HIGH FIDELITY ANALOG-TO-DIGITAL CONVERSION FOR SPACEBORNE APPLICATIONS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Charles Chang-I Wang
September 2009

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Umran S. Inan) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Ivan R. Linscott)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Boris Murmann)

Approved for the University Committee on Graduate Studies.

Abstract

Recent interest in wave-particle interactions in the Earth's Van Allen belts has spurred the construction of a new generation of plasma wave receivers, scientific instruments that measure electromagnetic signals while aboard a satellite flying through the upper atmosphere. Such receivers have stringent analog-to-digital conversion requirements, as they must simultaneously capture multiple signals spread over a broad frequency range (from 100 Hz to 1 MHz) and spanning a wide dynamic range (90 dB) while consuming minimal power. In addition, they must maintain this performance as they fly through the damaging radiation environment of the Earth's radiation belts.

This dissertation describes the specification, design, implementation, and testing of the SVADC-1, a radiation-hard, 12-bit, pipeline analog-to-digital converter that meets these requirements. A consideration of the spectrographic nature of plasma wave analysis, coupled with the common occurrence of simultaneous strong and weak phenomena, identifies the spurious-free dynamic range (SFDR) as a key metric for plasma wave receivers. A subsequent investigation of quantization shows that, to achieve the 90-dB SFDR required (assuming 100-Hz/bin spectral resolution), a conversion of just 12 bits is sufficient. In implementing such conversions, though, traditional pipeline converters suffer from quantization nonuniformities that reduce their SFDR. These nonuniformities are primarily due to the mismatch between the analog stages of the pipeline and the corresponding digital reconstruction. Hence this dissertation introduces a novel self-calibration technique based on DAC differencing that corrects for this mismatch in the more malleable digital domain. As a result, after self-calibration the SVADC-1 achieves a wideband peak SFDR of $\geq$90.9 dB while sampling at 5 MS/s and consuming just 48.8 mW.

In addition, the SVADC-1 is radiation-hard. Traditionally, radiation tolerant electronics rely upon specialized manufacturing processes to guarantee radiation hardness. However, these specialty processes are often expensive and hard to obtain. In constrast, the SVADC-1 is fabricated in a commercial, 0.25-μm CMOS manufacturing process, and employs radiation-hardness by design—including techniques such as enclosed terminal layouts for selected transistors, self-resetting architectures, selective analog overdesign, and the use of guard rings—to compensate for radiation-induced degradation and upsets. In total-dose radiation testing of the SVADC-1 by 50-MeV protons, it maintains a performance of $\geq$90.1-dB peak SFDR while sampling at 5 MS/s and consuming $\leq$60.2 mW up to a total dose of 1 Mrad(Si), experiencing a slight decrease to $\geq$88.2-dB peak SFDR and $\leq$60.5 mW up to 2 Mrad(Si) (the highest dose tested). And in single-event radiation testing of the SVADC-1 by 25-MeV/nucleon heavy ions, it displays no latchup through an LET of 63 MeV-cm^2/mg (the highest tested LET) at elevated supply (2.7 V) and temperature (131°C).

Acknowledgments

There's only one name on the front of this thing, but there's a whole cavalcade of names that helped put it there.

First, I would like to thank my advisors, Professor Umran S. Inan and Dr. Ivan R. Linscott. I am grateful for the rare opportunity of building a project from the ground up—from determining the specifications, to designing and testing the chip, to eventually seeing it incorporated into an actual satellite instrument—and am appreciative of their abiding commitment to research in these contexts. I don't think it a stretch to say that none of us came into this endeavor with much experience in analog ASIC design. But I've nonetheless learned a lot throughout my time here and, as I joked with Professor Inan in my last meeting with him, some of it even technical. That's certainly true, and—in my mind, at least—a unique boon.

I would like to thank Professor Boris Murmann, who was invaluable in answering many practical questions about ADC design and who graciously agreed to be the third reader of my dissertation. And I would like to thank Professor John T. Gill III, not only my oral defense committee chairman, but also my advisor during my undergraduate studies at Stanford. Asking Professor Gill questions has been one of the most enjoyable parts of my time at Stanford, partially because he's always willing to tackle them, but mostly because I invariably end up learning about whatever *he's* thinking about at the moment, which is usually orthogonal but much more interesting.

I would like to acknowledge a variety of people at National Semiconductor, which generously fabricated the SVADC-1 chip described in this dissertation. I'd like to acknowledge the aegis of Dr. Ahmad Bahai, and also Bijoy Chatterjee. In addition, during my graduate years at Stanford I functioned as the technical point of contact

with National Semiconductor for the BiCMOS manufacturing processes, and so would like to acknowledge my counterparts on the National side, who have been Michael Schwartz, Saeed Khosrowbeygi, Babak Matinpour, and Ali Kiaei. I would also like to thank Kirby Kruckmeyer and Tom Santiago, who conducted the 25-MeV heavy-ion radiation testing described in this dissertation. And finally, I would like to thank the late Michael Maher. The literature on radiation effects is vast and rather intimidating but, early on, I was fortunate to have Mike Maher help me sort through it, pointing out what I should (and, as importantly, should not) be worrying about. I am grateful he showed an interest—even going so far as to schedule a regular meeting on Friday mornings—and hope that my understanding lives up to his tutelage.

All the radiation testing experiments performed for this dissertation (except the aforementioned) were done in collaboration with The Aerospace Corporation. I would thus like to thank Dr. James L. Roeder, who personally coordinated much of this testing, and Dr. J. Bernard Blake, who provided the opportunity for the proton testing at LBNL. I would also like to thank a variety of folks at The Aerospace Corporation upon whose experience and advice I relied heavily in designing, executing, and interpreting the testing, including Everett King and Steve Witczak, the total-dose guys; Rocky Koga, Jeff George, and Van Tran, the single-event guys; Steve LaLumondiere, the laser guy; and Mark Zakrzewski, Mark Turpin, and Stephanie Brown, staff engineers. And honorable mention to Mike Johnson of LBNL. Both down in El Segundo and up at LBNL, I have found all these folks both exceptionally helpful and, much to their credit, enormously patient.

In my final years here, I was funded by the WIPER satellite project, and I would like to acknowledge Jack Doolittle, Clem Tillier, and Bob Bumala, who afforded me the chance not only to see how a satellite instrument actually gets built, but also to take part and get my hands dirty.

I would like to acknowledge the Stanford VLF group in general and some folks in particular. First, the silicon subdivision: Ben Mossawir (whose erudition and attention to detail I have benefited from in many a discussion), Kelin Lee, and Sarah Harriman. Given the aggressive tapeout schedules pursued by the VLF group,[1] we

[1] My personal tally was 4 designs taped out over 2 years.

all banded together to help draw the outgoing chips, including the SVADC-1 of this dissertation. I'd also like to thank the basement dwellers: Jeff Chang, Joe Payne, Robb Moore, Ben Cotts, and Ryan Said, who were always willing to lend a hand in the laboratory, or just to sit in the office and discuss everything from research[2] to more important things, like how to properly throw a curveball.[3] And I'd like to thank the VLF staff: Shaolan Min, Helen Niu, and Dan Musetescu, who keep this whole sprawling operation on the rails.

I've met a lot of people here at Stanford who've enriched my time here. The Professor Wooley group, which I leaned on heavily for advice on ASIC design in general and ADC design in particular: Sotirios Limotyrakis (the only other pipeline ADC designer on campus when I started), KiYoung Nam, Scott Kulchycki, and David Barkin (who I peppered with questions about how to set up and maintain the Cadence tools). The pMRI folks down the hall: Ross Venook, Nate Matter, Steve Conolly, and Grieg Scott, who always offered a different (and valuable) perspective on the graduate experience. Undergraduate classmates: Alex Tung, Mike Weimer, David Zhu, Howard Shen, Eric Barrett, and Doug Chung, who inspire me to do better. And the Packard folks: Keith Gaul and Pat Burke, who lent me equipment and issued me operating systems, respectively, in more than one pinch.

Finally, I would like to thank my family, my father Niantsu, my mother Alice, and my sister Shirley. Over my graduate studies, I have come to value their advice—both nontechnical and otherwise—more and more, perhaps because I realize that they are among the few that can be said to truly have my best interests at heart, unburdened by their own. I said it at my oral defense, but it bears repeating here: if ever I have acted with dignity and aplomb, it is because they have reminded me—and I often do forget and they often do remind me—that that's the way I should be.

Charles Chang-I Wang
Stanford University
9 September 2009

[2]There is at least occasional utility in explaining something to someone who has no idea what you're talking about, as I have learned from experiences on both sides of the conversation.

[3]Literally, as well.

This work was sponsored by NASA Grant #NAG5-10822 (New Technology Initiatives Program) under The Aerospace Corporation subcontract #4600001624-9, and by the United States Air Force Research Laboratory (ARFL) under contract #FA8178-05-C-0027-20. Additional support was provided by The Aerospace Corporation for radiation testing by 50-MeV protons at Lawrence Berkeley National Laboratory. All semiconductor fabrication was generously provided by National Semiconductor Corporation under a collaboration agreement between National Semiconductor Corporation and Stanford University.

Contents

List of Tables

List of Figures

Chapter 1

Plasma Wave Receivers

Ground-based observations of naturally occurring electromagnetic waves in the Earth's ionosphere date back to at least the late 1800s. The first published report is believed to be by *Preece* [1894] and described whistlers heard on long telegraph lines during a display of the aurora borealis [*Helliwell*, 1965, p. 11]. Further studies of these whistlers and other phenomena, coupled with other investigations of the upper atmosphere, led to discoveries of extensive magneto-ionic structure at great distances from the Earth. Indeed, by the end of the International Geophysical Year (1 July 1957 to 31 December 1958), a theory for whistlers—based on lightning-injected energy echoing back and forth roughly along the Earth's magnetic field lines and traversing a dispersive ionized medium—was well accepted [*Shawhan*, 1979, p. 215]. However, the mechanisms of many other phenomena were not as well understood.

The introduction of sounding rockets and scientific satellites in the late 1950s enabled more direct exploration of the space surrounding the Earth. The first satellite observations of whistlers were aboard Vanguard III in 1959 [*Helliwell*, 1965, p. 134]. While mostly consistent with the understood theory of the time, these observations also displayed unpredicted properties, spurning additional research. Many more instruments have flown since, contributing insights into the structure of the Earth's magnetosphere[1] and beyond. In addition to elucidating ground-based observations,

[1] That is, the region about the Earth predominantly organized by its magnetic field. A more extensive description of the magnetosphere is given in Section 2.1.1.

these experiments also recorded many more phenomena not seen on the ground.

It is now known that, above altitudes of ∼100 km, the Earth is surrounded by a plasma (that is, an ionized gas) of varying density, energy, and composition. This plasma is immersed in a magnetic field of primarily terrestrial origin. The propagation of electromagnetic waves in the upper atmosphere, then, is greatly modified by the presence of this magnetized plasma: the electric and magnetic fields of these waves drive oscillatory motion of the free ions and electrons, while the motion of these charged particles constitutes a time-varying current that then modifies the wave fields [*Inan and Inan*, 2000, Chap. 6]. The resulting electromagnetic waves are called plasma waves, and the Earth's magnetosphere and ionosphere are host to a wide variety of these phenomena [*Shawhan*, 1979].

Cyclotron resonant interactions between plasma waves and energetic particles provide an efficient means of transferring energy and momentum within the Earth's magnetosphere. In this way plasma waves strongly influence the dynamics of the upper atmosphere, including the radiation belts. With the increasing number of technological assets orbiting the Earth in these regions, interest in plasma wave dynamics has grown. For example, there has been recent research into the growth and loss rates of relativistic electron populations within the magnetosphere, with an eye toward mitigating enhanced radiation levels via controlled precipitation of these highly energetic particles [*Inan et al.*, 2003]. Satellite missions to further explore and understand plasma wave phenomena are an integral part of such efforts.

This dissertation, then, is concerned with the circuitry of a "plasma wave receiver": a scientific instrument that measures plasma wave phenomena in situ, that is, while aboard a satellite flying through the upper atmosphere. In particular, this dissertation focuses on the analog-to-digital converter (ADC) component of such receivers, especially its requirements, design, implementation, and testing.

This chapter provides a general introduction to plasma wave receiver architectures. It first presents a more complete description of the signals of interest. Many different architectures—reflecting not only different scientific objectives but also different solutions to the limitations of technology and telemetry—have been employed to capture these signals. These architectures are next summarized and examples of

flown instruments are presented. Naturally, these architectures bespeak different ADC requirements. Based on an understanding of the current state-of-the-art in space electronics, a target plasma wave receiver architecture is described, and its ADC requirements specified.

1.1 Specifications

The signals of interest are plasma waves, that is, electromagnetic waves whose characteristics (such as their generation or propagation) are significantly affected by the presence of the plasma surrounding the Earth.

1.1.1 Scientific Specifications

For scientific analysis it is desirable to simultaneously acquire the magnitude and direction of both the electric and magnetic field components of plasma waves. Ideally, three-axis measurements of both fields are made: 3 electric field antennas capture the E_x, E_y, E_z components, and 3 magnetic field antennas capture the B_x, B_y, B_z components.[2] Concurrent knowledge of both sets of fields enables derivation of many fundamental wave properties, such as the polarization of the wave and its propagation vector (e.g., *Lauben et al.* [2002]). In addition, each field individually offers unique information. For example, electric field measurements can capture quasi-electrostatic waves: waves with a measurable electric field, but whose magnetic field is commonly too small to be detected.

Cartoon depictions of some plasma wave phenomena measurable by electric and magnetic field antennas are shown in the power spectral densities of Figures 1.1[3] and

[2]Electric and magnetic field antennas are of different design. Electric field antennas are usually dipoles or spherical double probes (in the latter, the field is measured as the signal difference between two spheres placed far apart along an axis). Magnetic field antennas are commonly wire loops or search coils (in the latter, wire is tightly wound around a high permeability core oriented along an axis). While different antennas translate to different front-end circuitry, after this interface the architecture of both electric and magnetic field receivers is often very similar.

[3]For electric field antennas, the received signal voltage $V = E(L_{\text{eff}})$ where E is the electric field and L_{eff} is the effective antenna length. Hence the field itself has units of V/m and the power spectral density of Figure 1.1 has units of $\text{V}^2/(\text{m}^2\text{-Hz})$.

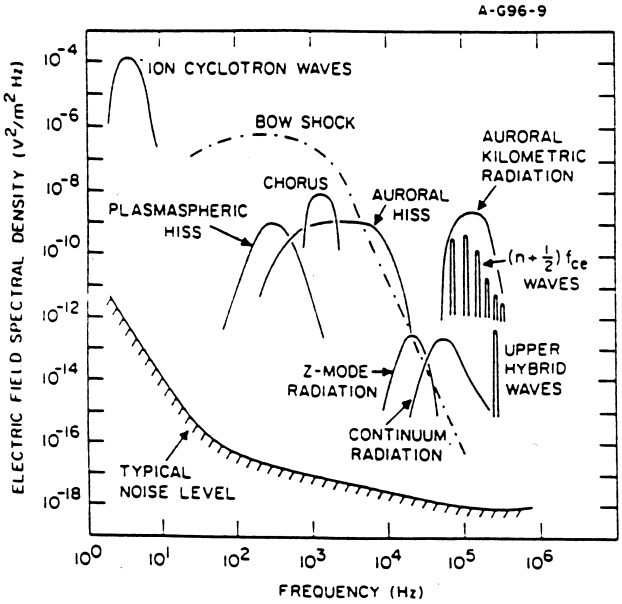

Figure 1.1: Power spectral densities of some plasma wave phenomena measurable by an electric field antenna, normalized by antenna length (in meters). Reproduced in total from *Gurnett* [1998].

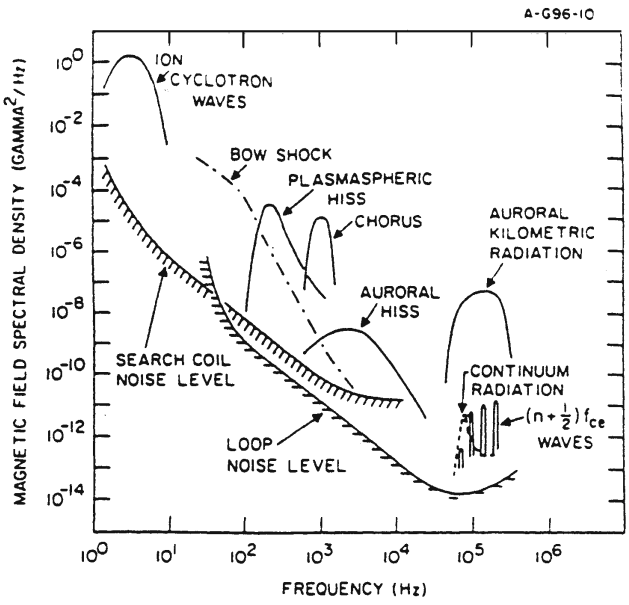

Figure 1.2: Power spectral densities of some plasma wave phenomena measurable by a magnetic field antenna. Note: 1 gamma equivalent to 1 nanoTesla. Reproduced in total from *Gurnett* [1998].

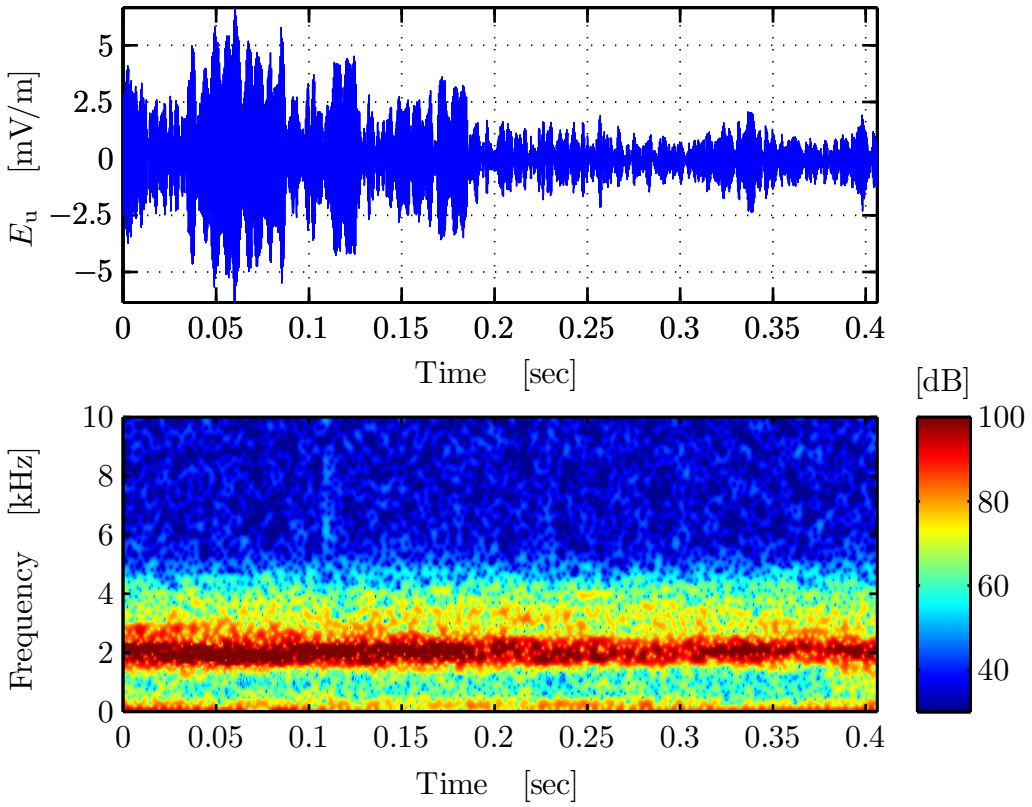

Figure 1.3: Electric field (E_u) measurements from the Polar plasma wave instrument. Both time domain data (above) and corresponding spectrogram (below) shown. Data measured by wideband receiver of said instrument.

