities are often low enough ($< 0.1\%$) to be negligible in overall analysis.

The device in this layout was added at the last possible minute and was unsimulated. The SPAD layout itself was an adaptation of the devices designed by Luya Zhang and fabricated in TSMC28 through the Berkeley Wireless Research Center for her MS thesis [87]. It is unknown how Luya created DRC clean circular elements for her layout.

C.6 Code Base and Cadence Locations

All HDL can be found on the BWRC Repo at https://bwrcrepo.eecs.berkeley.edu/SCuM/optical_bootloader. For Cadence Virtuoso libraries, all work was done on the Berkeley Wireless Research Center infrastructure using Cadence Virtuoso 6.1.5, though 6.1.8 has been verified functional. Design work was predominantly done in the library `optical_rx`, with the additional signoff elements (fill, I/O, seal ring) and reticle-level layout handled in `TAPEOUT_MAY2020_65LP`. All paths can be found in

`/tools/projects/lydiale/scum_v3/cds.lib`

but have been reiterated in Table C.7 for your convenience.

Virtuoso 6.1.7 encounters issues with the streamin/streamout process for TSMC65LP. Do not use it for any streamin/streamout with TSMC65LP on the BWRC infrastructure!

Library	Path
<code>optical_rx</code>	<code>/tools/B/lydiale/tsmc65lp/cadence/optical_rx</code>
<code>SPAD</code>	<code>/tools/B/lydiale/tsmc65lp/cadence/SPAD</code>
<code>TAPEOUT_MAY2020_65LP</code>	<code>/tools/B/lydiale/tsmc65lp/cadence/TAPEOUT_MAY2020_65LP</code>
<code>SPAD_dec17 (TSMC28)</code>	<code>/tools/projects/luya/proj/TSMC28/SPAD_dec17</code>
<code>BIOTHZ_dec17_top (TSMC28)</code>	<code>/tools/projects/ameri/proj/biothz/virtuoso/BIOTHZ_dec17_top</code>

Table C.7: Location of libraries in Cadence Virtuoso

Appendix D

BAG 2.0 Scripts

All BAG design work for this dissertation was done with the following:

- span_ion: https://github.com/PisterLab/span_ion
- bag2_analog: https://github.com/PisterLab/bag2_analog
- bag2_digital: https://github.com/PisterLab/bag2_digital

D.1 General Utilities

BAG2 as it's delivered does not come with a design equivalent of `gen_cell.py`. To fix this, we wrote `dsn_cell.py`, which can be found at https://github.com/PisterLab/bag2_analog/blob/master/scripts_dsn/dsn_cell.py. It should be copied into the same location as `BAG_framework/run_scripts/gen_cell.py` and can be invoked in a similar fashion for designing cells (see the README at https://github.com/PisterLab/bag2_xt018_workspace/tree/span-ion).

D.2 Scripts

Design scripts can be found in the associated repositories under “scripts_dsn”

Name	Description
attenuator	
attenuator2	
attenuator3	The resistive elements and mux of the attenuator used in this dissertation.
comparator*	Various fully differential and single ended stages; chains; and common mode feedback amplifiers.
dac_offset	The DAC used in for offset cancellation in the autozeroing scheme.
delay_sk_ord2	2nd order low pass filter Sallen-Key.
delay_tt1_ord2	2nd order Tow-Thomas type 1 filter.
delay_tt2_ord2	2nd order Tow-Thomas type 2 filter.
one_shot_nand	NAND-based one shot pulse generator.
one_shot_nand_tmr	NAND-based one shot pulse generator; with triple modular redundancy.
peak_detector_basic1/2/3	Peak detectors in various forms with and without current limiting resistors.
preamp	The preamplifier used in the full/main signal chain. Figure 3.21
scanchain	Scan chain with triple modular redundancy and DICE latches.
scanchain_cell	A single scan cell.
scanchain_cell_tmr	A scan cell; with triple modular redundancy.
scanchain_single	A scan chain without any triple modular redundancy.
voter_3	A 3x1 voter (NAND logic).
voter_3x3	A 3x3 voter (NAND logic).
watchdog	The watchdog described in Figure 2.11.

Table D.1: SPAN-Ion specific schematic generators.

Name	Description
amp_diff_mirr	5T differential pair with current mirror load; biased with a current source
amp_diff_mirr_bias	5T differential pair with current mirror load; biased with a gate voltage on the tail device
amp_folded_cascode	Folded cascode; biased with a gate voltage on the tail device
amp_gm_mirr	Gm amplifier; biased with a current source
amp_inv	Inverter
bandgap	Bandgap reference circuit. Note that this may need to replace the diode components depending on the PDK setup.
bandgap_startup	Startup circuit for the bandgap reference.
constant_gm	Constant gm circuit.
dac_r2r	R2R DAC.
dac_rladder	Resistive ladder DAC; constructed with a binary mux.
diffpair_n	Input NMOS differential pair.
diffpair_p	Input PMOS differential pair.
mirror_n	NMOS current mirror.
mirror_p	PMOS current mirror.
mux_bin	Binary mux; with single-ended select bit. Constructed as a tree of 1-bit muxes.
mux_bin_core	Binary mux; with differential select bits. Constructed as a tree of 1-bit muxes.
mux_bin_unit	1-bit binary mux.
mux_onehot_sel	One-hot mux; constructed as a bunch of switches with one ended shorted together.
nmos4_astack	
pmos4_astack	
r2r_core	The resistive elements of an R2R DAC.
regulator_ldo_series	Series low dropout regulator.
res_multistrip	Multistrip resistor.
res_trim_parallel	Trimming resistor where the resistor strips are in parallel and optionally switched in.
res_trim_series	Trimming resistor where the resistor strips are in series and optionally shorted out.
rladder_core	Resistive elements of a resistive ladder DAC.
switch	Do not use; Cadence does some specific name reservation.
switch_mos	MOS analog switch.