1.2, respectively.[4] In both cases, the waves occupy a wide gamut of frequencies (from a few Hertz to hundreds of kiloHertz) and span a broad breadth of power densities (over 12 and 13 orders of magnitude—or 120 dB and 130 dB—in the electric and magnetic cases, respectively). Missing from these depictions, though, is the dynamic nature of plasma waves in time. This property is illustrated in Figure 1.3, which shows (at top) the captured time domain waveform of a plasma wave measured by an electric field antenna flying aboard the Polar satellite. Clearly, the signal evolves over time.

As they simultaneously capture both the frequency and time evolution of plasma

[4]Figures 1.1 and 1.2 are indicative, not complete: there exist additional phenomena not shown. However, the Figures do capture the expected frequency and power ranges of the received signals.

waves, spectrograms are used extensively throughout plasma wave signal analysis (e.g., *Inan et al.* [2004]). An example is shown in Figure 1.3 (at bottom). Time is plotted along the horizontal axis, frequency along the vertical axis, and the power spectral density is indicated by the color. Every vertical column thus represents the frequency content of the signal over a short time interval. The spectrogram thus illustrates the time evolution of the signal in the frequency domain. Indeed, the spectrogram of Figure 1.3 identifies the captured signal as hiss, a plasma wave phenomenon believed to be a driver of energetic electron precipitation in the radiation belts [*Abel and Thorne*, 1998].

A more complete example of a spectrogram is shown in Figure 1.4, which again shows plasma waves measured by an electric antenna aboard the Polar satellite, but now over a 24-hour period. Indeed, many plasma wave signal characteristics important for receiver design are exhibited in Figure 1.4. First, the ample frequency range and power extent (i.e., dynamic range) of the signals are evident. Second, the dynamic nature of plasma wave signals is seen. For example, the insert shows bursts of a phenomenon known as "chorus" in the upper frequency band versus the persistent hum of hiss in the lower frequency band. Indeed, many temporal characteristics, such as onset time and growth rate, prove important in plasma wave signal analysis (e.g., *Gołkowski et al.* [2008]). Finally, Figure 1.4 shows that multiple plasma wave phenomena can occur at the same time, for example, the aforementioned concurrent chorus and hiss. Generally, the received signal can consist of multiple, simultaneous plasma wave phenomena of different frequencies and powers; often, the received signal consists of a strong intensity phenomenon accompanied by a variety of weaker intensity phenomena (e.g., *Platino et al.* [2004]). Overall, then, the signal space of plasma waves spans frequency, power, and time, each of which is important in scientific study.

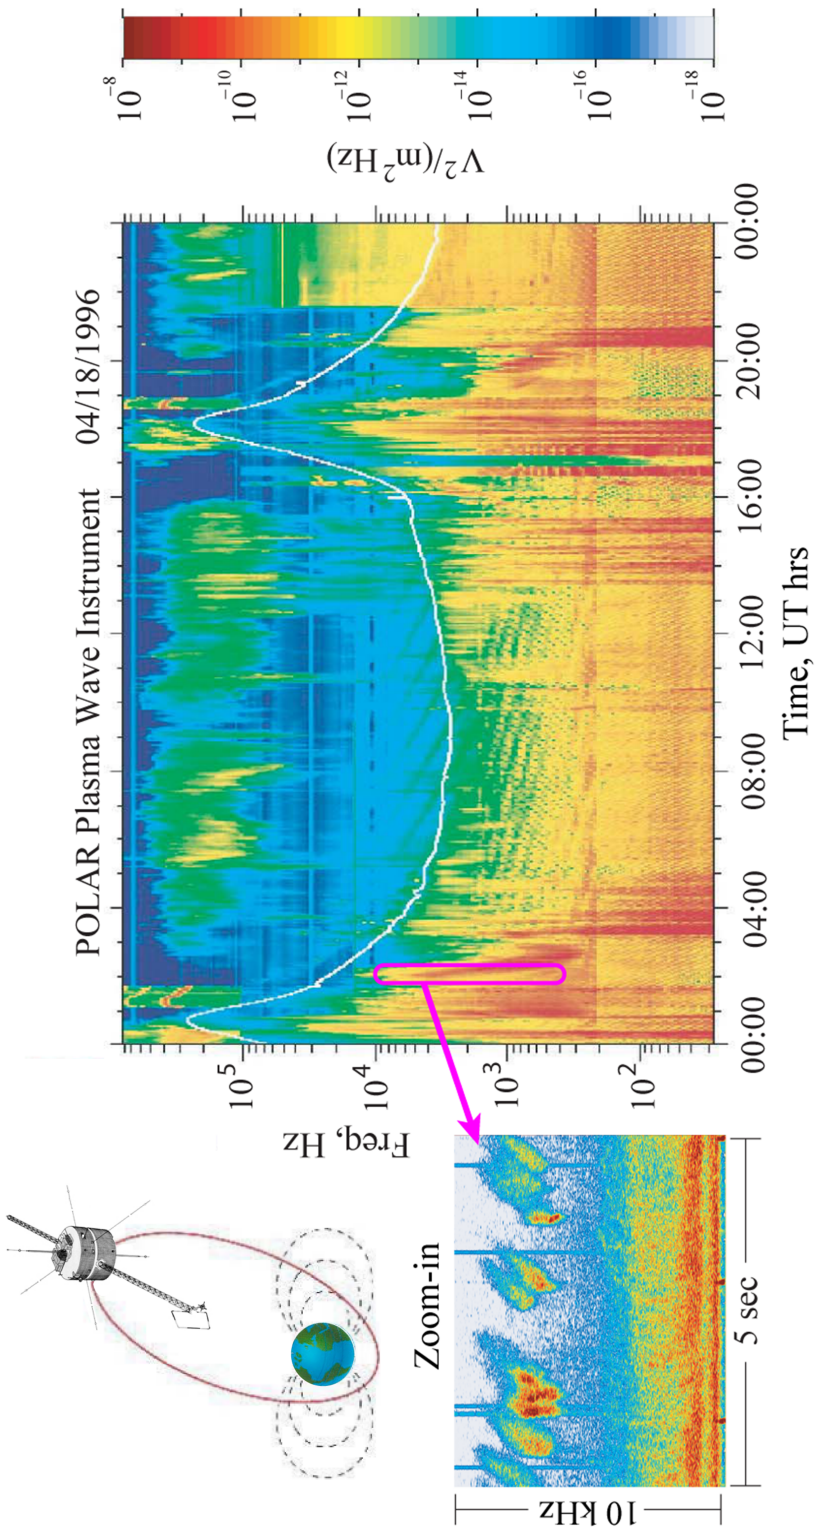

Figure 1.4: Spectrograms of electric field (E_u) measurements from the Polar plasma wave instrument. Complete 24-hour record shown to right, zoomed 5-second interval (including chorus and hum phenomena) shown in insert to left. Complete 24-hour record measured by sweep frequency receiver, and zoomed interval by wideband receiver, of said instrument.

1.1.2 Satellite Specifications

Not only must a plasma wave receiver capture the signal space, but also it must do so while satisfying a variety of specifications imposed by the spacecraft bus[5] and anticipated space environment. These additional specifications are important as they strongly influence the plasma wave receiver design and, by extension, the ADC requirements. Satellite-imposed specifications include resource allocation (mass, volume, power, telemetry, computation), environmental requirements (radiation tolerance, shock, vibration, temperature), and compatibility specifications (electromagnetic interference, electromagnetic cleanliness). While all are important in their own right, the requirements on power, telemetry, computation, and radiation tolerance often drive plasma wave receiver design more than the others.

Power is typically at a high premium on satellites, where flying additional solar panels or batteries incurs significant cost and complexity. This limitation is especially true for smaller spacecraft, or "microsatellites", whose smaller physical area and stricter mass requirement limit solar panel and battery use, respectively. However, even on larger spacecraft where more power is available, as the plasma wave receiver is typically but one of a suite of experiments, it is usually allocated but a small portion of the total bus power budget.

The receiver must also adhere to a telemetry budget. The task of transmitting data from a satellite in flight to the ground is a complicated technical problem subject to many concerns.[6] For this discussion, it suffices to say that telemetry is a shared resource on both the transmit (here the satellite) and receive (here the ground) sides.[7] Regarding the transmit side, on larger satellites multiple experiments must share the same capacity, whereas on smaller satellites the overall capacity is

[5]The spacecraft bus is defined as the portion of the spacecraft that supports the payload, the payload being the portion that directly accomplishes the mission objectives. The bus may include power, propulsion, attitude control, housekeeping, and telemetry systems. In common aerospace parlance, it also often refers to the physical orbiting vehicle that houses all these systems and the payload.

[6]These concerns include the transmission band, encryption overhead, and ground station coverage, among others.

[7]The inverse, where data is transmitted from the ground to the satellite, also occurs, but is typically insignificant in comparison to the sheer data volume of the satellite-to-ground telemetry in these applications.

less due to less available power. Regarding the receive side, reception is accomplished through ground station networks. Since operating and maintaining these networks is costly, networks are typically shared amongst many satellites of differing demands and priority, restricting receive-telemetry time for each spacecraft. The upshot is that plasma wave receivers can often only telemeter a limited amount of data to the ground.

To most efficiently make use of the available telemetry, plasma wave receivers often incorporate on-board signal processing to reduce data volume. Traditionally, such processing has been implemented within the receiver proper. However, modern spacecraft busses fly increasingly versatile digital microprocessors and memories, enabling use of increasingly sophisticated digital signal processing techniques in flight. This computation, though, is often a shared resource and the cycles afforded the receiver may be limited.

Finally, a satellite-based plasma wave receiver flying through the upper atmosphere is bereft of the radiation shielding enjoyed by terrestrial instruments. Thus the plasma wave receiver must be radiation tolerant, or "radiation hard" in aerospace parlance (the converse is "radiation soft"). The exact definition and degree of radiation tolerance depends on the particular orbit and bus and hence varies by flight.[8] Practically, though, it is notable that even at low tolerance levels, the realm of space-qualified, radiation-hard electronics is decidedly small, especially in comparison to its terrestrial counterpart. Many receiver designs are thus inevitably limited by part selection.[9]

[8]Shielding can be added to the bus to reduce radiation exposure of internal electronics. However, this protection is limited. First, shielding adds mass, a precious resource. Second, shielding only works up to certain particle energies: in a practical sense, it is often too expensive to shield against particularly high energy particles. Third, the effectiveness of shielding is limited as, due to interactions with the natural radiation environment, the shielding material itself functions as a secondary radiation source. Hence shielding efficiency typically drops dramatically past some thickness. (This decrease is demonstrated for several radiation environments in Appendix E of the *Handbook of Radiation Effects* [*Holmes-Siedle and Adams*, 2002, pp. 562–567].) Hence, while shielding can reduce the radiation exposure, in practice it often cannot eliminate it.

[9]It is possible to qualify terrestrial parts—called commercial-off-the-shelf (COTS) parts in aerospace vernacular—for satellite applications. The process, though, is not only expensive but also risky: as terrestrial parts are not explicitly designed for radiation environments, radiation testing may prove them radiation soft and thus unsuitable for flight.

1.2 Plasma Wave Instruments

Ultimately, the scientific and satellite specifications of a plasma wave receiver are uniquely determined by the objectives of a particular mission. Naturally, missions span a wide array of objectives and orbits, which has historically led to essentially custom plasma wave receiver designs. Nonetheless, over time common receiver architectures and approaches have emerged.

1.2.1 Canonical Architectures

In particular, three receiver architectures—the wideband receiver, the multichannel receiver, and the sweep frequency receiver [*Gurnett*, 1998]—are prevalent throughout much receiver design. While the cloth of any particular receiver is tailored for its specific mission, often it builds upon the patterns established by one or more of these canonical receivers.

To capture plasma wave signals, ideally a wideband, high dynamic range[10] receiver acquires the entire frequency and power range simultaneously, providing a very complete view of the signal. The architecture of a wideband receiver is shown in Figure 1.5. A low-noise amplifier (LNA) interfaces with the antenna to collect the signal in a low-noise manner. Wideband filtering, including anti-aliasing filtering, follows before an ADC digitizes the signal.[11]

Unfortunately, wideband receivers suffer from limitations both technological and telemetric. Technologically, wideband receivers require ADCs of high dynamic range

[10] Dynamic range here refers to the ability to capture signals of both large and small powers. In this dissertation, it is meant as a descriptive term and not a formal performance metric. Instead, the formal performance metrics are termed "instantaneous dynamic range" and "system dynamic range". The former refers to the dynamic range at a single instant in time and is given formal definition in Section 3.2.2.2. In the systems addressed here, it is typically set by the ADC. The latter refers to the dynamic range achieved by the system over time, which may include variable gain and log-compression stages, as is discussed in this section.

[11] Given the subject of this dissertation, these discussions concentrate on architectures that digitize the received signal with an ADC. Perhaps notably, such digitization is not strictly necessary: for example, both the plasma wave receiver [*Gurnett et al.*, 1978] and VLF receiver [*Bell and Helliwell*, 1978] aboard ISEE-A (launched October 1977) used a special analog telemetry link for transmitting wideband data.

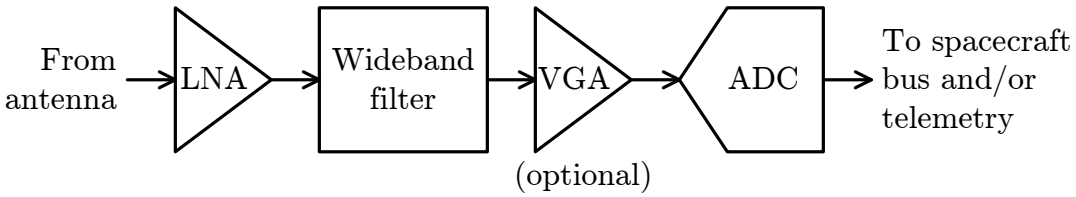

Figure 1.5: Wideband receiver architecture. Low-noise amplifier denoted as LNA, variable-gain amplifier as VGA, and analog-to-digital converter as ADC.

sampling at megasample-per-second (MS/s) rates.[12] Unfortunately, finding an ADC that achieves these requirements, and that is also radiation tolerant and can perform within the receiver power budget, can be quite difficult.[13] Hence, many wideband receivers employ a bandpass filter to exclude some phenomena and relieve the ADC dynamic range requirements. Another strategy introduces a variable-gain amplifier (VGA) before the ADC, which grants the receiver a wide system dynamic range though the instantaneous dynamic range is still ADC-limited. Telemetrically, directly transmitting a high-resolution, rapidly sampled digital output can rapidly overwhelm the receiver's telemetry budget. Hence, practically only short intervals of wideband receiver data can be transmitted. The consequent limited observation time not only compromises collection of longer time scale phenomena, but also complicates collection of shorter, burst phenomena as it necessitates close coordination between the manifestation of the phenomenon and the onset of acquisition.

To overcome these limitations, plasma wave receivers often employ architectures with increased on-board signal processing. Two architectures—the multichannel receiver and sweep frequency receiver—have emerged as common solutions over the years. Both implement a frequency decomposition of the signal via analog electronics upstream of the ADC. In doing so, they trade between time resolution and frequency resolution, the multichannel receiver favoring the former, the sweep frequency receiver favoring the latter.

The multichannel receiver, shown in Figure 1.6, achieves good time resolution,

[12] In this dissertation, ADC sampling rates are given in samples per second, or S/s, and prefixed as an SI unit.

[13] As discussed in more detail in Section 3.4.

but often diminished frequency resolution. The signal is divided into a series of frequency bands. Each band is averaged, often through root-mean-square (RMS) measurements. The bands are also often logarithmically compressed. The resulting low frequency amplitudes are sampled and then digitized by a rapidly cycled ADC. As shown in Figure 1.6, the sample-and-hold (S/H) circuits are often synchronized to maintain time coherence between bands.[14] Properly designed, a multichannel receiver uses a coarser ADC sampling at a slower rate, relieving both the technology and telemetry constraints. Naturally, many narrow bands are preferred. Unfortunately, the commensurate increase in hardware and power often restrains multichannel receivers to fewer, wider bands and hence impaired frequency resolution.

In contrast, the sweep frequency receiver maintains good frequency resolution, but diminished time resolution. Shown in Figure 1.7, the sweep frequency receiver heterodynes down[15] and narrowband filters the signal: in this way, it observes a small bandwidth of the original signal centered at the oscillator frequency. The oscillator frequency is then "swept" to scan through the frequency spectrum. Averaging, often through RMS measurements, and possibly with logarithmic compression, precedes digitization. Again, coarser, slower ADCs can be used, delivering technology and telemetry relief. Sweep frequency receivers often implement more and narrower frequency bands than multichannel receivers as doing so incurs minimal hardware cost.[16] However doing so also lengthens the frequency sweep times, decreasing time resolution.

Given their proclivities, multichannel receivers work well in providing continuous spectra with good time resolution but poor frequency resolution, whereas sweep frequency receivers work well in providing continuous spectra with good frequency

[14] Alternately, each band can sport its own ADC and a high speed digital multiplexer transfer the results to the bus. Yet another option is to forgo the S/Hs and rapidly cycle an ADC through the averager outputs: this choice sacrifices some time coherence for less hardware.

[15] Often single-sideband mixers are used to prevent interference from otherwise-mirrored negative frequencies.

[16] More bands requires adding frequencies to the local oscillator frequency sweep, which is easily done if the local oscillator is digitally synthesized. Narrower bands requires adjusting the narrowband filter, a one time cost. These changes are in contrast to the multichannel receiver, which requires adding an entire extra analog signal path for each additional band.

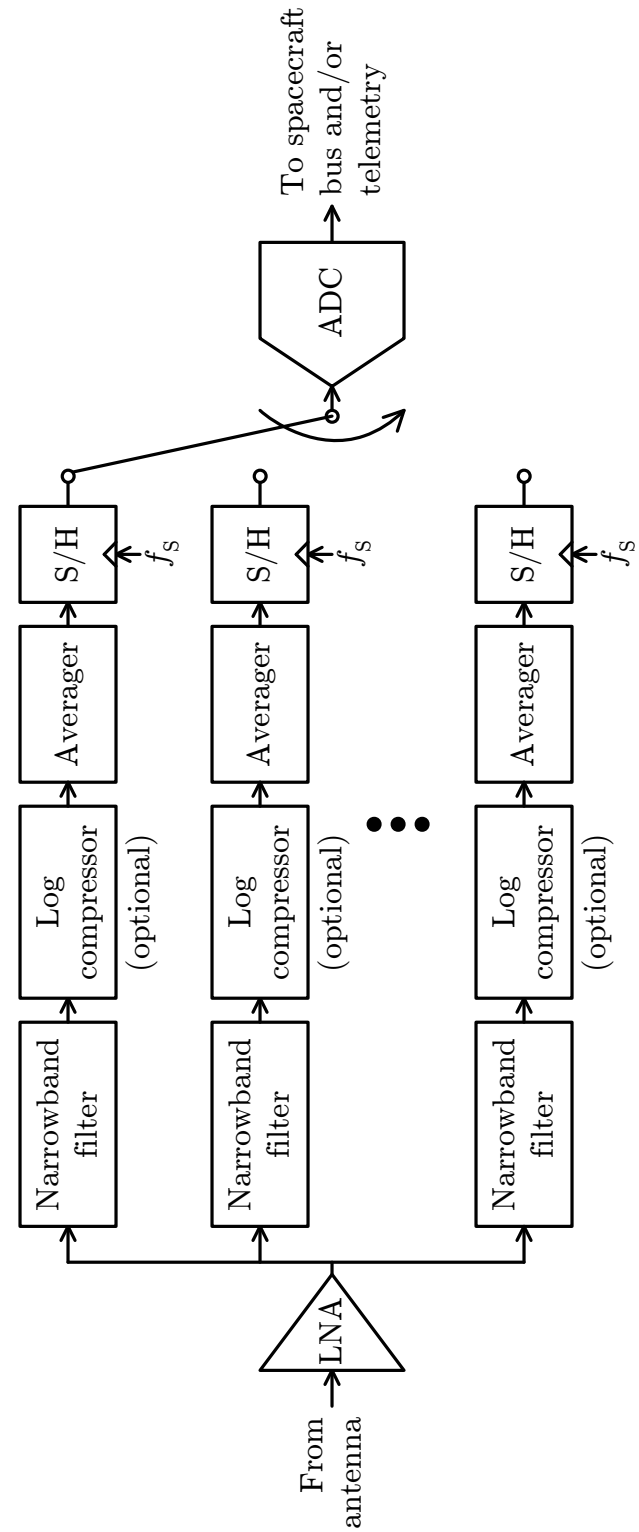

Figure 1.6: Multichannel receiver architecture. Sample-and-hold circuits denoted as S/H.

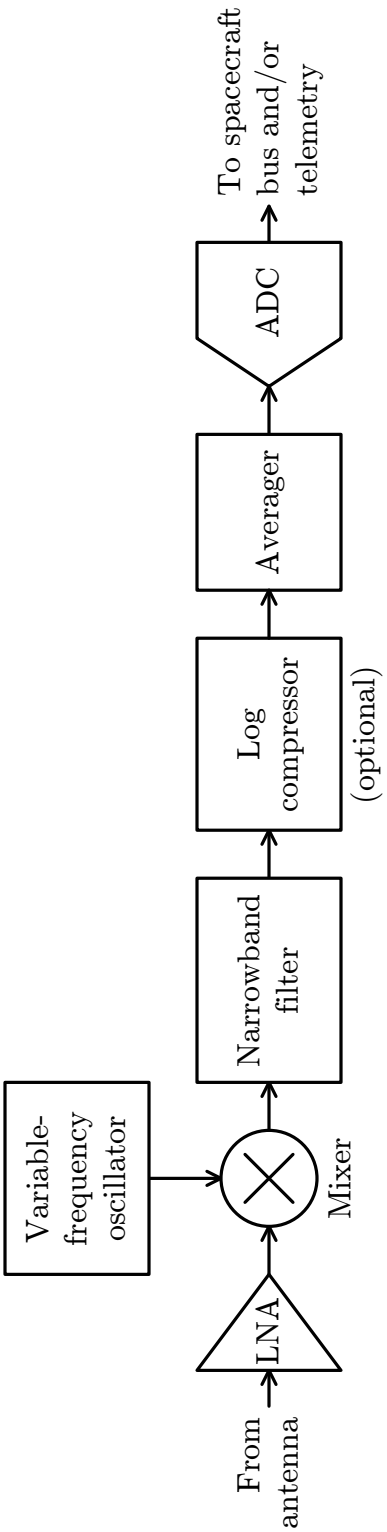

Figure 1.7: Sweep frequency receiver architecture.

resolution but poor time resolution [*Gurnett*, 1998, p. 134].[17]

1.2.2 Flown Examples

The three canonical architectures present different solutions to balancing signal resolution and instrument telemetry. In practice, many plasma wave experiments opt to fly a mix of all three to form a final plasma wave instrument: during flight, operators dynamically adjudge the best solution and activate the appropriate receiver(s). This method leads to highly flexible designs that accommodate a wide variety of signal capture modes.

An example is the plasma wave instrument flying aboard the Polar satellite (launched in February 1996), shown in Figure 1.8 [*Gurnett et al.*, 1995]. Signals from 7 antennas—3 electric and 4 magnetic—are passed through an analog switch network to route different signal combinations to different receiver combinations. In all, there are 7 receivers, including all three of the canonical types. The sweep frequency and multichannel receivers provide low-rate, log-compressed signal summaries for long portions of the orbit and feed a low-rate processor. On the other hand, the wideband receivers provide high resolution signal capture, but only for short periods of time, typically on the order of seconds.[18] Three wideband channels—the low frequency waveform receiver, the high frequency waveform receiver, and the wideband waveform receiver—offer various combinations of signal bandwidth, sampling rate, and bit resolution, allowing finer management of capture and telemetry tradeoffs. The low frequency waveform receiver uses 6 12-bit ADCs to sample 6 channels simultaneously at 100 S/s. The high frequency waveform receiver, which also uses 6 synchronized 12-bit ADCs, can sample at higher rates from 558 S/s to 71.43 kS/s, with appropriate bandlimiting implemented by a combination of analog and digital filtering.[19] And the

[17] Sweep frequency receivers, for example, have thus been used for tracking relatively steady narrow-band quantities, such as the upper hybrid resonance and electron cyclotron frequency [*Gurnett et al.*, 1978].

[18] Indeed, a sweep frequency receiver was used to capture the 24-hour record of Figure 1.4 (when compressed over long time periods, the data gives the appearance of being a continuous record), whereas a wideband receiver acquired the zoomed 5-second interval.

[19] That is, to accommodate anti-aliasing filter transition bandwidths, the analog signal is slightly oversampled and the transition bandwidth filtered out in the digital domain.

wideband waveform receiver offers modes with not only different ADC sampling rates, but also different ADC resolutions, from 8 bit to 1 bit.[20] All three are processed by a high-rate processor that conducts basic signal processing (including algorithms that automatically control the receiver gains) in addition to data formatting and storage. In all, this flexibility enables the operator to choose the best acquisition-telemetry balance in flight, based on the expected phenomena and desired science.

As modern satellites sport increasing amounts of digital signal processing and data storage, modern instruments opt to implement flexibility more in the digital than analog domain. A precursor to this new class of instruments is the fields instrument flying aboard the Fast Auroral SnapshoT (FAST) satellite (launched in August 1996) shown in Figure 1.9 [*Ergun et al.*, 2001]. Broadly, the instrument operates in either survey mode or burst mode: the former provides continuous-over-orbit data coverage but at low resolution, while the latter provides high-resolution data but only over a few minutes of orbit. Again, a variety of antennas can be routed to a variety of receivers. Here, though, the concentration is on high resolution wideband receivers. Survey data is acquired through the survey waveform receiver, which scrolls a 16-bit ADC between 18 channels for effective sampling rates of 512 S/s to 2048 S/s per channel, and a 3-channel sweep frequency receiver. Burst data is acquired through the burst waveform receiver, which contains 8 16-bit ADCs sampling simultaneously at 32.768 kS/s, and a high speed burst memory, which acquires 4 10-bit channels at 2 MS/s. The instrument relies heavily on on-board digital signal processing for data reduction. Much of this processing is performed by the digital signal processor (DSP): for example, the DSP handles signal cross-correlations for phase computations, performs fast Fourier transform (FFT) calculations for creating output spectra, and processes results from the plasma wave tracker (a modified channel of the sweep frequency receiver) to intelligently adjust the tracker's sweep range based on in-flight derivations of plasma parameters. To enable the DSP to concentrate on signal processing tasks, many housekeeping and interface functions are implemented by custom field-programmable gate arrays (FPGAs) (not shown), and data formatting

[20]Specifically, the wideband waveform receiver can sample at rates from 31.12 kHz to 249 kHz, relying on bandlimiting and subsampling to capture bandwidths of 11 kHz to 90 kHz over the full 600-kHz input frequency range.

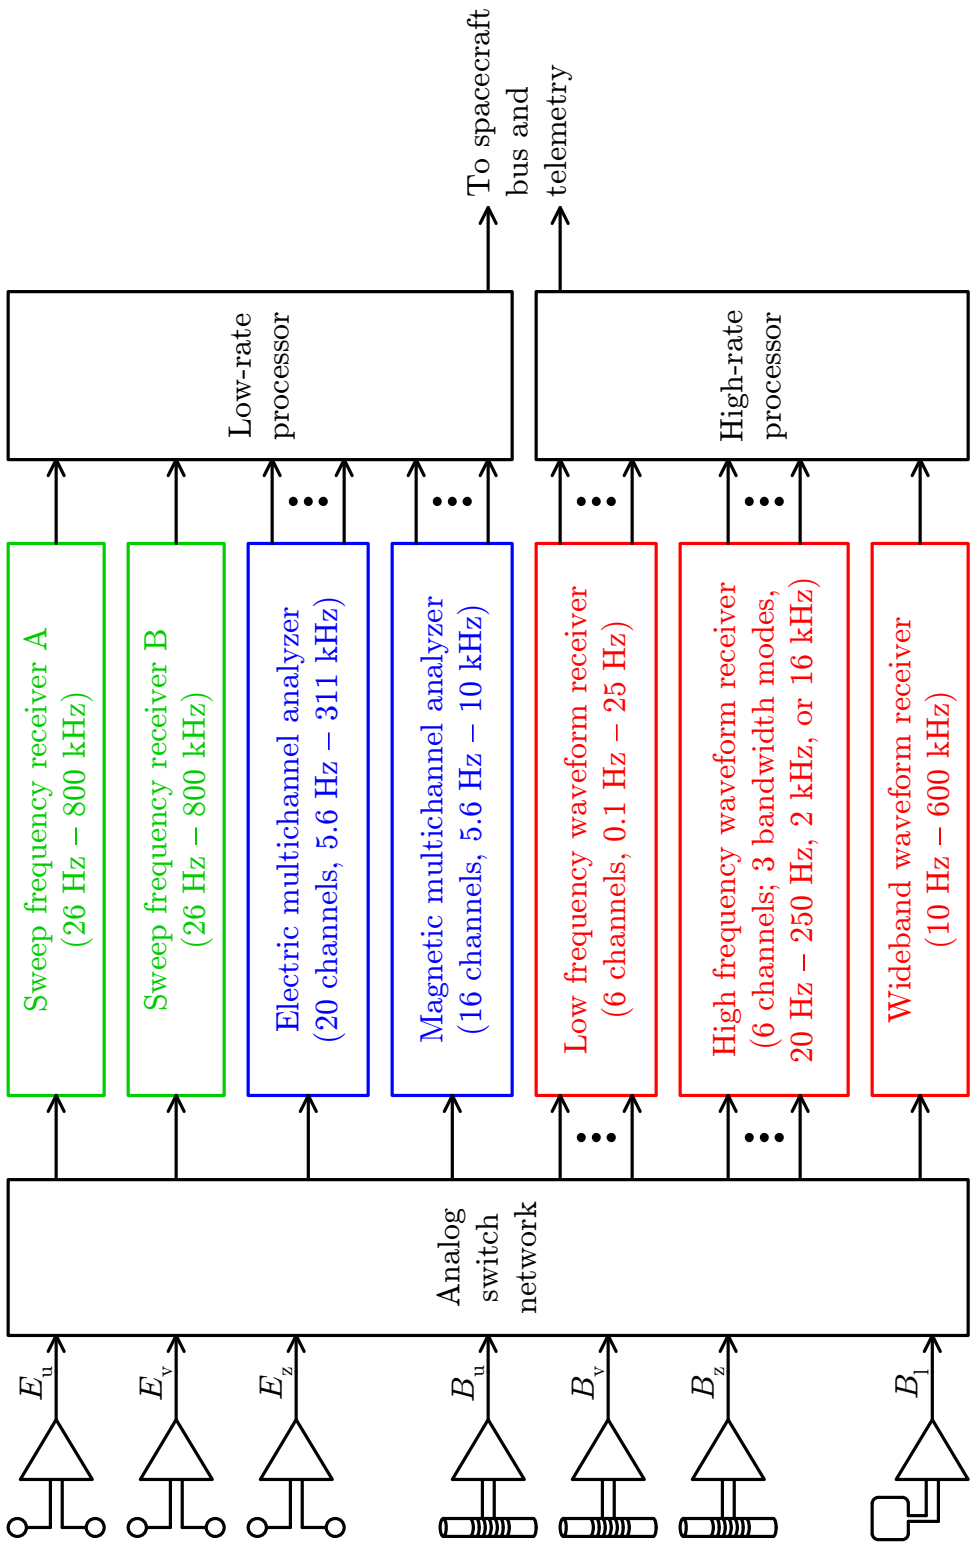

Figure 1.8: Plasma wave instrument aboard the Polar satellite (as described by *Gurnett et al.* [1995]). Multichannel receivers shown in blue, sweep frequency receivers in green, wideband receivers in red.

is handled by a central instrument data processing unit (IDPU) and 125 Mbyte on-board memory shared amongst the FAST satellite experiments.[21]

1.3 Target Plasma Wave Receiver

This dissertation concentrates on ADCs for next generation plasma wave instruments akin to that flown on FAST. On such instruments, increasing amounts of on-board digital signal processing and storage can be used for frequency decomposition, data compression, advanced data product calculation, or other scientific analyzes. The objective of the analog electronics is then to capture as much of the signal range as possible, leaving the implementation of signal analysis, instrument flexibility, and data reduction to the digital domain.

The SVADC-1 ADC of this dissertation thus targets a wideband receiver. This target receiver is designed to simultaneously capture the signal space enclosed within the dashed box of Figure 1.10, encompassing a broad frequency range from 100 Hz to 1 MHz, and the full 90-dB power range of signals over this bandwidth. Notably, the bandwidth does not extend to DC. DC electric field measurements entail unique challenges[22] and thus DC and AC measurements are typically handled by separate receivers: the target receiver focuses on AC measurements. As for the low frequency cutoff of 100 Hz, this limit arises from practical concerns. An electric antenna flying through a plasma accumulates DC charge. On a spinning satellite,[23] this DC component Doppler shifts to twice the spin frequency, resulting in a large interferer at a few to tens of Hertz. The 100-Hz cutoff allows sufficient transition bandwidth for suppressing such spin-induced signals. Finally, while the target receiver is designed to measure the electric field signals of Figure 1.10, it can also be adapted to measure

[21] For comparison, the low-rate and high-rate processors of the Polar instrument are 8085 and 8086 microprocessors, respectively [*Gurnett et al.*, 1995], whereas the DSP and IDPU of the FAST instrument are a 32-bit floating point ATT-DSP32C (clocked at 32 MHz) and a Sandia SA3300 32-bit microprocessor (clocked at 10 MHz), respectively [*Harvey et al.*, 2001].

[22] A good review of these challenges—from antenna impedance to paint choice—is given by *Maynard* [1998].

[23] As $V = E(L_{\text{eff}})$, longer antennas garner better signal strength. Thus many instruments fly antennas of tens of meters. However, satellites with such long antennas are difficult to attitude stabilize, and hence often allowed to spin.