Table D.2: General analog schematic generators. There may be some overlap between other pre-existing BAG libraries due to accidental parallel or simultaneous creation.

Name	Description
clkgen_nonoverlap	Nonoverlapping clock generator
delay_line_inv_starved	Inverter chain of current-starved inverters.
flipflop_DICE	D flip flop constructed of two DICE latches.
flipflop_D_inv	D flip flop constructed with inverters
flipflop_D_nand	D flip flop constructed with NAND gates
inv	2-transistor inverter
inv_chain	Chain of 2-transistor inverters
inv_tristate	Tri-state inverter
latch_D	D latch
latch_DICE_clk	DICE latch with single-ended clock
latch_DIE_clk_sel	DICE latch with differential clock
latch_DICE_tgate	DICE latch which uses transmission gates and has a single-ended clock
latch_DICE_tgate_sel	DICE latch which uses transmission gates and has a differential clock
latch_SbRb	SbRb latch
level_shift_lo2hi	Level shifter intended for low-to-high transitions
nand	NAND gate
nmos4_stack	
nor	NOR gate
pmos4_stack	

Table D.3: General digital schematic generators. There may be some overlap between other pre-existing BAG libraries due to accidental parallel or simultaneous creation.

Appendix E

Chip V1 Documentation

E.1 Infrastructure

This chip's design used the alcatraz server available through the Berkeley Sensor and Actuator Center (BSAC). Information on accessing the server can be found at <https://bsac.berkeley.edu/software>. This requires a CalNet login and appropriate BSAC access.

All design was done using Cadence Virtuoso 6.1.8. All test code should be able to run locally using a distribution of Python 3.6.5 or higher. The test code likely works with lower versions of Python 3, but this hasn't been verified.

E.2 IO

Pin	Pad Location (Counter Clockwise)
SCAN_INb	W0
VSS	W1
VDDAON	W2
x	W3
VDDHV	W4
x	W5
VSS	W6
x	W7
VDDAON	W8
x	W9
VSS	S0
VDDTEST	S1
VDDHV	S2
VBG	S3
RST_PK	S4
VIN_PK	S5
VOUT_PK	S6
x	S7
VDDMAIN	S8
VSS	S9
VDAC_MAIN	S10
VSS	E0
VDDMAIN	E1
OUT_MAIN	E2
VREF_ATTEN	E3
IIN	E4
VREF_PREAMP	E5
OUT_SMALL	E6
VINP_LED_SMALL	E7
VINN_ZCD_SMALL	E8
VINP_ZCD_SMALL	E9
VDAC_SMALL	N0
VSS	N1
VDDSMALL	N2
x	N3
VSS	N4
VDDHV	N5
VSS	N6
SCAN_OUTb	N7
VDDAON	N8
SCAN_CLK	N9
SCAN_LOADb	N10

Table E.1: Chip V1 I/O and associated pad locations. Unconnected pads are marked as x.

Pin Name	Description	Type	Domain
IIN	Anode connection for the CFD. Bias point nominally the same as VREF_PREAMP	Input	MAIN
OUT_MAIN	Digital output of the main signal chain.	Output	MAIN
OUT_SMALL	Digital output of the small signal chain.	Output	SMALL
RST_PK	Reset for the test peak detector. Note that the peak detector is purely for testing purposes. Input voltage should not exceed 1.98V.	Input	TEST
SCAN_CLK	Scan chain clock. Used to clock data into the scan chain shift register.	Input	AON
SCAN_INb	Input data for the scan chain (inverted). Used in conjunction with other SCAN_* pins to set the chip into a known configuration.	Input	AON
SCAN_LOADb	Scan chain load (inverted). Used to load data from the shift register to the rest of the chip to assign the chip into a known configuration.	Input	AON
SCAN_OUTb	Output scan data. Used for verifying scan chain functionality.	Output	AON
VBG	Nominally temperature-invariant bandgap voltage from a test bandgap. This bandgap is used for the current references for the test structures. Nominally $\approx 1.25V$ +/- 6% error across -55C to 125C with a 3.3V supply voltage across corners.	Output	HV
VDAC_MAIN	N-input of the leading edge detector of the main signal chain.	I/O	MAIN
VDAC_SMALL	N-input of the leading edge detector of the small signal chain.	I/O	SMALL
VDDHV	"High voltage supply. Expected 3.3V but can operate up to 5.5V. Unknown minimum capacitance requirements but recommend significant external capacitance ($\approx 10nF$)."	Input	HV
VDDMAIN	"Internally regulated supply voltage for the larger signal chain. Nominally 1.8V max 1.98V. Can be externally overridden via pad. Simulation suggests 10nF external capacitance to keep supply bounce $< 5mV_{pkpk}$."	I/O	MAIN
VDDSMALL	"Internally regulated supply voltage for the pared-down signal chain (the one without signal shaping). Nominally 1.8V max 1.98V. Can be externally overridden via pad. Simulation suggests $< 10nF$ external capacitance to keep supply bounce $< 5mV_{pkpk}$."	I/O	SMALL
VDDTEST	"Externally supplied supply voltage for low-voltage test structures. Nominally 1.8V max 1.98V."	Input	TEST
VIN_PK	Input for the test peak detector. Note that the peak detector is purely for testing purposes. Input voltage should not exceed 1.98V.	Input	TEST
VINN_ZCD_SMALL	"Input to the N-input of the zero crossing detector on the small signal chain. On the main signal chain this is the (peak held) and attenuated input."	Input	SMALL
VINP_LED_SMALL	"Input to the P-input of the leading edge detector on the small signal chain. On the main signal chain this is the (peak held) input."	Input	SMALL
VINP_ZCD_SMALL	"Input of the P-input of the zero crossing detector on the small signal chain. On the main signal chain this is the delayed input."	Input	SMALL
VOUT_PK	Output from test peak detector. Note that this peak detector is purely for testing purposes. Comparison should be relative to VIN_PK. Testing should include low and high speed signals.	Output	TEST
VREF_ATTEN	Attenuator reference voltage.	Input	MAIN
VREF_PREAMP	"Reference voltage for preamp output and input voltage biasing. Controlled via scan identically to VDAC_MAIN though it can be externally overridden."	I/O	MAIN
VSS	Ground.		
x	Unconnected.		