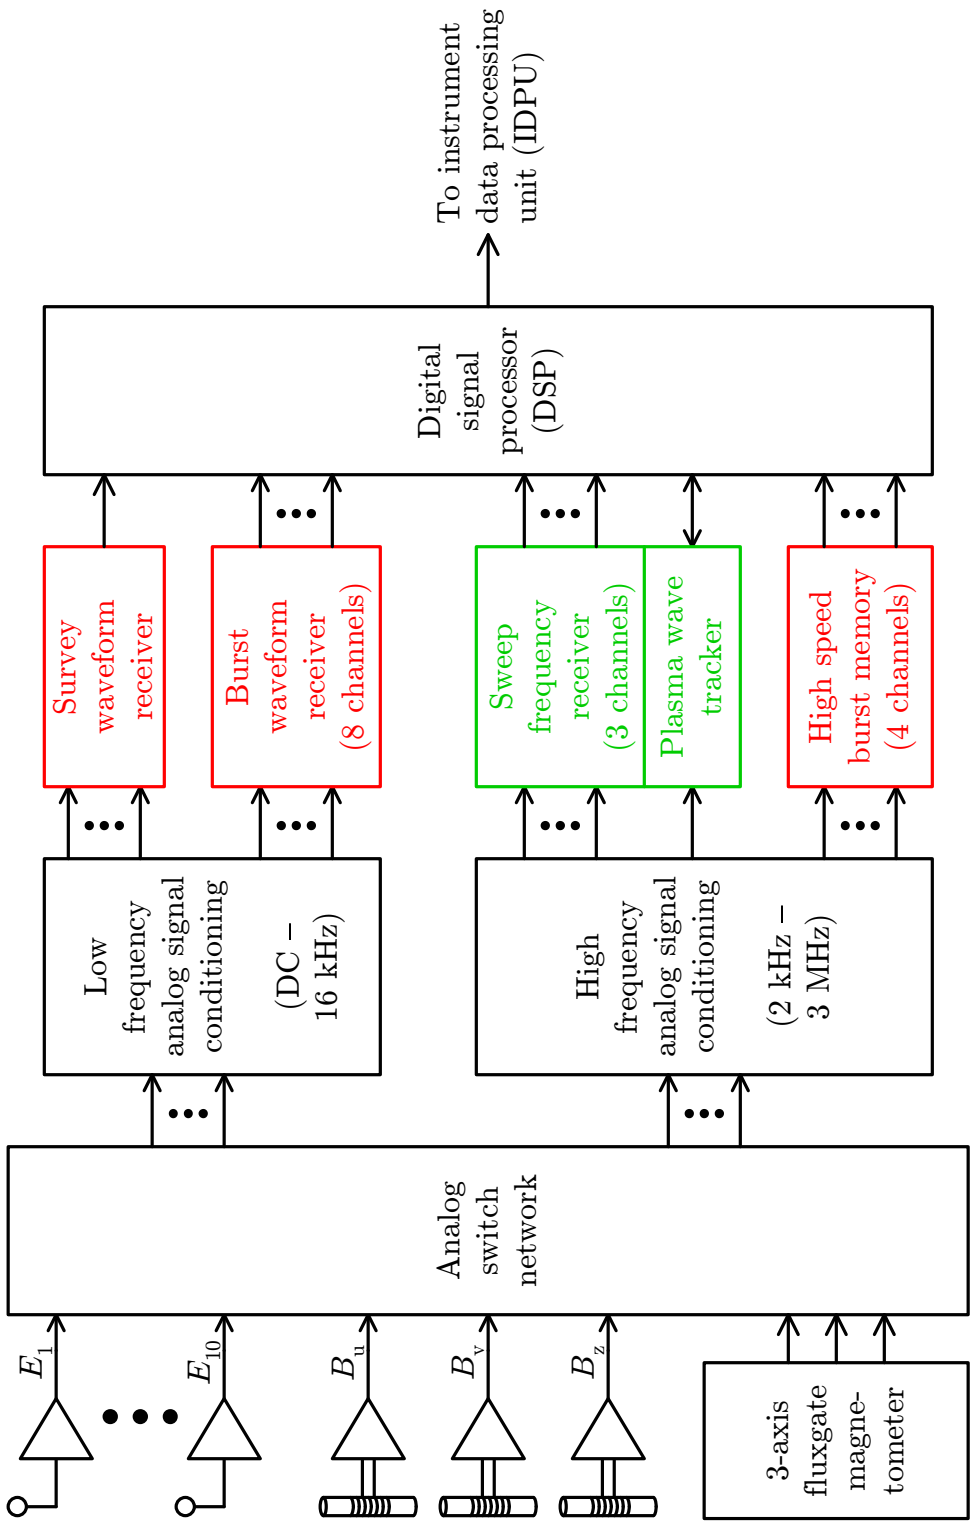

Figure 1.9: Fields instrument aboard the FAST satellite (as described by *Ergun et al.* [2001]). Multichannel receivers shown in blue, sweep frequency receivers in green, wideband receivers in red.

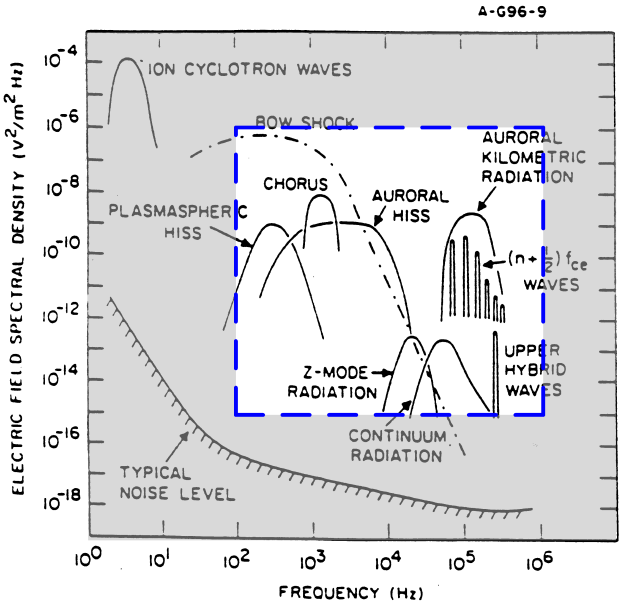

Figure 1.10: Signal space captured by the target plasma wave receiver. Modified from *Gurnett* [1998] (see Figure 1.1 for original).

a similar span of the magnetic field signals of Figure 1.2 with appropriate changes to its front-end circuitry.

The target wideband receiver architecture is shown in Figure 1.11. The input signal is acquired by an electric dipole antenna interfaced to an LNA; magnetic field measurements can be acquired with appropriate antenna and LNA modifications. The signal is then passed through an anti-aliasing filter (AAF) in preparation for analog-to-digital conversion. The ADC digitizes the signal, both sampling it in time and quantizing it in value. The digital signal is processed by a dedicated FPGA that implements simple signal processing in addition to the standard tasks of system control and data formatting. The processed signal can then be sent to an on-board computer for more advanced digital signal processing, or directly to telemetry.

Key specifications for the various blocks are given in Figure 1.11. To capture the full range of powers in Figure 1.10, a specification of 90-dB spurious-free dynamic

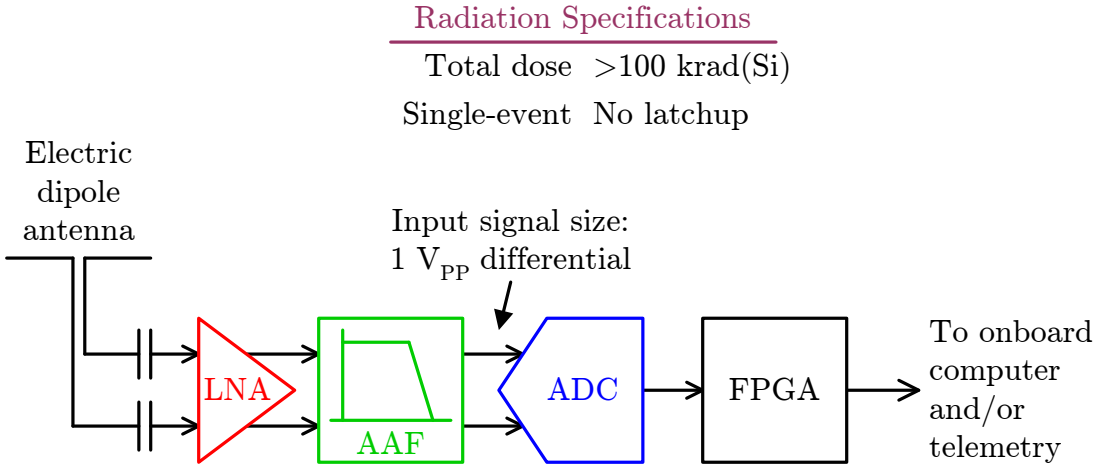

LNA Key Specifications

Bandwidth	100 Hz – 1 MHz
Input-referred noise	300 nV/Hz$^{1/2}$
Input impedance	1 GΩ ‖ 100 fF
SFDR	90 dB
Power	5 mW

ADC Key Specifications

Sampling rate	5 MS/s
SFDR	90 dB
Power	60 mW

AAF Key Specifications

Architecture	6th-order Chebyshev (Type I)
Cutoff frequency	30 kHz, 180 kHz, or 1080 kHz
SFDR	90 dB
Power	35 mW

Figure 1.11: Target plasma wave receiver with key specifications labeled. Anti-aliasing filtered denoted as AAF. All components should satisfy stated radiation specifications.

range (SFDR) is imposed throughout the system.[24] Concentrating on the ADC,[25] the upstream gain is set so that its input signal range is 1 V_{PP} fully differential. To capture the 1-MHz signal bandwidth of Figure 1.10, the AAF cutoff of 1080 kHz is selected[26] and the ADC sampling rate set to 5 MS/s to allow sufficient bandwidth for the AAF to roll off to 96-dB suppression by the edge of the first aliased spectral island at 4 MHz. Note that the ADC is thus oversampled: for evaluating ADC performance, only the frequencies of the 100 Hz to 1 MHz signal bandwidth are considered. The ADC is allocated a total power consumption of just 60 mW. This budget is based on the construction of a complete 6-channel plasma wave instrument with power consumption comparable to the FAST fields instrument (see Appendix G), however, it is notable that this specification also renders the ADC viable for microsatellite applications where power is severely constrained. Finally, the ADC should be radiation hard, specifically, it (and, indeed, the entire receiver) should maintain performance up to a total dose of at least 100 krad(Si) and display no latchup. Notably, 100 krad(Si) is a typical 1-year dose for an orbit through the radiation belts:[27] maintaining performance to higher total doses is thus desirable as it translates to longer device lifetimes.

1.4 Contributions

Searches of available ADCs reveal that, currently, there are no ADCs that meet the requirements of Figure 1.11 (see Section 3.4). Hence, this dissertation describes the

[24] For the ADC, the SFDR is computed assuming an FFT bin width of 100 Hz. A complete definition of, and explanation for, the SFDR specification is given in Sections 3.2.2.2 and 3.2.3, respectively.

[25] More detailed examinations of the LNA and AAF specifications can be found in the dissertation of fellow Ph.D. student Benjamin J. Mossawir, which describes the design of a custom ASIC that implements these functions [*Mossawir*, in preparation].

[26] The multiple AAF cutoff frequencies accommodate the different bandwidth needs of different missions [*Mossawir*, in preparation]. If desired, the ADC can be appropriately downclocked for lower bandwidth applications. This dissertation, though, focuses on the full, 1-MHz (1080-kHz AAF cutoff) application.

[27] Curves of dose for various 1-year orbits are given in Appendix E of the *Handbook of Radiation Effects* [*Holmes-Siedle and Adams*, 2002].

design of the SVADC-1, an application-specific integrated circuit (ASIC) custom-designed to meet these specifications. This dissertation traces the entire development of the SVADC-1, from derivation of its required specifications, to the fashioning of its architecture, to the implementation of its circuits in a commercial 0.25-μm CMOS process, to its measured performance including its radiation performance. The primary contribution of this dissertation is thus:

- Design, implementation, and testing of an SVADC-1 chip that achieves at least 90-dB SFDR (assuming 100-Hz/bin spectral resolution) while sampling at 5 MS/s and consuming 60 mW. The SVADC-1 is fabricated in a commercial SiGe BiCMOS process, although only the 0.25-μm CMOS layers are used. While not manufactured in a radiation-hard process, the SVADC-1 nonetheless maintains the aforementioned performance up to a total dose of 1 Mrad(Si), and remains functional up to a total dose of 2 Mrad(Si) (the highest tested dose). It also displays no latchup up to an effective LET of 63 MeV-cm^2/mg (the highest tested LET) at elevated temperature (131°C) and supply voltage (2.7 V). (Discussed in Chapters 4, 5, and 6.)

The development of the SVADC-1 entails other contributions as well:

- Formulation of the analog-to-digital conversion requirements for wideband electromagnetic wave measurements. (Discussed in Chapter 3.)

- Description—by both analytical and numerical methods—of the effect of circuit noise on the SFDR of quantizers. (Discussed in Chapter 3.)

- Formulation of a novel self-calibration technique to improve pipeline ADC linearity by estimation of the discontinuity heights in the transfer function of the front-end stages via DAC differencing. This self-calibration technique allows the SVADC-1 to achieve high SFDR without incurring high power consumption. (Discussed in Chapter 4.)

- Description of a two-gain technique for assessing the effect of operational amplifier nonlinearity on switched-capacitor amplifiers. (Discussed in Chapter 4.)

- Demonstration of radiation-hardness by design at the architecture, design, and layout levels for the implementation of radiation-hard circuits in a commercial manufacturing process. Adoption of this design philosophy affords the SVADC-1 its high level of radiation tolerance. (Discussed in Chapter 5.)

While designed for plasma wave receiver applications, it should be clear that both the SVADC-1 and the principles of its design are applicable wherever high-fidelity, radiation-hard analog-to-digital conversion is needed.

1.5 Organization

This chapter provided a basic introduction to the science of plasma waves and the principles of plasma wave receiver design. It then formulated a wideband plasma wave receiver for next-generation plasma wave measurements, leveraging the increasing amounts of on-board digital signal processing and storage available on modern satellites. The key specifications of this receiver were given, including those for the ADC component. Construction of the SVADC-1 to meet these specifications, then, is the central work of this dissertation.

The ADC requirements can be broadly divided into radiation requirements and analog-to-digital conversion requirements. These requirements are addressed in more detail in Chapters 2 and 3, respectively. Specifically, Chapter 2 provides a basic introduction to the effect of radiation, both short-term and long-term, on MOSFET-based electronics. Chapter 3 then considers analog-to-digital conversion and derives the SVADC-1 specifications in more detail. Given a better appreciation for the challenges of designing electronics for radiation environments, and a better understanding of the analog-to-digital conversion requirements of the SVADC-1, Chapter 3 concludes by considering the field of currently available, radiation-tolerant, high-fidelity ADCs. No currently available ADC is found that meets all the requirements, validating construction of the SVADC-1.

To perform the necessary conversion, the SVADC-1 employs a pipeline architecture. This architecture is reviewed in Chapter 4. Pipeline converters can achieve high conversion rates, but are often limited in their SFDR due to manufacturing tolerances

and matching. However, calibration techniques can overcome these nonidealities. Such techniques are reviewed in Chapter 4, and a novel DAC-differencing self-calibration technique proposed for the SVADC-1.

The design and implementation of the SVADC-1 chip is the subject of Chapter 5. Naturally, the SVADC-1 incorporates the novel self-calibration technique. In addition, since the SVADC-1 is fabricated in a commercial process, a philosophy of radiation-hardness by design is adopted to compensate for performance loss from radiation effects. The measured results of the fabricated SVADC-1 chip, including both baseline and radiation performance, are presented in Chapter 6. These results confirm the efficacy of both the self-calibration technique and the radiation-hardness-by-design approach. Finally, Chapter 7 concludes by reviewing the major contributions of this dissertation and postulating future directions.

A number of appendices are attached to this dissertation. The first set of appendices concentrate on radiation testing, with Appendix A presenting a general introduction to the procedures of such testing. The next five appendices then present the specifics of the radiation testing undertaken in support of this dissertation. Specifically, Appendix B presents the 50-MeV proton total-dose testing of the SVADC-1, Appendix C presents the 10-MeV heavy-ion single-event testing of the SVADC-1, Appendix D presents the 25-MeV heavy-ion single-event testing of the SVADC-1, and Appendix E presents the pulsed-laser single-event testing of the SVADC-1. Finally, Appendix F addresses the ^{60}Co total-dose testing of a series of test NMOS transistors fabricated on the SVADC-1 chip alongside the converter proper.

The remaining appendices appear in the order in which they are quoted in the dissertation. These appendices include Appendix G, which details the derivation of the ADC power specification presented earlier in this chapter; Appendix H, which presents a statistical consideration of the impact of Gaussian noise on nhSFDR assuming an N-point FFT; Appendix I, which gives a general introduction to $\Sigma\Delta$-modulators; Appendix J, which describes a series of test circuits included on the SVADC-1 chip for radiation-testing purposes; and Appendix K, which presents the results of the burn-in testing of the SVADC-1.

Chapter 2

Radiation Effects

Denied the degree of shielding afforded terrestrial electronics by the Earth's dense neutral atmosphere, spaceborne electronics are exposed to radiation during flight. This radiation can cause both transient and permanent changes in the functionality and performance of these electronics, from the sensors and passive components to the active devices and integrated circuits. While the radiation responses of all these components are important, this chapter focuses on the impact of radiation on the last—the integrated circuits—and specifically those fabricated in a CMOS technology.

The chapter begins with a brief introduction to the near-Earth radiation environment and describes the effects of this radiation. In particular, the chapter focuses on ionization damage, the dominant damage mechanism for MOSFETs. The consequent radiation effects are divided into two categories: single-event effects and total-dose effects which, broadly, describe the immediate and long-term effects, respectively.[1] Each category is then discussed with an emphasis on those effects that most affect circuit functionality and performance. Notably, the description of total-dose effects includes measured radiation testing results for the manufacturing process used to fabricate the SVADC-1. Of particular interest is the measured radiation evolution of the leakage current of NMOS devices, which shows a pronounced increase

[1] For a more thorough introduction to the radiation effects discussed in this chapter, the reader is recommended to the book *Ionizing Radiation Effects in MOS Devices and Circuits*, edited by *Ma and Dressendorfer* [1989]. While an older text, much of its material remains relevant for modern CMOS technologies.

of about 6 orders of magnitude with dose. As is discussed in Chapter 5, such increases can subsequently compromise the performance of switched-capacitor circuitry.

The majority of these effects are described at the device level. It should be noted that the discussions thus describe a moving target: CMOS technologies are constantly evolving their process steps, structures, and materials, all of which can affect radiation response. This chapter, then, is restricted to those effects most pertinent to submicron CMOS processes.

2.1 The Radiation Environment

Electronics are deployed throughout a variety of radiation environments, from space to nuclear reactors to high-energy physics accelerators. For plasma wave receivers, the space environment—including the space environment near the Earth or other planets—is most relevant. Thus, for this dissertation, the focus is on near-Earth space-based applications.

2.1.1 Near-Earth Radiation Environment

Near-Earth space is home to a vast array of energetic particles, which form the radiation environment. This radiation is often classified into two primary elements: trapped radiation (especially the radiation belts) and transiting radiation.

The magnetic field of the Earth forms a semi-permeable obstacle to the plasma pervading space. The flow of the solar wind[2] past the Earth distorts the shape of the magnetic field, and thus a vast geomagnetic cavity, known as the magnetosphere, is formed. The magnetic field configuration of the Earth's magnetosphere, shown in Figure 2.1, supports long-term trapping of energetic ions and electrons. Note that Figure 2.1 shows but a typical magnetosphere: in reality, the exact extent and internal structure is sensitive to changes in the solar wind density and velocity, as well

[2]The solar wind is a plasma of energetic electrons and protons that streams continuously outward from the Sun.