Table E.2: Chip V1 I/O descriptions.

E.3 Scan Bits

Signal Name	Description	Bits (MSB:LSB)
preamp_res	Binary. Controls the feedback resistor value; which in turn controls DC gain and time constant. $R = (\text{code}+1)*1\text{k}\Omega$.	13.8
delay_res	Binary. Controls the resistor values for the delay filter where the resistor values increase linearly vs. code.	11.1
watchdog_res	Controls the time constant which is used to determine if the CFD is stuck high due to a single event effect. 0 = minimum time constant; 1 = maximum time constant.	30
attenuator_sel	Binary. Sets the attenuation of the signal relative to the DC reference voltage to be compared vs. the delayed; unattenuated signal. $V_{\text{out}} = V_{\text{in}} * (\text{code}+1)/8$	31.12.9
dac_sel	Binary. Selects between 256 contiguous sections of a 512-element resistor ladder DAC; starting at index 71. $V_{\text{out}} = \text{FSR} * (\text{code}+71)/512$	5.26.17.4.25.18.3.2
az_main_gain	Thermometer. Determines the magnitude of offset correction gain for the main amplifier in the autozeroing comparator. 000 corresponds to no offset correction. 100; 010; and 001 are functionally identical.	23.20.1
az_aux_gain	Thermometer. Determines the magnitude of offset correction gain for the auxiliary amplifier in the autozeroing comparator. 000 corresponds to no offset correction. 100; 010; and 001 are functionally identical.	22.21.0
oneshot_res	Binary. Controls the resistor (and ergo; the time constant) which is used in the output one-shot pulse generator/monostable multivibrator. 00 = minimum time constant; 11 = maximum time constant.	24.19
vref_preamp	Binary. DAC control for the preamplifier reference voltage. Identical to signal chain resistor ladder DAC described in dac_sel; $V_{\text{out}} = \text{FSR} * (\text{code}+71)/512$	29.14.7.28.15.6.27.16
vdd_aon	DAC control for the always-on regulator. Nominal voltage range is 1.8V-2.1V; it was set high to function across corners. Higher code = higher voltage. Note that this requires a 1.5nF cap with a 50k Ω parallel resistor.	38.37.36.35.34
vdd_signal	DAC control for the signal chain regulators. Nominal voltage range is 1.8V-2.1V; it was set high to function across corners. Higher code = higher voltage. Note that this requires an off-chip 10nF cap to maintain <5mV bounce across corners.	43.42.41.40.39
en_main	Enables the LDO associated with the full signal chain.	32
en_small	Enables the LDO associated with the pared-down signal chain.	33

Table E.3: Chip V1 scan bits.

E.4 Test Setup

All test code, raw data, and PCB files can be found at <https://github.com/PisterLab/span-ion>. For organization, the PCB files and code have been separated into submodules. The board design for this chip can be found at <https://github.com/PisterLab/span-ion-board/tree/master/span-ion-cfd>. Unfortunately the submodule did not preserve the commit history of the board design because of file size restrictions. The full board design commit history can be found in the history of the primary repository <https://github.com/PisterLab/span-ion>. The code base used to test the chip can be found at <https://github.com/PisterLab/span-ion-code/tree/master>, commit

08fb728893532a3aeb5ea94b921fe70aa7aeb14

made on May 22, 2022. The reason for the specificity of the commit is that the code base was updated to include more than a single channel per board (see Appendix F.4). The result is that the most recently-used code base cannot guarantee compatibility with this particular board due to slight changes in function arguments. Minimum working examples of all of the following can be found in `testing/run_me.py`. Teensy connections can be found at `testing/teensy/cfd.ino`.

E.4.1 Scan Chain Function

This determines if the data clocked into the chip's scan chain is equivalent to the data clocked out of the chip's scan chain. A minimum working example can be found in `testing` → `/run_me.py` on lines 18-40.

Required Connections:

- VDDHV + VSS
- VDDAON (optional, but should be verified as functional)
- SCAN_CLK, _INb, _LOADb, _OUTb

Equipment:

- Teensy 3.6
- DC power supply (Keysight E3631A)

Relevant Code:

- `scan.py`

E.4.2 Bandgap Voltage vs. Temperature

This is a characterization of the output voltage of the test bandgap reference structure against ambient temperature. A minimum working example can be found in `testing` → `run_me.py` on lines 70-88.

Required Connections:

- VDDHV + VSS
- VBG

Equipment:

- Teensy 3.6
- DC power supply (Keysight E3631A)
- temperature chamber (TestEquity Model 107)
- digital temperature sensor (TMP102, optional)

Relevant Code:

- bandgap.py
- temp_chamber.py

Procedural Notes:

- The temperature chamber's sweep must be started using the front panel of the chamber, and to our knowledge cannot be done programmatically.
- Communication with the temperature chamber is done via RS-232.

E.4.3 Peak Detector Static Error

This is a characterization of the static error associated with the peak detector for a slow input step, produced by a second Teensy 3.6 with its analogWrite. For each measurement, the voltage is reset and the output of the peak detector is also reset. A minimum working example can be found in `testing/run_me.py` on lines 111-134.

Required Connections:

- VDDHV + VSS
- VDDTEST
- VIN_PK + VOUT_PK + RST_PK

Equipment:

- Teensy 3.6 (2×)
- DC power supply (Keysight E3631A)

Relevant Code:

- teensy/cfd_aux

E.4.4 Voltage DAC Characterization

This characterizes the DAC's output voltage versus code by programming the scan chain setting of the relevant DAC, then measuring the output voltage. A minimum working example can be found in `testing/run_me.py` on lines 42-65. Code for calculating the full scale

range, DNL, INL, etc. is also included.

Required Connections:

- VDDHV + VSS
- VDDAON (optional, but should be verified as functional)
- SCAN_*
- VDAC_SMALL + VDDSMALL (if testing the no-shape/small chain's DAC)
- VDAC_MAIN + VDDMAIN (if testing the full/main signal chain's DAC)
- VREF_PREAMP + VDDMAIN (if testing the preamp DAC)

Equipment:

- Teensy 3.6 (2×)
- DC power supply (Keysight E3631A)

Relevant Code:

- dac.py
- spani_globals.py

E.4.5 No-Shape/Small Signal Chain

This code can be used to get the time difference of arrival between an input trigger and the output of the no-shape/small signal chain. A minimum working example can be found in `testing/run_me.py` on lines 187-237.

Required Connections:

- VDDHV + VSS
- VDDAON (optional, but should be verified as functional)
- SCAN_*
- OUT_SMALL
- VDAC_SMALL (optional)
- VDDSMALL (optional, but should be verified as functional)
- VINN_ZCD_SMALL + VINP/N_LED_SMALL

Equipment:

- Teensy 3.6 (2×)
- DC power supply (Keysight E3631A)
- pulse generator (DG535)
- attenuator (Kay Elemetrics 839)
- delay line (long coaxial cable)

- 1.8V latch
- level shifter (MAX13003EEUE, on the board)
- TDC (TDC7200PWR, on the board)

Relevant Code:

- `tdc.py`
- `testing.py/test_tdiff_small()`

Procedural Notes:

1. Connect the output of the DG535's T0 to branch three ways—the attenuator (ZCD_N), the delay (ZCD_P), and unaltered (LED_SMALL).
2. The Teensy START signal that goes to the TDC can be used as the trigger signal for the DG535, or the code can be modified to configure the DG535 in single shot mode with a programmatic trigger.
3. The output of the chip is at the core voltage, whereas the TDC communicates with 3.3V. For this, we use a latch and a level shifter (shown in Figure 3.6). Between each measurement, the latch is reset by the Teensy.
4. The example assumes communication with the DG535 via a Prologix ethernet connection. Code for GPIB and Prologix connections can be found in `testing/gpib.py`.
5. To find the IP address of the Prologix, use the free Netfinder application from <https://prologix.biz/>.

E.4.6 Full/Main Signal Chain

This code can be used to get the time difference of arrival between an input trigger and the output of the no-shape/small signal chain. A minimum working example can be found in `testing/run_me.py` on lines 337-394.

Required Connections:

- VDDHV + VSS
- VDDAON (optional, but should be verified as functional)
- SCAN_*
- OUT_MAIN
- VDAC_MAIN + VREF_PREAMP (optional)
- VDDMAIN (optional, but should be verified as functional)
- VREF_ATTEN
- IIN

Equipment:

- Teensy 3.6 (2×)

- DC power supply (Keysight E3631A)
- pulse generator (DG535) connected with a Prologix ethernet connection
- 1.8V latch
- level shifter (MAX13003EEUE, on the board)
- TDC (TDC7200PWR, on the board)

Relevant Code:

- tdc.py
- testing.py/test_tdiff_main()

Procedural Notes:

1. The Teensy START signal that goes to the TDC can be used as the trigger signal for the DG535, or the code can be modified to configure the DG535 in single shot mode with a programmatic trigger.
2. The output of the chip is at the core voltage, whereas the TDC communicates with 3.3V. For this, we use a latch and a level shifter (shown in Figure 3.6). Between each measurement, the latch is reset by the Teensy.
3. To find the IP address of the Prologix, use the free Netfinder application from <https://prologix.biz/>.