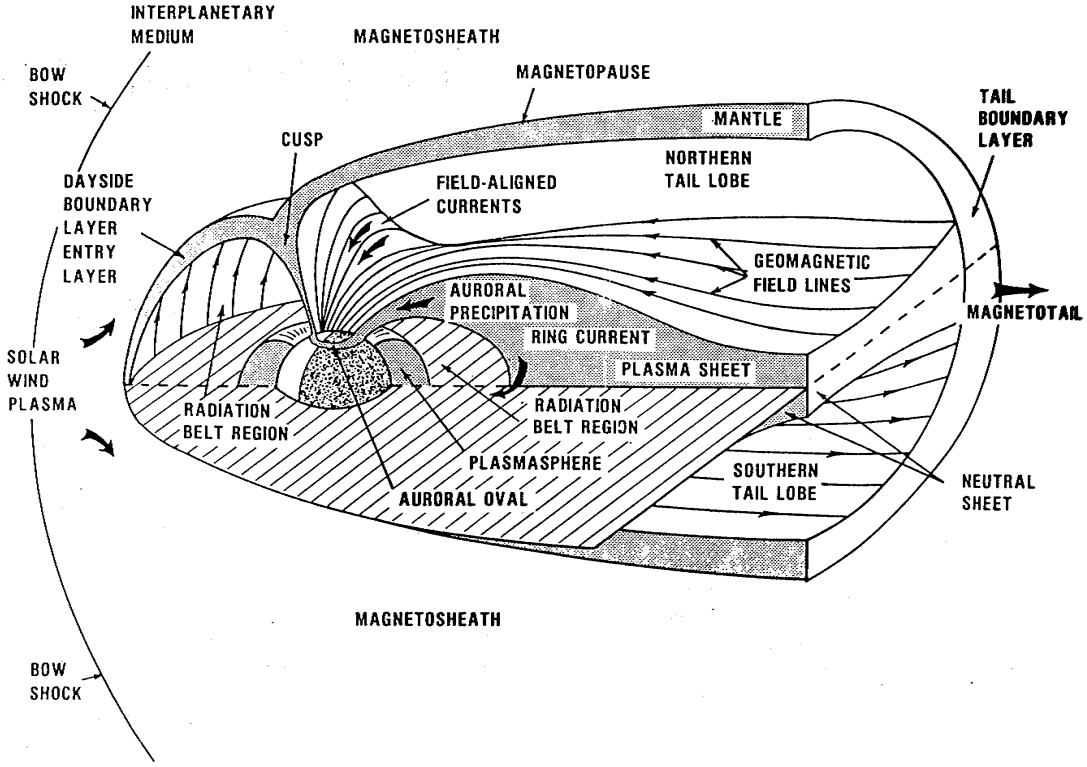

Figure 2.1: Illustration of the Earth's magnetosphere highlighting several internal structures, including the radiation belts. Earth (encircled by plasmasphere) dayside and nightside shown: Sun (not shown) is to the left. Reproduced in total from *Tascione* [1994].

as changes in the Earth's magnetic field.[3] Variations occur with season, solar cycles, and geomagnetic storms, to name a few dependencies [*Stassinopoulos and Raymond*, 1988]. Hence, the radiation environment is not only a function of latitude, longitude, and altitude, but also of time.

Especially pertinent for space electronics are the radiation belts (also called the Van Allen belts), depicted schematically in Figure 2.2. These belts are regions of

[3]The magnetosphere is also subject to cosmic ray flux and man-made effects. An oft-given example of the latter is the Starfish incident, an exoatmospheric nuclear explosion experiment conducted by the United States on 9 July 1962. Coupled with a series of similar Soviet experiments later that year, the strength of the radiation belts was sufficiently increased to cause failure of multiple assets, including the Telstar 1 communications satellite [*Teague and Stassinopoulos*, 1972; *Hughes*, 1989, p. 47].

significant fluxes of high-energy trapped particles, including mostly electrons (up to a few MeV) and protons (up to several hundred MeV) [*Holmes-Siedle and Adams*, 2002, p. 18], but also heavy ions (up to a few MeV) [*Walt*, 1994, p. 81]. For electrons, two altitude regimes of greater flux are identified as the inner and outer belts, which extend to roughly 2.4 R_E and span roughly 2.8 R_E to 10 R_E, respectively[4] [*Walt*, 1994, p. 80]. The highest energy electrons tend to occur in the outer belt. For protons, both flux and energy decrease with distance from Earth; an outer boundary of about 3.8 R_E is sometimes given [*Holmes-Siedle and Adams*, 2002, p. 18]. For heavy ions, flux peaks around 3 R_E [*Walt*, 1994, p. 81]. Heavy ions include mostly helium and oxygen, and are believed to originate out of the upper atmosphere. Indeed, the atmosphere presents a lower boundary for the radiation belts: owing to increased interactions with constituents of the increasingly dense atmosphere, the lower extent of the belts is at altitudes of 200–1000 km, depending on latitude and longitude [*Walt*, 1994, p. 4].[5] Many space assets pass through the belts in the course of their orbits around the Earth. For example, the Hubble Space Telescope flies at altitudes of 500–600 km, subject to strong fluxes of mostly protons but also inner belt electrons [*Holmes-Siedle and Adams*, 2002, p. 22]. And the geostationary orbit, a popular location for communications satellites, is approximately 6.6 R_E and subject to high energy outer belt electrons though relatively small proton flux [*Holmes-Siedle and Adams*, 2002, p. 25].

Satellites in geostationary orbits are also strongly subject to transiting radiation, especially galactic and solar cosmic rays, which are energetic particles originating from outside the solar system and from the Sun, respectively. While of less flux than their trapped counterparts, this radiation is nonetheless energetic enough to cause disruptions in spacecraft electronics. Galactic cosmic rays consist mostly of protons and heavy ions of high energies: near Earth, the greatest flux is for particles of 1 GeV per nucleon, but can extend well past 10 GeV per nucleon (although rarely) [*Stassinopoulos and Raymond*, 1988]. Solar cosmic rays are dependent on solar

[4] R_E is the mean radius of the Earth: 1 $R_E = 6370$ km [*Walt*, 1994, p. 2]. Distances given in R_E are measured from the center of the Earth, so that 1 R_E refers to the Earth's surface.

[5] For example, in the South Atlantic Anomaly, the belts "dip" downward to nearly 200 km; otherwise, they are typically higher.

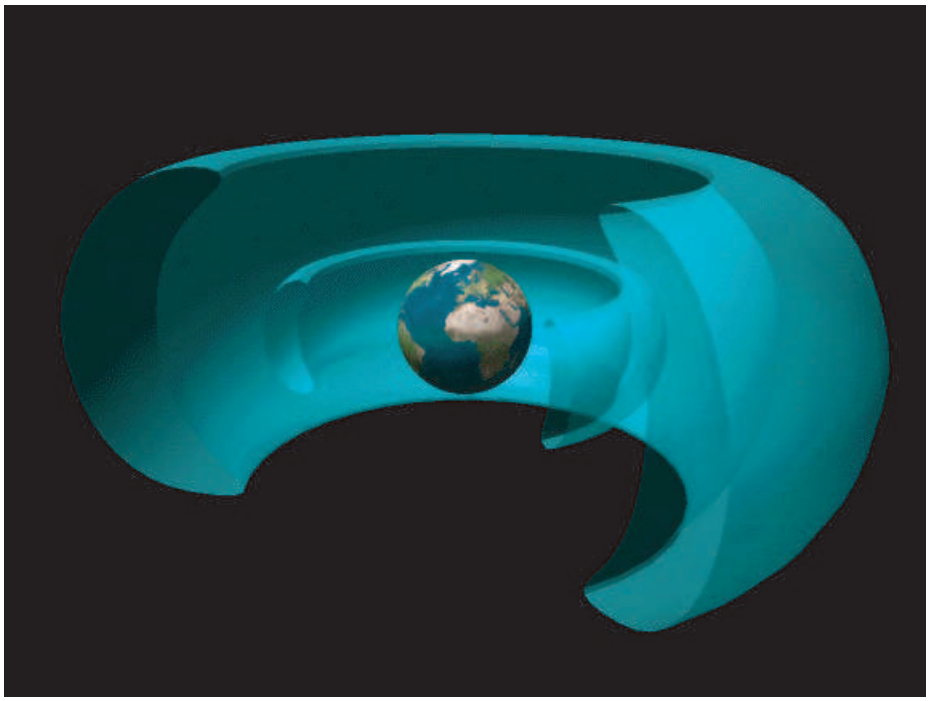

Figure 2.2: Cutaway depiction of the radiation belts that encircle the Earth. Both inner and outer radiation belts shown.

activity, but most are sourced by solar flares. Solar flares generate intense bursts of UV-rays and X-rays and also protons and heavy ions. These bursts are typically of less energy than galactic solar rays, but the increased radiation flux can dominate the natural radiation environment for several days [*Stassinopoulos and Raymond*, 1988].[6]

It should be noted that the above discussion describes only the natural radiation environment, and that at best briefly.[7] In practice, other sources of radiation are also present. For example, the spacecraft bus itself is a radiation source, emitting secondary radiation from interactions of its structures with the aforementioned primaries. The exact radiation environment encountered by a spacecraft is thus flight-specific and a strong function of both the orbit traversed and the type of bus flown.

[6]Indeed, solar flares play such a dominant role in the near-Earth radiation environment that they are often afforded their own category alongside trapped and transiting radiation. Certainly much near-Earth radiation environment modeling includes an extensive solar component.

[7]A more extensive presentation can be found in a variety of sources; a good starting point is *Stassinopoulos and Raymond* [1988].

2.1.2 Damage Mechanisms and Effects

A wide diversity of particles composes the space radiation environment, from atomic particles (energetic protons, electrons, and neutrons) to heavy ions (including iron and oxygen) to photons (including γ-rays and X-rays). These particles in turn interact with matter through a wide variety of processes. Many create products that become secondary radiation, leading to complex interaction profiles.[8] For semiconductor devices, though, two dominant radiation damage mechanisms arise: displacement damage and ionization damage.

Displacement damage is the displacement of atoms from their normal lattice positions by incident radiation. This process generates vacancy-interstitial pairs, introducing lattice defects. Indeed, if the incident particle imparts sufficient energy, the displaced atom can in turn displace other atoms, leading to entire defect clusters.[9] Lattice imperfections are particularly important in the silicon bulk where they can act as recombination and scattering centers that decrease minority carrier lifetime and mobility, and increase bulk resistivity [*Rasmussen*, 1988]. This process can have ramifications for bipolar transistors, for example, whose current gain can be compromised by displacement damage in the base regions.[10]

For MOS technologies, though, displacement damage is typically not significant compared to ionization damage. Ionization damage is the creation of free electron-hole pairs due to incident radiation. As an incident particle passes through matter, it loses energy. This energy can liberate electrons from their host atoms; the vacated position forms the hole. Such an ionizing strike, then, leaves a plasma track of electron-hole pairs in its wake. Once created, these charges migrate depending upon

[8] For a working introduction to radiation-matter interactions, the reader is recommended to *Srour and McGarrity* [1988]. Those interested in more complete descriptions, including derivations of the underlying particle physics, are recommended to *Evans* [1955].

[9] There are a variety of other lattice imperfections that can arise from displacement damage. Many of these structures involve interactions between the generated Frenkel pair and surrounding imperfections. Hence phenomena such as di-vacancies and di-interstitials have been observed. In addition, in doped regions, defect-impurity complexes, formed when the vacancy or interstitial is adjacent to a dopant atom, can arise [*Srour and McGarrity*, 1988].

[10] Shrinking of the base region, though, mitigates this effect, reducing its significance in modern bipolar technologies [*Kerns et al.*, 1988].

local conditions, including surrounding structures, materials, and electric fields.[11]

The remainder of this chapter, then, focuses on the effect of ionization damage on CMOS technologies. These effects are typically organized into two broad classes: single-event effects and total-dose effects. Single-event effects are immediate effects resulting from a small number of damage mechanisms, typically a single ionizing strike. In integrated circuits, these effects often lead to functionality failures. Total-dose effects are long-term effects resulting from accumulated radiation damage over an absorbed dose. These effects can change device behavior, often initially degrading circuit performance before eventually compromising circuit functionality when behavior changes become too large.

2.2 Single-Event Effects

For purposes of this discussion, single-event effects are immediate disruptions in circuit operation caused by a single ionizing strike. While initially transient, these effects can become permanent, and even cause physical damage to the integrated circuit. Over the years, much taxonomy has arisen to describe single-event effects. This dissertation adopts the taxa of "soft errors" and "hard errors" (a scheme also popular in terrestrial applications). Soft errors are those wherein, though a device or circuit is upset and its functionality disrupted, the device or circuit is not itself permanently damaged. In hard errors, on the other hand, the device or circuit is itself also permanently damaged.[12]

2.2.1 Soft Errors

Soft errors manifest as upsets within normal circuit operation. For digital circuits, soft errors translate to upset bits in logic, registers, or memory. If upsets in the first

[11] *Dodd and Massengill* [2003] present a good overview of the charge deposition and collection processes and their modeling.

[12] To some extent, the division is nebulous: a circuit may be partially damaged such that, under certain conditions, it is effectively permanently damaged, whereas under others, it operates properly. In this capacity, the effect mimics more a parameter shift reminiscent of a total-dose effect. Nonetheless, the various effect categories serve as useful organizational aids here.

are latched, or upsets in the latter two are uncorrected, this effect can result in data loss or inadvertent state changes. For analog circuits, soft errors translate to sudden glitches in the voltage or current at a node. The impact and permanence of these glitches on and in the overall circuit can be quite complicated, depending on the magnitude and duration of the glitch and the role of the node in the operation of the circuit. For analog circuits especially the definition of a soft error is thus highly circuit dependent.

In spaceborne applications, soft errors are typically characterized by their cross section versus linear energy transfer curves. Conceptually, the cross section is a measure of the rate of occurrence of a soft error. Cross section is tabulated by counting the number of errors and normalizing by the number of incident particles. This tabulation is repeated for varying incident particle energies, where the incident particle energy is characterized by its linear energy transfer (LET) to the target material. The resulting cross section versus LET curve usually increases rapidly for low LET before saturating at higher LET—see Figure 6.21 for two examples—and the knee of the curve is often called the LET threshold of the soft error. While this threshold is sometimes sufficient for describing the soft error rate, often the cross section versus LET curve is instead fit to a Weibull distribution for a more complete characterization.[13] The consequent Weibull parameters can then be combined with the expected radiation environment for a given orbit to predict the soft error rate for a particular mission.

Distinguishing what constitutes an ADC soft error, though, is an open question. This dissertation adopts the method proposed by *Turflinger and Davey* [1990]. Broadly, this method differentiates ADC output code errors as either noise errors or offset errors. Both are identified statistically: the former manifest as a broadening of the noise distribution of the ADC output, whereas the latter manifest as large code shifts well outside the noise distribution. Hence this scheme not only defines, but also categorizes, soft errors. The more appropriate error type can then be used to estimate the upset rate for a given application.

[13]The Weibull distribution arises from reliability analysis, where it well describes situations wherein a system fault can be modeled as a fault in the weakest link of many competing fault processes [*Petersen et al.*, 1992; *Tobias and Trindade*, 1986, pp. 70–72].

More complete explanations of the cross section versus LET curve, and the Turflinger method, can be found in Appendix A and Section 6.3.2.2, respectively.

2.2.2 Hard Errors

Hard errors are permanent damages to the devices or circuits of an integrated circuit. Typically, these errors are closely associated with particular device structures.

2.2.2.1 Gate Rupture

Radiation-induced gate rupture occurs when an ionizing strike causes the electric field across a transistor gate oxide to exceed a critical breakdown field E_C. Oxide breakdown occurs, resulting in a permanent short-circuit through the oxide.

This effect is mostly associated with vertical power MOSFETs [*Allenspach et al.*, 1996]. However, the increasing gate oxide electric field of modern CMOS technologies—arising from aggressive gate oxide thinning and not-as-aggressive supply voltage lowering—suggest gate rupture to be a possible concern for conventional CMOS technologies as well, especially those with ultra-thin gate oxides. Experiments, though, have strongly shown otherwise, including studies of SiO_2 gate oxides of 6-nm to 18-nm thickness [*Sexton et al.*, 1997], and studies of SiO_2 oxides and high-κ dielectrics of sub-5-nm thicknesses [*Massengill et al.*, 2001]. In both experiments it was found that E_C actually increases for thinner oxides. This increase mitigates the impact of the increased nominal electric field of thinner oxide devices. Practically, then, gate rupture requires either gate voltage biases in excess of the rated supply voltage of the process, or very high LET. The upshot is that gate rupture is not considered to be a concern in modern CMOS technologies.[14]

2.2.2.2 Latchup

Latchup is a catastrophic failure mechanism caused by activation of a parasitic PNPN structure. Such structures readily occur in CMOS integrated circuits: an example

[14]Indeed, *Sexton et al.* [1997] go so far as to argue that "advanced technologies should become less susceptible to SEGR [single-event gate rupture] as gate oxide thickness decreases".

is shown in Figure 2.3, which simply depicts adjacent NMOS and PMOS devices. As shown, the PNPN structure can be understood as a feedback loop composed of two bipolar devices: a vertical PNP Q_1 (composed of the P+ PMOS source/drain diffusion, N-well, and P-substrate) and a lateral NPN Q_2 (composed of the N+ NMOS source/drain diffusion, P-substrate, and N-well).[15] To better illustrate this feedback loop, the circuit is redrawn in Figure 2.4.[16] Normally Q_1 and Q_2 are off (V_b is drawn to the supply and V_a to ground) and the loop is inactive: the circuit draws no current. This state is called the blocking state. In a radiation environment, though, the circuit can be activated by current injection from an ionizing strike.[17] Once activated, the positive feedback of the loop[18] propels it towards the latchup state, which is a stable state wherein the circuit effectively shunts the supply to ground. The resulting high current condition, if left unchecked, can physically damage the devices and the chip; for example, it can vaporize on-chip metal traces [*Voldman*, 2007, pp. 2–3]. As the latchup state is stable, its correction typically requires removal of the power supply.

Latchup is typically described by considering the I-V curve of the circuit of Figure 2.4, as sketched in Figure 2.5 [*Johnston*, 1996]. Most analysis focuses on two points: the switching point (V_S, I_S) and the hold point (V_H, I_H). The switching point is the point where the circuit leaves the blocking state and corresponds to activation of the loop. Following the switching point, owing to the positive feedback, the circuit exhibits a negative resistance that propels it to the hold point, which is the stable latchup state itself: note the low-voltage, high-current condition.

To securely prevent latchup, the circuit should remain squarely in the blocking

[15] While other examples of PNPN structures exist, given its prevalence, the PNPN structure created by these two particular parasitic devices is the subject of much study.

[16] For simplicity, in Figure 2.4 V_p is driven to the supply and V_n to ground. In reality, latchup can occur for a variety of V_p and V_n conditions so long as V_p and V_n connect to high-current-capable sources: Figure 2.4 simply connects them to two readily available such sources.

[17] For example, if a strike generates electron-hole pairs near the N-well/P-substrate junction, the electrons are collected by the N-well and drive down V_b, while the holes are collected by the P-substrate and drive up V_a [*Voldman*, 2007, p. 179].