E.5 Cadence Locations

This section assumes you have access to the TSMC180nm PDK available on the alcatraz server. Table E.4 has the location of the top level libraries used for Chip V1 design and tapeout. The library used for reticle assembly is tapeout_2021.07.21. The library used for reticle streaming in and out is tapeout_2021.07.21_streamin. The library used for Chip V1 is TAPEOUT_20210721_SPANI_techFixed, placed within a 2.5mm×2.5mm seal ring. Many of the BAG-generated components were placed into separate libraries, all starting with “ZZ.KEEP”. The cds.lib containing all the references to the various generated libraries can be found at

```
/home/eecs/lydialeee/tsmc180_virtuoso6/bag_workspace_tsmc180/cds.lib
```

Library	Path
tapeout_2021.07.21	/home/local/git/tsmc18/tapeout_2021.07.21
TAPEOUT_20210721_SPANI_techFixed	/home/eecs/lydialeee/tsmc180_virtuoso6/TAPEOUT_20210721_SPANI_techFixed
tapeout_2021.07.21_streamin	/home/eecs/lydialeee/tsmc180_virtuoso6/tapeout_2021.07.21_streamin

Table E.4: Top-level Cadence libraries for Chip V1.

E.6 Known Idiosyncrasies

The list here includes things only seen when testing Chip V1. All weirdness seen in Appendix Section F.6 could conceivably apply here as well, but weren't observed at the time of testing.

- Missing references in libraries: If you open a schematic—particularly a test bench—and find that a block appears to be missing, check the library in question and ensure that the case of the instance name matches. During the tapeout process, it was discovered that the LVS rules did not distinguish between cell names with different cases. For example, if there were two different cells instantiated in a single block, one of type “cell_test” and the other of type “CELL_TEST”, LVS would construct the schematic netlist using the subcircuit of only one, rather than distinguishing between the two. As such, some cells within libraries needed to be renamed to not have any capital letters in the cell name, causing their instantiations in subblock test benches to be replaced by blinking empty boxes with a warning about missing references.

Appendix F

Chip V2 Documentation

F.1 Infrastructure

This chip's design used the alcatraz server available through the Berkeley Sensor and Actuator Center (BSAC). Information on accessing the server can be found at <https://bsac.berkeley.edu/software>. This requires a CalNet login and appropriate BSAC access.

All design was done using Cadence Virtuoso 6.1.8. All test code should be able to run locally using a distribution of Python 3.6.5 or higher. The test code likely works with lower versions of Python 3, but this hasn't been verified.

Issue tracking for individual boards and chips, as well as chip locations for future return to TSMC can be found at <https://docs.google.com/spreadsheets/d/1RqJV0E82mo0agj-ajrNGloXr8MB5m/edit?usp=sharing>. You must have a berkeley.edu email to view this.

F.2 IO

Pin	Pad Location (Counter Clockwise)
SCAN_OUTb	E0
VSS	E1
OUT_CFD_SMALL_LV	E2
OUT_LED_SMALL_LV	E3
OUT_SMALL_HV	E4
OUT_SMALL_LV	E5
VINP_ZCD_SMALL	E6
VINN_ZCD_SMALL	E7
VINP_LED_SMALL	E8
VDDSMALL	E9
OUT_CFD_MAIN_LV	E10
OUT_LED_MAIN_LV	E11
VDDMAIN	E12
VSS	E13
VSS	N1
OUT_MAIN_LV	N3
OUT_MAIN_HV	N4
IIN	N5
VSS	N6
VREF_ATTEN_MAIN	N7
VDDAON	N10
VDDHV	N11
VSS	N14
SCAN_INb	S0
SCAN_LOADb	S1
VDDHV	S2
VDDAON	S5
VSS	S6
VBG	S7
VOUT_PK	S8
VDDTEST	S9
VIN_PK	S10
VIN_DELAY	S1
VOUT_DELAY	S12
VRST_PK_TEST	S13
SCAN_CLK	S14
VREF_PREAMP	W0
VDAC_MAIN	W1
VDAC_SMALL	W5
VDDAON	W7
VSS	W8
VDDHV	W9
RST_PK	W10
VDDAON	W13
VSS	W14

Table F.1: Chip V2 I/O and associated pad locations. Unconnected pads are not included.

Pin Name	Description	Type	Domain
IIN	Anode connection for the CFD. Bias point nominally the same as VREF_PREAMP	Input	MAIN
OUT_CFD_MAIN_LV	Digital output of the CFD branch of the main signal chain.	Output	MAIN
OUT_CFD_SMALL_LV	Digital output of the CFD branch of the small signal chain.	Output	SMALL
OUT_LED_MAIN_LV	Digital output of the LED branch of the main signal chain.	Output	MAIN
OUT_LED_SMALL_LV	Digital output of the LED branch of the small signal chain.	Output	SMALL
OUT_MAIN_HV	Digital output of the main signal chain level shifted to VDDHV.	Output	HV
OUT_MAIN_LV	Digital output of the main signal chain.	Output	MAIN
OUT_SMALL_HV	Digital output of the small signal chain level shifted to VDDHV.	Output	HV
OUT_SMALL_LV	Digital output of the small signal chain.	Output	SMALL
RST_PK	Reset for the test peak detector. Note that the peak detector is purely for testing purposes. Input voltage should not exceed 1.98V.	Input	MAIN
SCAN_CLK	Scan chain clock. Used to clock data into the scan chain shift register.	Input	AON
SCAN_INb	Input data for the scan chain (inverted). Used in conjunction with other SCAN_* pins to set the chip into a known configuration.	Input	AON
SCAN_LOADb	Scan chain load (inverted). Used to load data from the shift register to the rest of the chip to assign the chip into a known configuration.	Input	AON
SCAN_OUTb	Output scan data. Used for verifying scan chain functionality.	Output	AON
VDAC_MAIN	N-input of the leading edge detector of the main signal chain. Is connected to an on-chip resistive DAC. Can be externally overridden.	I/O	MAIN
VDAC_SMALL	N-input of the leading edge detector of the small signal chain. Is connected to an on-chip resistive DAC. Can be externally overridden.	I/O	SMALL
VDDAON	Internally regulated supply voltage for the scan chain. Nominally 1.8V max 1.98V. Can be externally overridden via pad.	I/O	AON
VDDHV	High voltage supply. Expected 3.3V but can operate up to 5.5V. Unknown minimum capacitance requirements but recommend significant external capacitance (10nF).	Input	HV
VDDMAIN	Internally regulated supply voltage for the larger signal chain. Nominally 1.8V max 1.98V. Can be externally overridden via pad. Simulation suggests 10nF external capacitance to keep supply bounce <5mVpkpk.	I/O	MAIN
VDDSMALL	Internally regulated supply voltage for the pared-down signal chain (the one without signal shaping). Nominally 1.8V max 1.98V. Can be externally overridden via pad. Simulation suggests <10nF external capacitance to keep supply bounce <5mVpkpk.	I/O	SMALL
VDDTEST	Externally supplied supply voltage for low-voltage test structures. Nominally 1.8V max 1.98V.	Input	TEST
VIN_DELAY	Input for the test delay filter. Note that this delay filter is purely for testing purposes. Input voltage should not exceed 1.98V.	Input	TEST
VIN_PK	Input for the test peak detector. Note that the peak detector is purely for testing purposes. Input voltage should not exceed 1.98V.	Input	TEST
VINN_ZCD_SMALL	Input to the N-input of the zero crossing detector on the small signal chain. On the main signal chain this is the (peak held) and attenuated input.	Input	SMALL
VINP_LED_SMALL	Input to the P-input of the leading edge detector on the small signal chain. On the main signal chain this is the (peak held) input.	Input	SMALL
VINP_ZCD_SMALL	Input to the P-input of the zero crossing detector on the small signal chain. On the main signal chain this is the delayed input.	Input	SMALL
VOUT_DELAY	Output from the test delay filter. Note that this delay filter is purely for testing purposes. Given a lack of likely inadequate buffering on the output recommend testing with relatively slow signals using an oscilloscope.	Output	TEST
VOUT_PK	Output from test peak detector. Note that this peak detector is purely for testing purposes. Comparison should be relative to VIN_PK. Testing should include low and high speed signals.	Output	TEST
VREF_ATTEN_MAIN	Attenuator reference voltage. Recommend placing this at or below the preamp reference voltage.	Input	MAIN
VREF_PREAMP	Reference voltage for preamp output and input voltage biasing. Controlled via scan nominally 0.5V +/- <<double-check>>V though it can be externally overridden.	I/O	MAIN
VRST_PK_TEST	Voltage that the output of the test peak detector is pulled to when the peak detector's reset is held high.	Input	TEST
VSS	Ground.	Input	

Table F.2: Chip V2 I/O descriptions.