[18] To see that the feedback is positive, assume Q_1 and Q_2 are on and consider, say, an increase in the Q_2 collector current. The increase in Q_2 collector current in turn increases the current flowing into the base of Q_1, increasing Q_1's collector current. But the increase in Q_1 collector current in turn increases the current flowing into the base of Q_2, increasing Q_2's collector current. Thus the loop is closed in positive reinforcement.

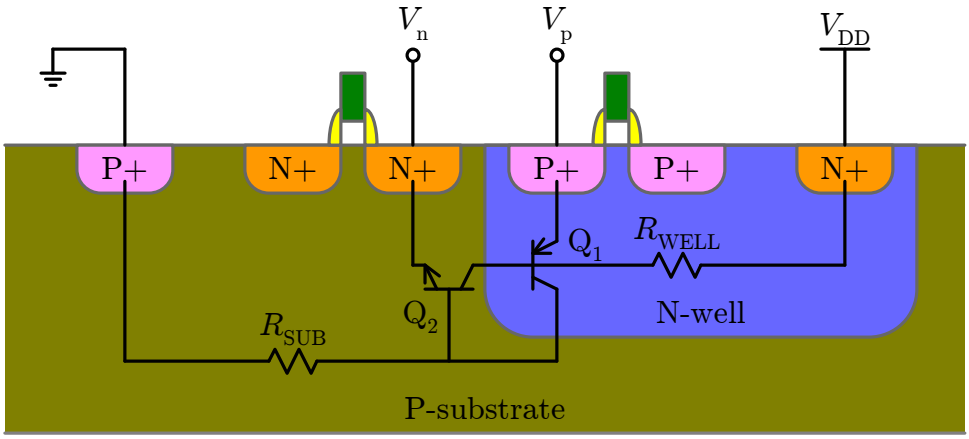

Figure 2.3: Silicon cross-section of adjacent NMOS and PMOS devices. Potential latchup circuit superposed.

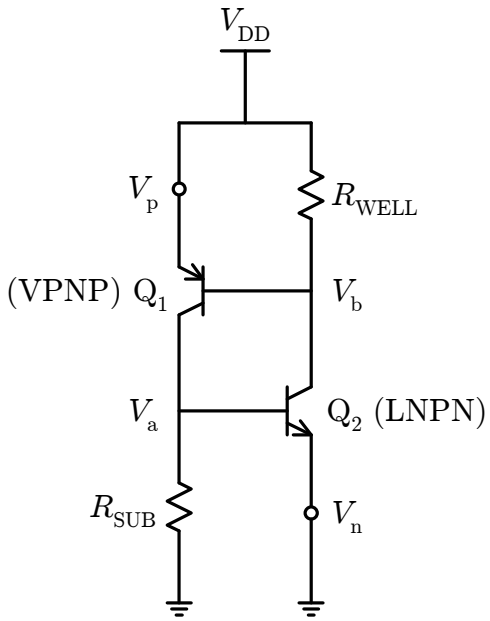

Figure 2.4: Latchup circuit for cross-section of Figure 2.3. Transistor types also indicated (VPNP denotes vertical PNP, LNPN denotes lateral NPN).

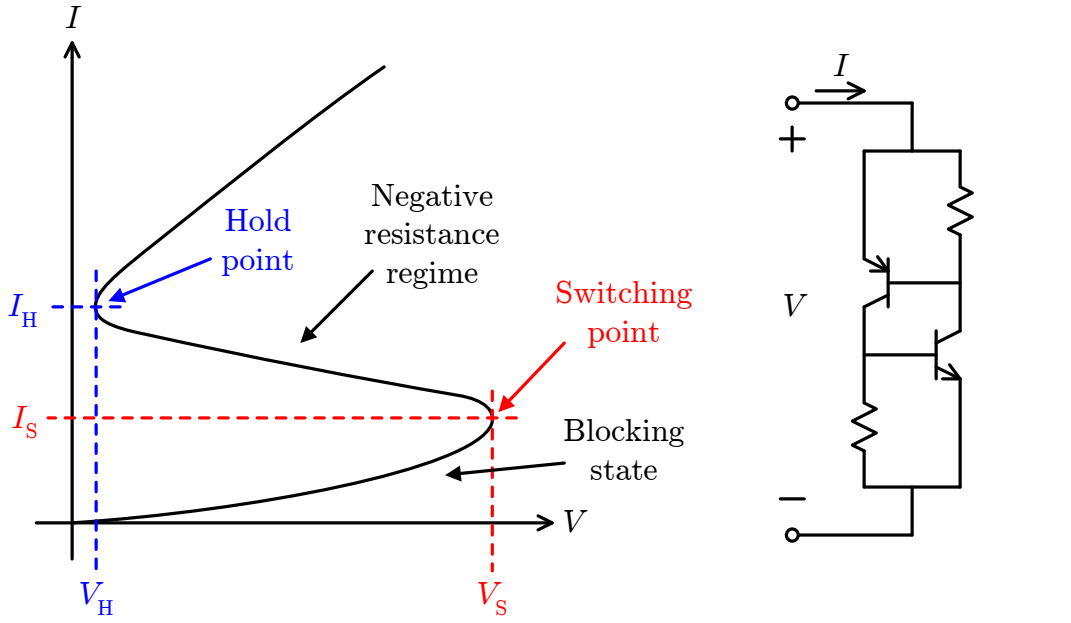

Figure 2.5: Sketch of I-V characteristic for latchup circuit of Figure 2.4. Latchup circuit redrawn at right for reference. Switching point labeled in red, hold point in blue.

state [*Troutman*, 1986, pp. 204–205].[19] From a stability analysis, it can be shown that such is the case so long as [*Troutman*, 1986, p. 72]:

$$\frac{\alpha_{\mathrm{fs,Q_1}}}{1+\dfrac{R_{\mathrm{e,Q_1}}}{R_{\mathrm{WELL}}}}+\frac{\alpha_{\mathrm{fs,Q_2}}}{1+\dfrac{R_{\mathrm{e,Q_2}}}{R_{\mathrm{SUB}}}}<1 \tag{2.1}$$

where $\alpha_{\mathrm{fs,Q_1}}$ and $\alpha_{\mathrm{fs,Q_2}}$ are the small-signal current gains [*Troutman*, 1986, p. 61]:

$$\alpha_{\mathrm{fs,Q_1}}=\frac{\mathrm{d}I_{\mathrm{c,Q_1}}}{\mathrm{d}I_{\mathrm{e,Q_1}}} \quad , \quad \alpha_{\mathrm{fs,Q_2}}=\frac{\mathrm{d}I_{\mathrm{c,Q_2}}}{\mathrm{d}I_{\mathrm{e,Q_2}}} \tag{2.2}$$

[19] This strategy seeks to ensure that the circuit never reaches the switching point. An alternate proposition is to prevent latchup by ensuring that the circuit never reaches the hold point. However, the switching-point requirement is favored here; this decision is discussed in more detail in Section 5.2.1.

and R_{e,Q_1} and R_{e,Q_2} the small-signal emitter resistances [*Troutman*, 1986, p. 64]:

$$R_{e,Q_1} = \frac{dV_{be,Q_1}}{dI_{e,Q_1}} \quad , \quad R_{e,Q_2} = \frac{dV_{be,Q_2}}{dI_{e,Q_2}} \tag{2.3}$$

Conversely, then, under the condition:

$$\frac{\alpha_{fs,Q_1}}{1 + \dfrac{R_{e,Q_1}}{R_{WELL}}} + \frac{\alpha_{fs,Q_2}}{1 + \dfrac{R_{e,Q_2}}{R_{SUB}}} \geq 1 \tag{2.4}$$

the loop activates and the circuit proceed towards latchup. Note that equality in Equation (2.4) corresponds to the switching point.

For most CMOS technologies, α_{fs,Q_1} and α_{fs,Q_2} are close to their large signal counterparts α_{f,Q_1} $(= I_{c,Q_1}/I_{e,Q_1})$ and α_{f,Q_2} $(= I_{c,Q_2}/I_{e,Q_2})$, respectively, in the blocking state, which typically sum to less than unity [*Troutman*, 1986, p. 68]. Instead, it is usually the denominators of Equations (2.1) and (2.4) that destabilize the blocking state. Note that R_{WELL} and R_{SUB} act as shunting resistors that siphon current away from the loop, preventing its activation. However, should the resistance ratios $R_{e,Q_1}/R_{WELL}$ and $R_{e,Q_2}/R_{SUB}$ become sufficiently small, the loop can consume sufficient current and activate. Generally, then, to prevent latchup, the resistance ratios should be as large as possible; from a circuit design perspective, this condition can be realized by reducing R_{WELL} and R_{SUB}, for example.

Finally, two general latchup trends are notable. First, as the switching point requires a voltage elevation (in addition to the current elevation), increasing the supply voltage increases the chance of latchup. Second, latchup is strongly temperature dependent. In terms of LET, at higher temperatures latchup exhibits both a lower LET threshold and a larger saturation cross section than at lower temperatures [*Johnston*, 1996]. The dominant causes are a decrease in the base-emitter junction forward voltage[20] and an increase in R_{WELL} with increasing

[20] That is, in the V_{be} applied to the base-emitter junction to achieve a given collector current I_c. From basic transistor theory, it is known that this V_{be} decreases by about 2 mV for each 1°C increase in temperature [*Sedra and Smith*, 1998, p. 239]. Hence, at lower temperatures, the same V_{be} voltage achieves much greater I_c; intuitively, the transistor can be thought of as being easier to turn on.

temperatures [*Johnston*, 1996]. As latchup is a catastrophic failure mechanism capable of not only destroying the chip, but also collapsing the entire power supply system, it is thus often tested under elevated temperature and supply conditions.

2.3 Total-Dose Effects

Whereas single-event effects occur from a single ionizing strike, total-dose effects result from an accumulated radiation dose. Broadly, radiation damage causes changes in MOSFET behavior, affecting circuit performance and possibly functionality.

Ionization damage to insulators is the dominant cause of total-dose effects in MOSFETs. A MOSFET contains at least two radiation-susceptible insulators: the gate oxide and the field oxide.[21] Both are typically SiO_2 in submicron processes.[22] The gate oxide is the thin oxide that forms the MOS structure of the nominal transistor. The field oxide is the thicker interdevice oxide that isolates devices one from another. The field oxide is often implemented as a local oxidation of silicon (LOCOS) oxide or, more recently, as a shallow trench isolation (STI) oxide. Note that the field oxide borders the transistor channel. In this region—often characterized as the polysilicon gate overhang—the gate and field oxides abut (as can be seen in Figure 2.13) and radiation damage to the field oxide can influence nominal transistor behavior.

Naturally, radiation damage is indiscriminate: it also alters the silicon substrate. As previously seen, this damage is important in single-event effects. However, ionization damage in the substrate does not accumulate—the created charges migrate to the supplies through relatively low resistance paths (at least compared to the oxides)—and hence is not significant in total-dose effects of CMOS technologies. This discussion thus focuses on ionization damage to the oxides.

[21] Some technologies sport more than these two oxides, such as silicon-on-insulator (SOI) technologies which employ a third, buried oxide. In these technologies, all three oxides can be important in MOSFET radiation response; see footnote 4 of Chapter 5. For the discussions of this chapter, only total-dose effects stemming from gate and field oxide radiation damage are considered, although many of the principles can be applied to the buried oxide as well.

[22] Deep-submicron processes often use high-κ dielectrics, instead of standard SiO_2, for the gate oxide. The discussions of this chapter only address SiO_2 oxides, although many of the principles expounded are applicable to high-κ dielectrics as well.

2.3.1 Oxide Ionization Damage

To understand oxide ionization damage, consider an ionizing strike through the oxide of a MOS structure as depicted in Figure 2.6.[23] Doping polarities befitting an NMOS device are shown, although the physics described here are readily applied to PMOS devices as well. Note that a bias is placed across the MOS structure resulting in an electric field across the oxide.

The ionizing strike creates a trail of electron-hole pairs, illustrated as process (1) in Figure 2.6. Being formed in close proximity, many of the electrons and holes recombine, annihilating each other. The remaining charges undergo transport under the influence of the applied electric field: given the bias of Figure 2.6, electrons ascend to the gate, while holes descend to the substrate. Owing to their higher mobility, many of the electrons that escape recombination are rapidly swept out of the oxide. What remains, then, is a net yield of less mobile holes, illustrated as process (2) in Figure 2.6. The exact hole population depends strongly on the type and energy of the incident radiation (which determines the initial density of pairs) and the magnitude of the electric field (which separates the pairs, reducing recombination) [*Oldham and McLean*, 2003].

The net hole population of process (2) forms on picosecond time scales at room temperature [*Oldham and McLean*, 2003]. The population then transports subject to the applied electric field, illustrated as process (3) in Figure 2.6. The parameters of this transport depend on the field, temperature, oxide thickness, and oxide quality [*Oldham and McLean*, 2003]. The transport is often modeled mathematically by a continuous-time random walk; the underlying transport mechanism is believed to be small polaron hopping[24] between localized shallow trap states [*Oldham and*

[23] A more thorough presentation of the processes described in this section can be found in *Oldham and McLean* [2003]; indeed, the paper is cited extensively in these discussions. Readers desirous of yet more rigor can consult *Oldham* [1999], a complete text on the subject.

[24] As a hole moves, it distorts the surrounding medium. If this distortion is strong enough, it creates a region of low system potential, which in turn effectively traps (i.e., localizes) the hole, decreasing its mobility (or equivalently, increasing its effective mass). The hole is thus said to be self-trapping. As the hole and lattice distortion move in accord, the combination is regarded as a single particle called a small polaron.

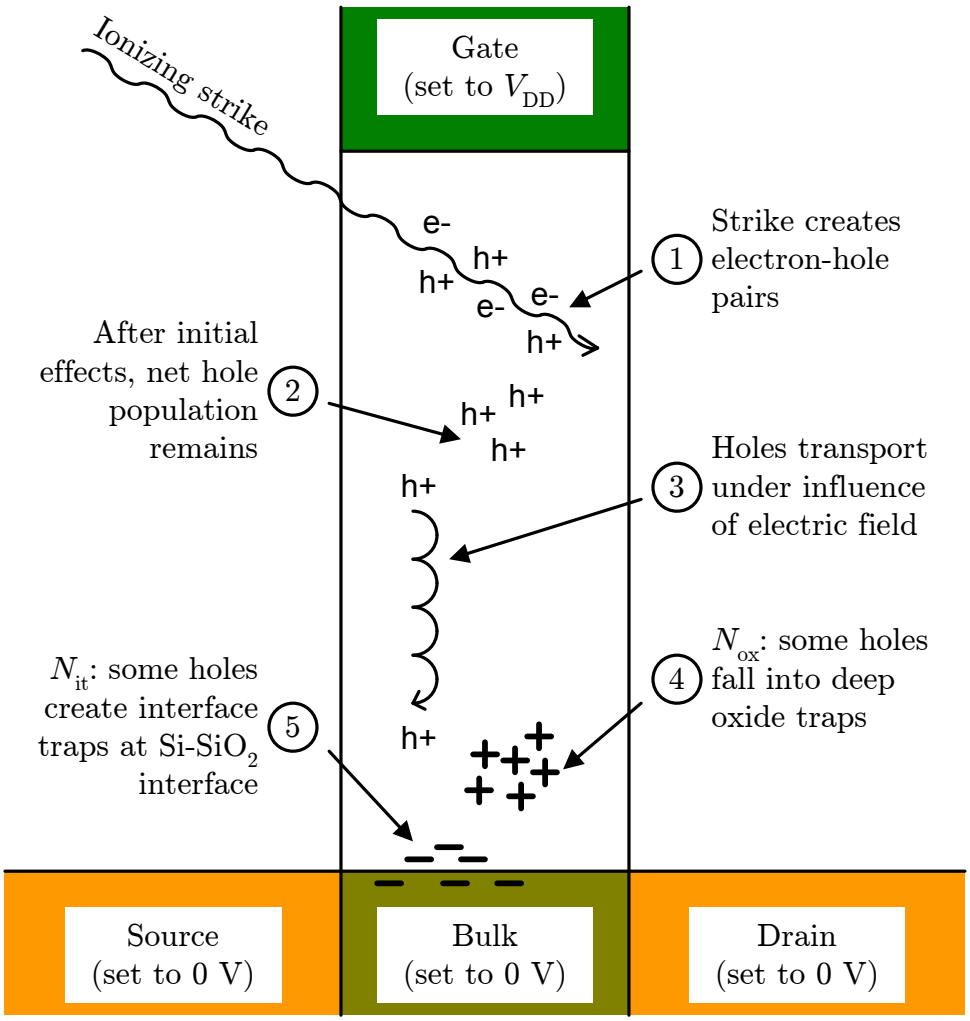

Figure 2.6: Ionization damage processes within the oxide of an MOS structure. Biased NMOS device shown. Drawing not to scale.

McLean, 2003].[25] At room temperature for oxide thicknesses as expected of modern technologies, the hole transport typically completes within one second of the irradiating strike.

As the holes reach the oxide edge, many escape into the channel, annihilated by the available electron population. Some holes, however, become more permanently trapped, falling into deep, long-lived trap states near the Si-SiO_2 interface, illustrated as process (4) in Figure 2.6. These deep hole traps are associated with defects near the interface,[26] and form a population of oxide trapped charge N_{ot}. Oxide trapped charge can persist for hours to years, with a complex dependence on temperature and field. The trapped charge is removed by annealing, either through thermal processes or by tunneling from free electrons in the substrate; near room temperature, tunneling dominates [*Oldham and McLean*, 2003]. It has been shown, though, that some of this annealing is just compensation and not elimination: application of electric fields after annealing can reactivate the oxide trapped charge [*Schwank et al.*, 1984].

In addition to falling into deep traps, transporting holes can also give rise to increased populations of interface traps at the Si-SiO_2 interface, illustrated as process (5) in Figure 2.6. These traps are typically acceptors above midgap and donors below. The existence of interface traps is a well-known phenomenon in MOSFET manufacturing attributed to dangling Si bonds at the Si-SiO_2 interface.[27] These traps must be passivated lest MOSFET functionality be compromised. Passivation is typically accomplished by bonding hydrogen to the sites: the added Si-H bond neutralizes the trap.[28] Under irradiation, though, these sites can become

[25] Given a polaron in a trap site, if the potential of a nearby shallow trap is sufficiently lowered through thermal activation, the polaron can tunnel to the new location, enabling the polaron to move through the lattice. The probability of such movement is determined by the probability of creating a suitable site (which is strongly dictated by temperature) and the wave function overlap of the two sites (which dictates whether tunneling occurs) [*Anelli*, 2000, pp. 42–43].