F.3 Scan Bits

Signal Name	Description	Bits (MSB:LSB)
preamp_res	Binary. Controls the feedback resistor value; which in turn controls DC gain and time constant. $R = (\text{code}+1)*5\text{k}\Omega$.	102. 101
delay_res	Binary. Controls the resistor values for the delay filter where the resistor values increase linearly vs. code.	71. 48
watchdog_res	Controls the time constant which is used to determine if the CFD is stuck high due to a single event effect. 0 = minimum time constant; 1 = maximum time constant.	144
en_stuck	Active high enable to use the stuck signal. The stuck circuitry is always sensing; but setting this low makes it so the rest of the hardware never uses it.	120
attenuator_sel	Binary. Sets the attenuation of the signal relative to the DC reference voltage to be compared vs. the delayed; unattenuated signal. $V_{\text{out}} = V_{\text{in}} * (\text{code}+1)/8$	60. 59. 143
dac_sel	Binary. Selects between 256 contiguous sections of a 512-element resistor ladder DAC; starting at index 71. $V_{\text{out}} = \text{FSR} * (\text{code}+71)/512$	187. 184. 145. 142. 103. 100. 61. 58
oneshot_res	Binary. Controls the resistor (and ergo; the time constant) which is used in the output one-shot pulse generator/monostable multivibrator. 00 = minimum time constant; 11 = maximum time constant.	185. 186
vref_preamp	Binary. DAC control for the preamplifier reference voltage. Identical to signal chain resistor ladder DAC described in dac_sel; $V_{\text{out}} = \text{FSR} * (\text{code}+71)/512$	188. 183. 146. 141. 104. 99. 62. 57
en_pullup_p_led	Enable for the p-side pull-up for LED comparator stages. Each bit corresponds to a single stage. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	94. 53. 179. 138. 97. 56. 182
en_pullup_n_led	Enable for the n-side pull-up for LED comparator stages. Each bit corresponds to a single stage. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	109. 66. 192. 149. 106. 63. 189
en_pullup_p_cfd	Enable for the p-side pull-up for CFD primary comparator stages. Each bit corresponds to a single stage. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	74. 36. 28. 20. 16. 7. 4
en_pullup_n_cfd	Enable for the n-side pull-up for CFD primary comparator stages. Each bit corresponds to a single stage. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	87. 35. 29. 19. 17. 6. 5
en_pulldown_p_led	Enable for the p-side pull-down for LED comparator stages. Each bit corresponds to a single stage. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	155. 82. 124. 39. 92. 164. 135
en_pulldown_n_led	Enable for the n-side pull-down for LED comparator stages. Each bit corresponds to a single stage. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	174. 121. 133. 80. 111. 165. 152
en_pulldown_p_cfd	Enable for the p-side pull-down for CFD primary comparator stages. Each bit corresponds to a single stage. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	41. 126. 84. 157. 46. 159. 117
en_pulldown_n_cfd	Enable for the n-side pull-down for CFD primary comparator stages. Each bit corresponds to a single stage. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	161. 89. 42. 127. 85. 158. 75
ctrl_pullup_led	Code setting for the pull-up DAC on the LED comparator. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	177. 194. 52. 67. 136. 151. 178. 193. 95. 108. 137. 150. 54. 65. 96. 107. 180. 191. 55. 64. 139. 148. 181. 190. 98. 105. 140. 147
ctrl_pulldown_led	Code setting for the pull-down DAC on the LED comparator. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	90. 113. 125. 132. 162. 167. 40. 79. 49. 70. 91. 112. 154. 175. 163. 166. 81. 122. 50. 69. 123. 134. 153. 176. 51. 68. 93. 110
ctrl_pullup_cfd	Code setting for the pull-up DAC on the CFD comparator. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	45. 32. 33. 34. 37. 38. 31. 30. 27. 26. 25. 18. 21. 22. 23. 24. 15. 14. 13. 12. 8. 9. 10. 11. 3. 2. 1. 0
ctrl_pulldown_cfd	Code setting for the pull-down DAC on the CFD comparator. Note that this is buffered to the correct voltage domain; since it is applied to both the small and main signal chains.	168. 78. 173. 83. 114. 131. 72. 156. 77. 119. 169. 47. 130. 172. 115. 160. 118. 73. 76. 88. 171. 170. 129. 43. 86. 116. 44. 128
vdd_aon	DAC control for the always-on regulator. Nominal voltage range is 1.8V-2.1V; it was set high to function across corners. Higher code = higher voltage. Previously this was used with a 1.5nF cap and a 50k Ω parallel resistor.	206. 195. 205. 196. 204
vdd_signal	DAC control for the signal chain regulators. Nominal voltage range is 1.8V-2.1V; it was set high to function across corners. Higher code = higher voltage. Note that this requires an off-chip 10nF cap to maintain <5mV bounce across corners.	197. 203. 198. 202. 199
en_main	Enables the LDO associated with the full signal chain.	201
en_small	Enables the LDO associated with the pared-down signal chain.	200

Table F.3: Chip V2 scan bits.

F.4 Test Setup

All test code, raw data, and PCB files can be found at <https://github.com/PisterLab/span-ion>. This is the same repository used in Appendix E.4. For organization, the PCB

files and code have been separated into submodules. The board design for this chip can be found at

<https://github.com/PisterLab/span-ion-board/tree/master/span-ion-cfd-respin-noMuxDemux>

Unfortunately the submodule did not preserve the commit history of the board design because of file size restrictions. The full board design commit history can be found in the history of the primary repository <https://github.com/PisterLab/span-ion>. The code base used to test the chip can be found at <https://github.com/PisterLab/span-ion-code/tree/master>, where the most recent commit is

785e0bcb16fc820ed747dc578fe0d5d537667b8c

made on June 23, 2022. This test setup differs from that of Chip V1 in Appendix E.4 in several significant ways:

- Each board contains two chips. Between the two chips, the board can be configured via a surface-mount 0Ω resistor in a specific direction (see Figure F.1) to either test an individual chip—called “SINGLE”—or to test two chips as they relate to one another—called “DUAL”. There are an accordingly increased number of TDCs.
- The microcontroller for testing was changed from the Teensy 3.6 to the Arduino Due.
- The latch after the main output is no longer placed in the signal chain leading up to the TDC, and instead the 3.3V level shifted output is fed to the input of the TDC.




Figure F.1: The L-resistor footprint. Only one direction (N/S or E/W) is populated.