[26] Specifically, nominally Si atoms in SiO_2 are separated by an O between them. However, near the interface, oxygen vacancies (or alternately, Si excesses) can form, resulting in weak Si-Si bonds, sans O. This bond can be broken by a hole; in doing so, the hole becomes trapped [*Oldham and McLean*, 2003].

[27] Specifically, in the insulator and semiconductor Si atoms are typically bonded to four other atoms: four O atoms in the SiO_2 insulator, four other Si atoms in the substrate. However, near the Si-SiO_2 interface, trivalently bonded Si atoms can occur. The fourth position, then, is unbonded, and forms the trap [*Plummer et al.*, 2000, p. 353].

[28] Indeed, this passivation is essential for consistent manufacturing: *Pierret* [1990, p. 111] goes so

unpassivated, reactivating the electrically active defect and affecting transistor operation.[29] The resulting interface trap population N_{it} depends on many factors, including electric field, temperature, oxide thickness, and oxide quality [*Oldham and McLean*, 2003]. The traps themselves have proven stable at room temperature over the time scales of most experiments, although some annealing has been observed at elevated temperatures [*Oldham*, 1999, pp. 157–158]. Generally, then, N_{it} tends to persist much longer than N_{ot}.

These two populations—the oxide trapped charge N_{ot} and the interface traps N_{it}—accumulate over the course of multiple strikes and account for the vast majority of the total-dose changes in transistor behavior.[30]

2.3.2 Damage Dependencies

The previous discussion makes clear that radiation-induced N_{ot} and N_{it} depend on a variety of factors. Before discussing the impact of N_{ot} and N_{it} on MOSFET behavior, though, it is worthwhile to address a few of these factors in more detail.

First, radiation damage depends heavily on oxide thickness: in general, thinner oxides accumulate less radiation damage. For N_{ot}, if the oxide is thinner, a greater volume is available for electron tunneling from both the substrate and gate to mitigate the oxide trapped charge [*Oldham*, 1999, pp. 99–100]. For N_{it}, measurements have shown that thinner oxides yield reduced interface trap populations [*Oldham*, 1999,

far as to say that, without it, "MOS devices would merely be a laboratory curiosity".

[29] The exact process by which radiation gives rise to interface traps has long been debated. The current belief is that transporting holes free hydrogen trapped in the SiO_2 oxide. This hydrogen, in the form of protons, transports subject to the electric field. Should these protons encounter passivated Si-H bonds, they can react with the H to form more stable H_2, unpassivating the Si and restoring the trap. When this reaction occurs near the Si-SiO_2 interface, an interface trap is formed. This model explains the vast majority of interface traps, although some second-order effects remain unaccounted for [*Oldham and McLean*, 2003].

[30] In mathematical expressions, the oxide trapped charge N_{ot} and interface traps N_{it} are typically assumed to be volumetric densities with units of particles/cm^3. However, their definitions can vary depending on context, for example, N_{it} is occasionally assumed to be an areal density with units of particles/cm^2. Readers are thus cautioned to check the definition assumed by any particular reference. For this dissertation, whenever N_{ot} and N_{it} are used in mathematical expressions, the assumed definition and units are stated explicitly, otherwise, the terms refer to the populations in general.

p. 161]. From a total-dose perspective, submicron CMOS technologies thus benefit from the ever-thinning gate oxides of CMOS scaling.[31] However, while the thin gate oxide may be fairly radiation hard, the thicker field oxide may not be, and field oxide damage may dominate MOSFET radiation response.

Second, radiation damage is a strong function of both the direction and magnitude of the electric field imposed across the oxide. For example, as previously mentioned, the larger the electric field magnitude, the greater the separation of electrons and holes following the initial strike: this separation reduces recombination, and hence larger electric field magnitudes increase the net hole yield of process (2) of Figure 2.6. And the electric field direction must be oriented to direct holes towards the Si-SiO_2 interface to give rise to N_{it}. Hence, the accumulated radiation damage depends heavily upon device bias. Most presentations of MOSFET radiation effects thus assume "worst-case bias" during irradiation. For an NMOS device, worst-case bias means the gate voltage V_G is set at the rail voltage V_{DD}, and the source, drain, and bulk voltages V_S, V_D, and V_B, respectively, are set to ground. For a PMOS device, V_G is set to ground, and V_S, V_D, and V_B set to V_{DD}.

Finally, it should be stressed that radiation damage is strongly dependent on the manufacturing process. Processes offering similar commercial performance may differ dramatically in their radiation response owing to subtle differences in process steps. Furthermore, as CMOS technologies advance, different radiation damage processes may come to dominate. For example, experiments with ultra-thin gate oxides of only a few nanometer thickness have shown increases in gate-to-channel leakage current due to electron tunneling through radiation-induced neutral trap sites in the gate oxide [*Ceschia et al.*, 1998],[32] an increase not predicted by the physics described so far. Thus, while the discussions herein address the dominant total-dose effects expected of a generic submicron CMOS technology, designers should be aware that the actual effects manifested by their particular process may vary.

[31] Aggressive oxide thickness scaling is less prevalent in the deep-submicron regime, where the use of high-κ dielectrics enables gate capacitance scaling without necessitating thinner gate oxides (and without incurring the consequent gate leakage from increased electron tunneling).

[32] In commercial settings this leakage current is called stress-induced leakage current: radiation exacerbates the effect. In contrast to N_{ot} and N_{it}, this effect is worst when there is no electric field across the oxide [*Ceschia et al.*, 1998].

2.3.3 Transistor Effects

The following subsections consider total-dose changes to MOSFET behavior along select parameters. Both qualitative descriptions and quantitative results are presented. Regarding the qualitative descriptions, it should be noted that the effects described are often true of N_{ot} and N_{it} populations in general, including any baseline populations that occur as a result of the manufacturing process itself. From a circuit design perspective, though, it is the effect of the contributions to N_{ot} and N_{it} caused by radiation, and the consequent changes in device behavior, that are of most interest. The latter, then, are referred to as radiation-induced N_{ot} and N_{it}, or N_{it} and N_{it} damage, when distinguished from the general N_{ot} and N_{it} populations.

The quantitative results presented are primarily composed of radiation testing results of the National Semiconductor Corporation BiCMOS8iED technology used to implement the SVADC-1. This process—made available by a collaboration agreement between National Semiconductor Corporation and Stanford University—is a 0.25-μm CMOS technology that includes both low-voltage 2.5 V (LV) and high-voltage 3.3 V (HV) devices with gate oxide thicknesses of $t_{ox} = 5.3$ nm and $t_{ox} = 6.4$ nm, respectively. A more complete description of this process is given in Section 5.1, but suffice to say here that BiCMOS8iED is a single-well process fabricated on a non-epitaxial substrate with STI interdevice separation oxide. To assess its radiation response, test transistors fabricated in this process[33] are total-dose radiation tested—by exposure to γ-rays from a ^{60}Co source—to a maximum dose of 2 Mrad(Si). This testing was conducted jointly by Stanford University and The Aerospace Corporation: further details are available in Appendix F. Notably, unless otherwise noted, for all the measured results presented here the BiCMOS8iED test devices are irradiated under worst-case bias.

To present a broader perspective of radiation response, in many cases the measured BiCMOS8iED results are augmented the published results of *Manghisoni et al.* [2002], which describes the radiation response of two commercial CMOS processes manufactured by STMicroelectronics. The first is a 0.18-μm process with a gate oxide thickness of $t_{ox} = 4$ nm, the second a 0.35-μm process with a gate oxide thickness of

[33] These test devices are included at the bottom of the SVADC-1 chip, as can be seen in the chip micrograph of Figure 5.40.

$t_{ox} = 7.2$ nm. Notably, the 0.18-μm devices are irradiated under both worst-case bias (referred to as Group A) and under a more moderate "analog operating bias" (referred to as Group B), and the 0.35-μm devices are irradiated under an analog operating bias.[34] All devices are irradiated via γ-rays from a ^{60}Co source.

It is conventional to express accumulated radiation dose in units of rad (or Gray, where 1 Gy = 100 rad). The rad unit is material dependent, but rad(Si) is typically used.[35] In addition, the results often present an anneal dose step: following the final dose, devices are usually allowed to anneal in an unirradiated state for a set duration and then re-measured. Insofar as the anneal mitigates radiation damage, this dataset can provide insight into the underlying damage mechanisms.

Further information regarding the terms and methods of total-dose testing can be found in Appendix A.

2.3.3.1 Threshold Voltage

N_{ot} and N_{it} affect the threshold voltage V_T of both NMOS and PMOS devices. N_{ot}, being positive, works in concert with an ascending gate voltage in an NMOS device, and in opposition against a decreasing gate voltage in a PMOS device. Hence, N_{ot} causes a decrease in $|V_{T,N}|$, and an increase in $|V_{T,P}|$, the absolute value of the threshold voltage of an NMOS and PMOS device, respectively.

The impact of N_{it} is more complicated. Radiation-induced interface traps tend to be acceptors above the midgap E_i and donors below. An acceptor trap state is neutral when above the Fermi level E_F and becomes negative (by "accepting" an electron) when below, whereas a donor is neutral below E_F and becomes positive (by "donating" an electron) when above. Consider, then, a rising E_F as shown in Figure 2.7 for an NMOS channel transitioning from flatband (left) to inversion (right). In flatband, acceptor traps are neutral whereas donor traps are positive or neutral. The

[34] Specifically, for the three classes of devices, the NMOS/PMOS irradiation biases are:
- 0.18-μm, Group A: $V_{GS} = 1.6$ V/0 V, $V_{DS} = 0$ V, $V_{BS} = 0$ V
- 0.18-μm, Group B: $V_{GS} = 0.4$ V/-0.4 V, $V_{DS} = 0.8$ V/-0.8 V, $V_{BS} = 0$ V
- 0.35-μm: $V_{GS} = 0.6$ V/-0.65 V, $V_{DS} = 1.5$ V/-1.5 V, $V_{BS} = 0$ V

[35] On occasion, the unit of rad(SiO_2) is used instead. This practice is typically found in references focusing on the underlying particle physics of radiation-damage in SiO_2 oxides (as an example, see *Oldham* [1999], which presents a survey of results, some quoted in rad(Si) and some in rad(SiO_2)).

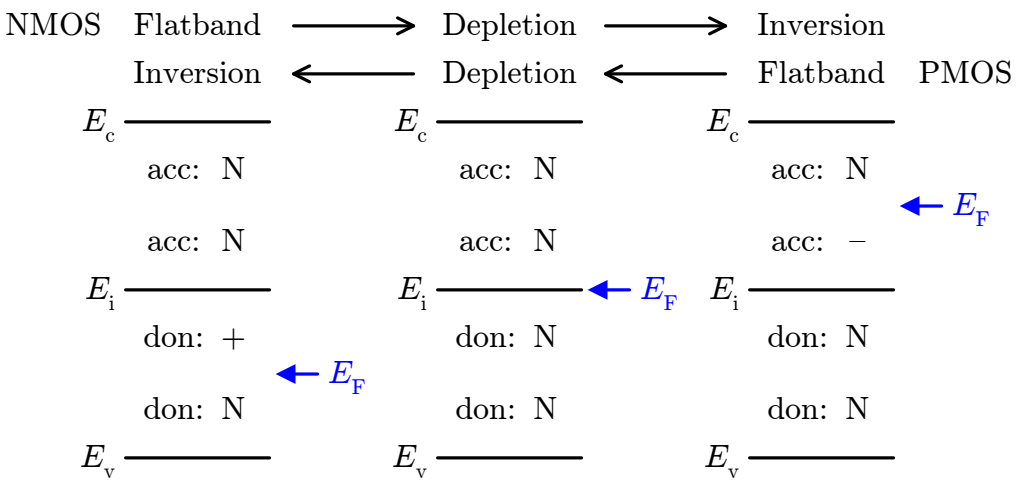

Figure 2.7: Interface trap states for varying E_F levels for radiation-induced acceptor (acc) and donor (don) traps. Trap states indicated by N (neutral), or + or − for positive and negative states, respectively. Corresponding NMOS and PMOS channel state indicated above.

positive donors become neutral as E_F rises so that all traps are neutral by depletion. As E_F continues to rise, the channel moves to inversion and acceptors start becoming negative. Hence, from flatband to depletion there is a loss of positive charge, and from depletion to inversion there is a gain of negative charge. Both of these tendencies act in opposition to the increasing gate voltage working to invert the channel. Hence, radiation-induced interface traps increase $|V_{T,N}|$. Traversing Figure 2.7 right to left yields a similar argument for PMOS devices: radiation-induced interface traps also increase $|V_{T,P}|$.[36]

In total, then, for an NMOS device, $|V_{T,N}|$ may decrease or increase with dose, depending on the relative strengths of radiation-induced N_{ot} and N_{it}, respectively. In contrast, for a PMOS device, $|V_{T,P}|$ only increases with dose.

Figure 2.8 shows the radiation evolution of $V_{T,N}$ for BiCMOS8iED NMOS devices of varying type and size irradiated under worst-case bias. Generally, $V_{T,N}$ initially

[36]While the vast majority of radiation-induced interface traps are acceptor above E_i and donor below, small populations of donors above E_i and acceptors below have been reported [*McWhorter et al.*, 1988]. However, an extension of the preceding argument shows that these traps qualitatively contribute to threshold voltage in the same fashion as their more plentiful counterparts.

decreases (suggesting strong N_{ot} damage) before recovering at high doses (suggesting relatively stronger N_{it} damage) and/or in anneal (due to annealing of N_{ot} damage). Overall, $V_{T,N}$ changes are small, being at most 5% of the original value, suggesting that the thin gate oxide of BiCMOS8iED is fairly radiation hard. This result is expected: the conventional wisdom is that "oxides less than 10 nm in thickness show almost no radiation induced threshold voltage shift" [*Hughes and Benedetto*, 2003]. Finally, note that the results of Figure 2.8 are bias dependent: for comparison, Figure 2.9 shows the radiation evolution of $V_{T,N}$ for the same devices biased with non-zero V_{DS} during irradiation (specifically, $V_{DS} = 0$ V in Figure 2.8 versus $V_{DS} = 1.75$ V in Figure 2.9). To aid comparisons, the axes in both figures are the same. It is clear that applying different biases during irradiation results in different $V_{T,N}$ shifts, with the worst-case bias indeed displaying the greatest shifts. The existence of such bias dependence can influence circuit design, for example, in the construction of current mirrors [*Edwards et al.*, 1999].

To place these $V_{T,N}$ shifts in context, Figures 2.10 and 2.11 show the radiation evolution of $V_{T,N}$ and $V_{T,P}$, respectively, for various devices from commercial 0.18-μm and 0.35-μm processes [*Manghisoni et al.*, 2002]. Regarding the NMOS devices of Figure 2.10, note that the 0.18-μm Group A device shows an increase in $V_{T,N}$ with dose, highlighting that $|V_{T,N}|$ may increase or decrease, even initially, with dose. And regarding the PMOS devices of Figure 2.11, note that for all devices $|V_{T,P}|$ only increases with dose ($V_{T,P}$ is defined negative here). Overall, the 0.18-μm and 0.35-μm processes demonstrate $V_{T,N}$ changes by 5% to 14%, and $V_{T,P}$ by 10% to 19%, respectively. As expected, the 0.18-μm process changes less, owing to its thinner gate oxide. Overall, the V_T changes demonstrated by the 0.18-μm, 0.25-μm, and 0.35-μm processes, though small, can be significant in analog design, for example, they can contribute to analog biasing inaccuracy or increased offsets.[37]

[37] In light of the V_T shifts shown in these processes, the conventional wisdom of *Hughes and Benedetto* [2003] quoted previously should perhaps be qualified. With micron technologies, V_T shifts were much greater: indeed, for NMOS devices, the shift would sometimes be sufficient to change the sign of $V_{T,N}$, effectively rendering an enhancement-mode device depletion-mode [*Dressendorfer*, 1989a, p. 267]. Certainly in that light, changes of even 20% are small.

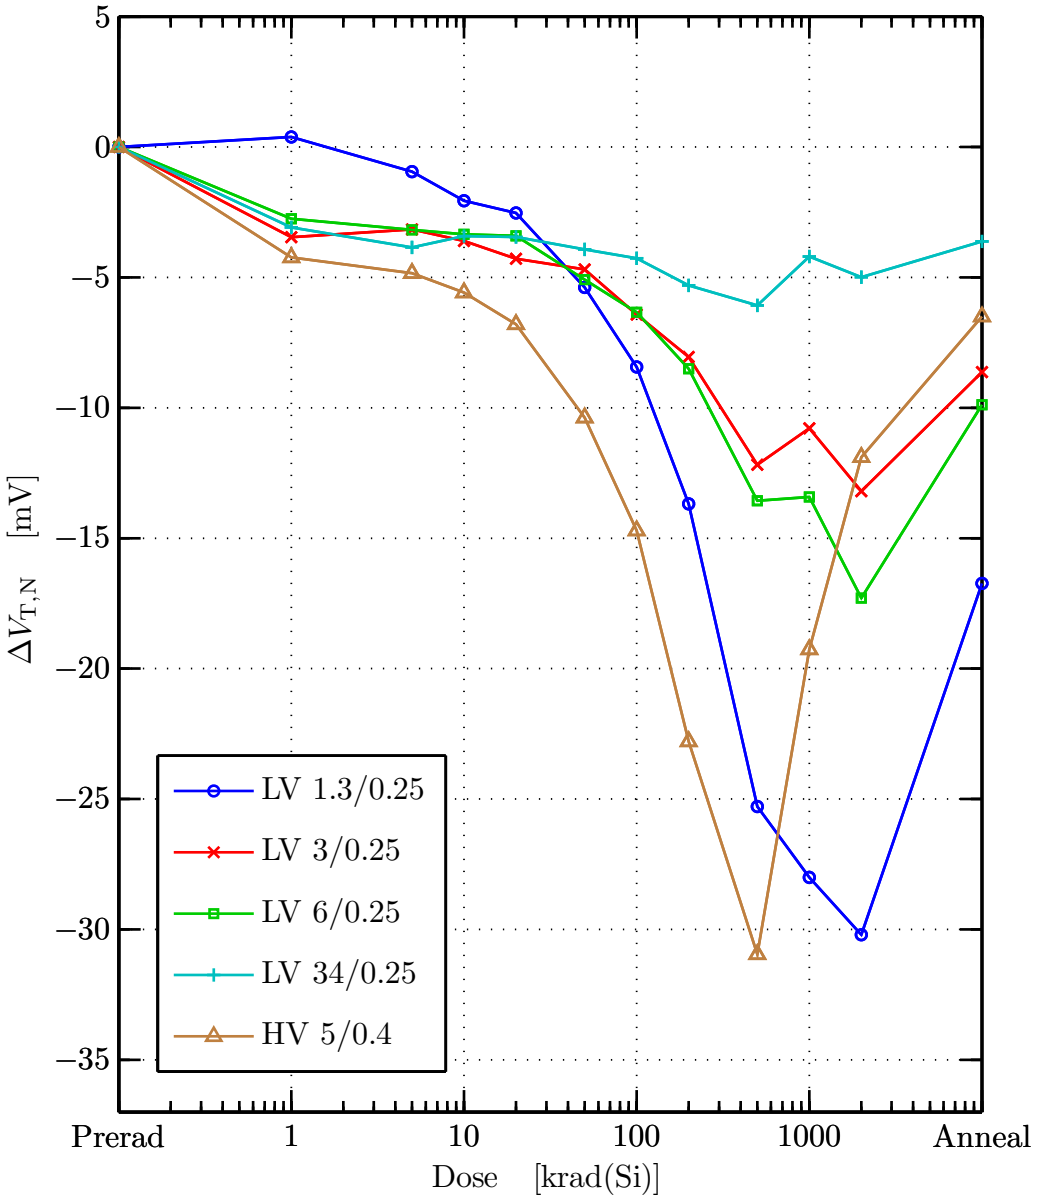

Figure 2.8: Measured threshold voltage change $\Delta V_{T,N}$ versus total dose for multiple BiCMOS8iED NMOS devices irradiated under worst-case bias. Device type and W/L given in legend. Change from pre-irradiation $V_{T,N}$ shown; pre-irradiation $V_{T,N}$ (in order as presented top-to-bottom in legend): 0.58 V, 0.60 V, 0.60 V, 0.62 V, 0.69 V. All devices irradiated under biases of $V_G = 2.5$ V (LV)/3.3 V (HV) and $V_S = V_D = V_B = 0$ V. Markers indicate measured data points.