Minimum working examples of all of the following can be found in `testing/run_me.py`. Arduino Due connections can be found at `testing/arduino-due/cfd.ino`. The code setup for Chip V2 is nearly identical to that of Chip V1, with minor modifications made to adjust to having multiple channels on a single board. So far, only code testing the scan chain, bandgap test code, peak detector characterization, and voltage DAC characterization code have been verified as functional with the new boards. All of these have minimum working examples in `testing/run_me.py`, and the connections are largely the same as those found in Appendix E.4.

F.5 Cadence Locations

This section assumes you have access to the TSMC180nm PDK available on the alcatraz server. Table F.4 has the location of the top level libraries used for Chip V2 design and tapeout. The library used for seal ring assembly with the full reticle is TAPEOUT_20221116_SPANI_sealring. The library used for my chip streaming in and out is TAPEOUT_20221116_SPANI_streamin3_finalrun. The library used for Chip V2 design is TAPEOUT_20221116_SPANI. Many of the BAG-generated components were placed into separate libraries, all starting with “ZZ_KEEP” or “ZZZ_KEEP”. The cds.lib containing all the references to the various generated libraries can be found at

`/home/eecs/lydialeee/tsmc180_virtuoso6/bag_workspace_tsmc180/cds.lib`

Library	Path
TAPEOUT_20221116_SPANI	<code>/home/eecs/lydialeee/tsmc180_virtuoso6/TAPEOUT_20221116_SPANI</code>
TAPEOUT_20221116_SPANI_sealring	<code>/home/eecs/lydialeee/tsmc180_virtuoso6/TAPEOUT_20221116_SPANI_sealring</code>
TAPEOUT_20221116_SPANI_streamin3_finalrun	<code>/home/eecs/lydialeee/tsmc180_virtuoso6/TAPEOUT_20221116_SPANI_streamin3_finalrun</code>

Table F.4: Top-level Cadence libraries for Chip V2.

F.6 Known Idiosyncrasies

The list here includes things only seen when testing Chip V2. All weirdness seen in Appendix Section E.6 still applies.

- Prologix not appearing on netfinder: Redownload the software from `prologix.biz`. I don’t know why this works every time with a fresh download.
- Prologix IP address issues: Sometimes you have to set the Prologix IP address manually in order for your machine to connect to it. Many thanks to Robert Abiad for helping figure out a fix. What worked for me on a machine running Windows 10:

1. Control Panel → Network Connections → Ethernet → Properties → Internet Protocol Version 4.
2. Manually set the IP address (192.168.1.100/255.255.255.0/192.168.1.1 worked for me).
3. Manually set the DNS server address preference (192.168.1.1 worked for me with the aforementioned IP address).
4. Close and restart the netfinder software.

To find the appropriate IP address, subnet mask, etc.,

1. In the terminal, type

```
ipconfig
```

and check under Ethernet what the IPv4 address, subnet mask, and default gateway are.

2. Use the subnet mask and default gateway when setting the IP address for the Prologix. For the Prologix, just make sure the first 3 numbers are the same; we don't know what the last number means.

Appendix G

Miscellaneous

This appendix is devoted to the tribal knowledge I've accumulated in my time as a graduate student. Ideally a lot of these tables would be in a place that's editable by the people currently in the group to keep things up-to-date, but based on experience that's asking for the file to vanish. Instead, I'm going to put this here with the asterisk that the information here is accurate only as of the time this dissertation was written.

Subject	Location	Additional Notes
XFAB on BSAC infrastructure	https://bsac.berkeley.edu/software	Must be a BSAC member to view
TSMC on BSAC infrastructure		
Swarm pick-and-place surface mount soldering in Swarm	https://bamlab.berkeley.edu/wiki/swarm_lab	
Running BAG scripts	https://github.com/PisterLab/bag2_xt018_workspace/tree/span-ion	Must fall under the Berkeley NDA for XFAB to view
Overleaf with raw data for this dissertation	https://www.overleaf.com/read/dkvwcsfrmtqq	
dissertation talk	https://drive.google.com/drive/folders/18xUbt051_JMmkbrVbaGyksrns1615bzQ?usp=sharing	

Table G.1: Locations of useful tribal knowledge.

- TestEquity Model 107 checksum failures: You may encounter errors involving a checksum when attempting to communicate with this machine via RS-232. This is likely because the port on the back has 3 of 9 pins remaining, all of which are barely hanging on.
- PCB fab: Advanced Circuits is based out of the US and expensive, though their customer service has historically been fantastic. JLCPCB is based out of China and extremely cheap and fast.

- Chip packaging: QPTechnologies, formerly QuikPak, has done our packaging and chip-to-board bonding.
 - Email the point of contact (currently Marthus Victoria) directly. They have historically been much faster to respond this way than if contacted through the form on their website.
 - Some of their packages or epoxies cannot withstand higher surface mount soldering temperatures. Without precise knowledge of what temperatures are acceptable (we know that 255°C is a no-go), we hold the package in place with UV cure epoxy, then apply silver epoxy (cure temp $\approx 100^{\circ}\text{C}$) to form electrical connections between the package and pads.
 - Other groups have reported issues with packaging integrity. Specifically, two different packages of the same chip (switched because of an epoxy shortage) have vastly different power consumption and measurable impedance differences between pins.
- PCB assembly: Digicom, based out of Oakland.
 - Mo Ohady has been our point of contact.
 - As of 2023, the highest temperature they report using for surface mount soldering is 255°C .