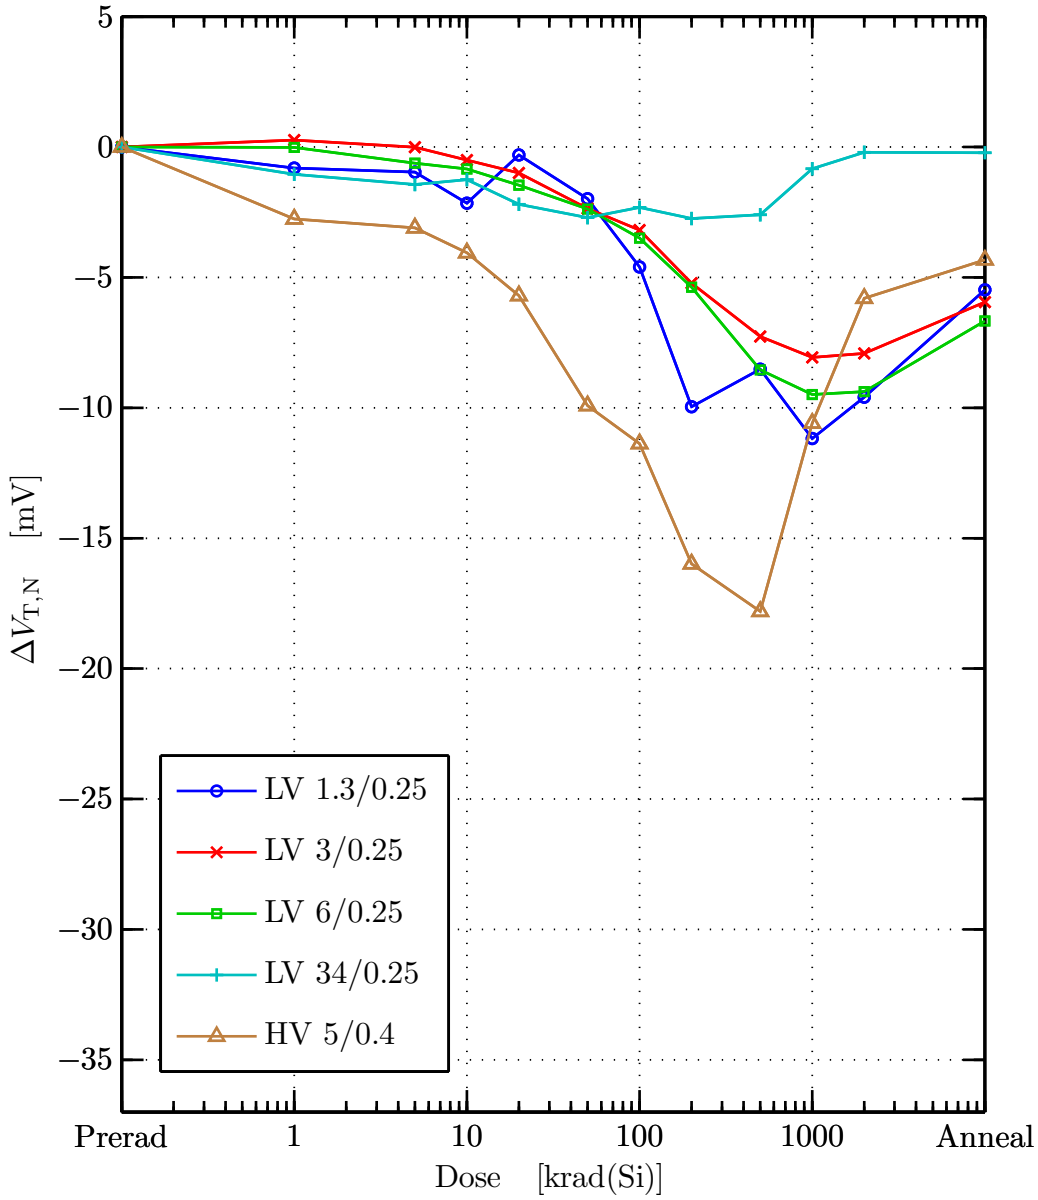

Figure 2.9: Measured threshold voltage change $\Delta V_{\mathrm{T,N}}$ versus total dose for multiple BiCMOS8iED NMOS devices irradiated under non-worst-case bias. Device type and W/L given in legend. Change from pre-irradiation $V_{\mathrm{T,N}}$ shown; pre-irradiation $V_{\mathrm{T,N}}$ (in order as presented top-to-bottom in legend): 0.66 V, 0.64 V, 0.66 V, 0.59 V, 0.66 V. All devices irradiated under biases of $V_{\mathrm{G}} = 2.5$ V (LV)/3.3 V (HV), $V_{\mathrm{S}} = V_{\mathrm{B}} = 0$ V, and $V_{\mathrm{D}} = 1.75$ V. Markers indicate measured data points.

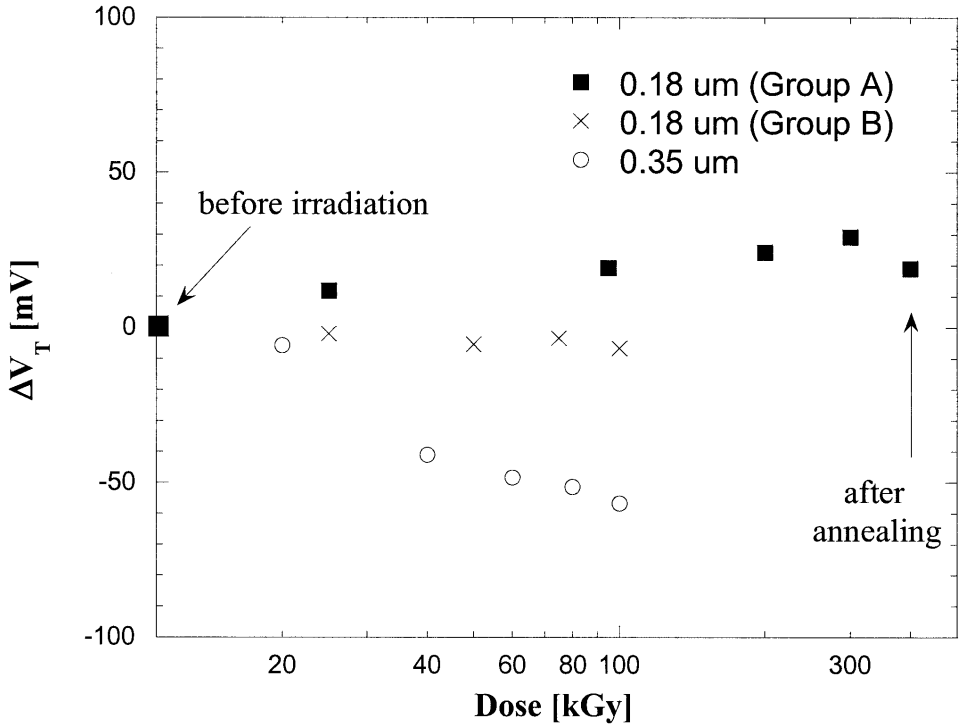

Figure 2.10: Reported threshold voltage change $\Delta V_{\mathrm{T,N}}$ versus total dose for NMOS devices from commercial 0.18-μm and 0.35-μm processes. Change against pre-irradiation $V_{\mathrm{T,N}}$ shown. Pre-irradiation $V_{\mathrm{T,N}}$ is 0.44 V for 0.18-μm process, and is 0.61 V for 0.35-μm process. Dose given assuming Si as incident material. Reproduced in total from *Manghisoni et al.* [2002].

2.3.3.2 Leakage Current

Radiation can activate significant leakage current flow between the source and drain of a nominally "off" NMOS device, that is, one whose V_{G} is set to 0 V.[38]

The impact of radiation on the drain-to-source leakage current I_{LEAK} is illustrated in Figure 2.12, which depicts the drain current versus gate-to-source voltage characteristic, or I_{D}-V_{GS} curve, for a MOSFET. In such plots, I_{D} is typically plotted on a logarithmic scale. The leakage current I_{LEAK} is defined as I_{D} when $V_{\mathrm{GS}} = 0$ V.

[38] A variety of other leakage currents can also activated, including gate leakage and interdevice leakage. The former is described briefly in Section 2.3.2. The latter operates through principles similar to that discussed in this section, but acting through a parasitic device composed of the polysilicon connecting different devices over the field oxide (such as a polysilicon line connecting the NMOS and PMOS gate terminals in an inverter) [*Shaneyfelt et al.*, 1998].

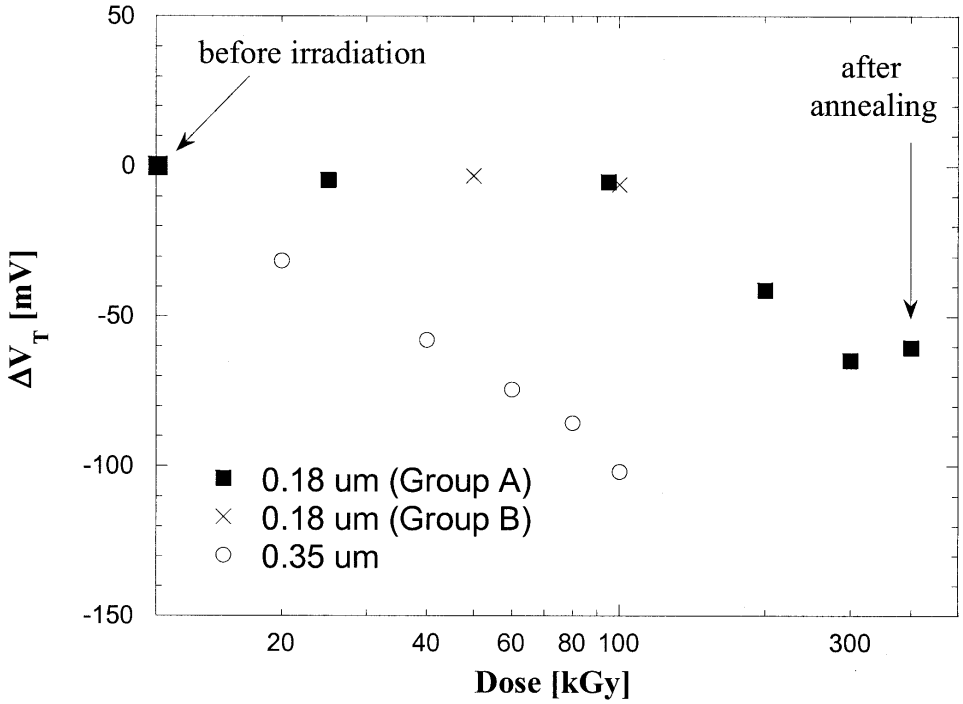

Figure 2.11: Reported threshold voltage change $\Delta V_{\mathrm{T,P}}$ versus total dose for PMOS devices from commercial 0.18-μm and 0.35-μm processes. Change against pre-irradiation $V_{\mathrm{T,P}}$ shown. Pre-irradiation $V_{\mathrm{T,P}}$ is -0.47 V for 0.18-μm process, and is -0.59 V for 0.35-μm process. Dose given assuming Si as incident material. Reproduced in total from *Manghisoni et al.* [2002].

Two radiation-induced processes can alter I_{LEAK}. First, changes in V_{T} cause shifts of the I_{D}-V_{GS} curve toward higher or lower V_{GS}. A shift toward lower V_{GS}, which occurs for an NMOS device with dominant N_{ot} damage, is shown in process (1) of Figure 2.12: such a shift increases I_{LEAK}. Second, N_{it} increase causes a softening of the subthreshold slope. In the subthreshold regime (that is, when $V_{\mathrm{GS}} < V_{\mathrm{T}}$) the drain current varies exponentially with V_{GS}: since I_{D} is plotted in logarithmic scale versus V_{GS}, this variation manifests as a straight line in Figure 2.12. The slope of this line is the subthreshold slope. The inverse of this slope, dubbed the subthreshold swing S_{swing}, can be expressed as [*Sze*, 1981, p. 447]:

$$S_{\mathrm{swing}} = \frac{k_{\mathrm{B}}T}{q}\ln(10)\left(1 + \frac{C_{\mathrm{D}} + C_{\mathrm{it}}}{C_{\mathrm{i}}}\right) \tag{2.5}$$

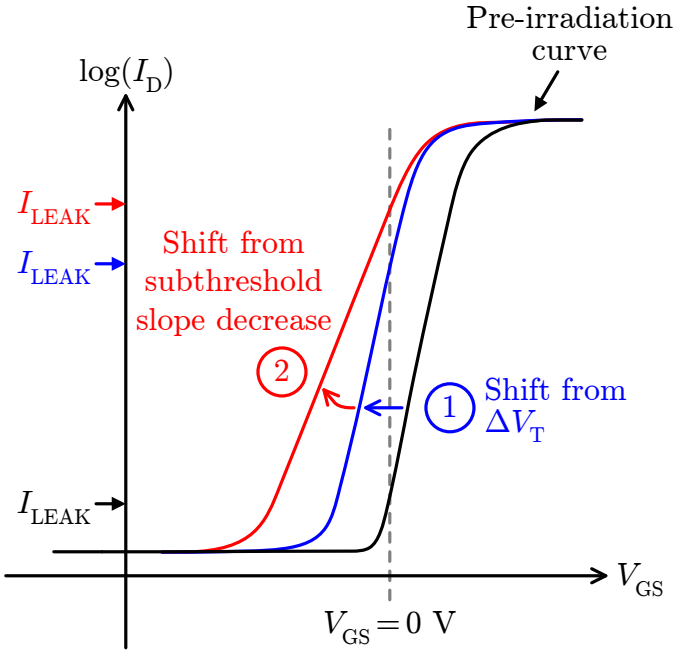

Figure 2.12: Impact of radiation on I_D-V_{GS} curve as regards leakage current. For process (1), V_T shift befitting an NMOS device with dominant N_{ot} damage shown.

where C_D and C_i are the capacitances per unit area of the depletion region in the silicon and of the gate oxide, respectively, and C_{it} is the capacitance associated with the interface traps, that is:

$$C_{it} = qN_{it} \tag{2.6}$$

for N_{it} the interface-trap density in particles per unit area. Radiation-induced increases in N_{it} thus increase S_{swing}: the resulting decrease in the subthreshold slope ultimately raises I_{LEAK}, as shown in process (2) of Figure 2.12.

These effects act not only on the gate oxide, but also act on the adjacent field oxide of a MOSFET, as illustrated in Figure 2.13(a) for a LOCOS field oxide, and Figure 2.13(b) for an STI field oxide. In both cases, parallel to the nominal transistor under the thin gate oxide (shown in green) is a series of parasitic devices formed from the polysilicon overhang and growing field oxide (shown in orange). These parasitic devices are often called FOXFETs.[39] In modern CMOS technologies, the gate oxide is

[39]FOXFET for Field OXide Field Effect Transistor.

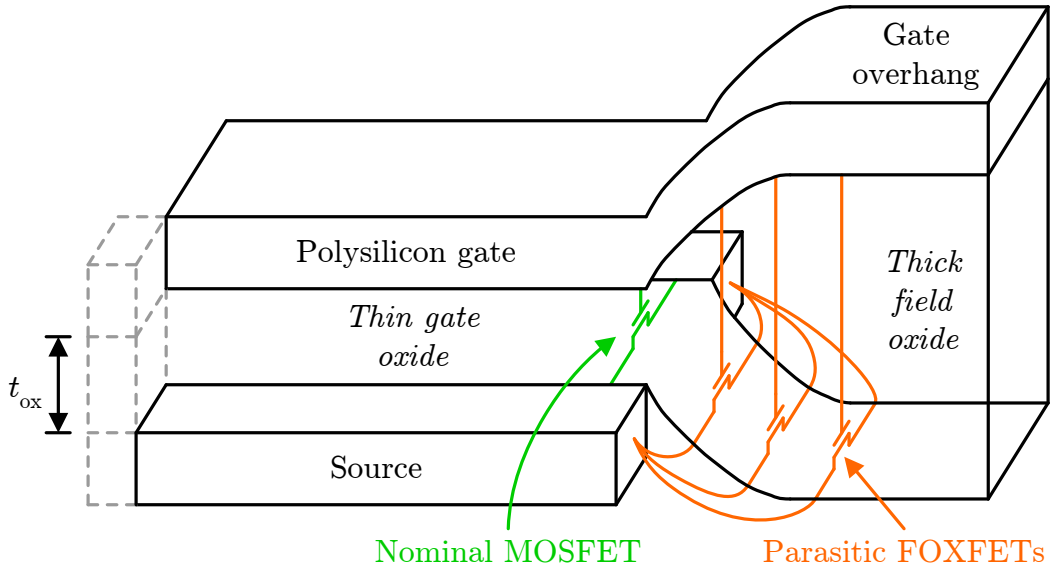

(a) LOCOS field oxide.

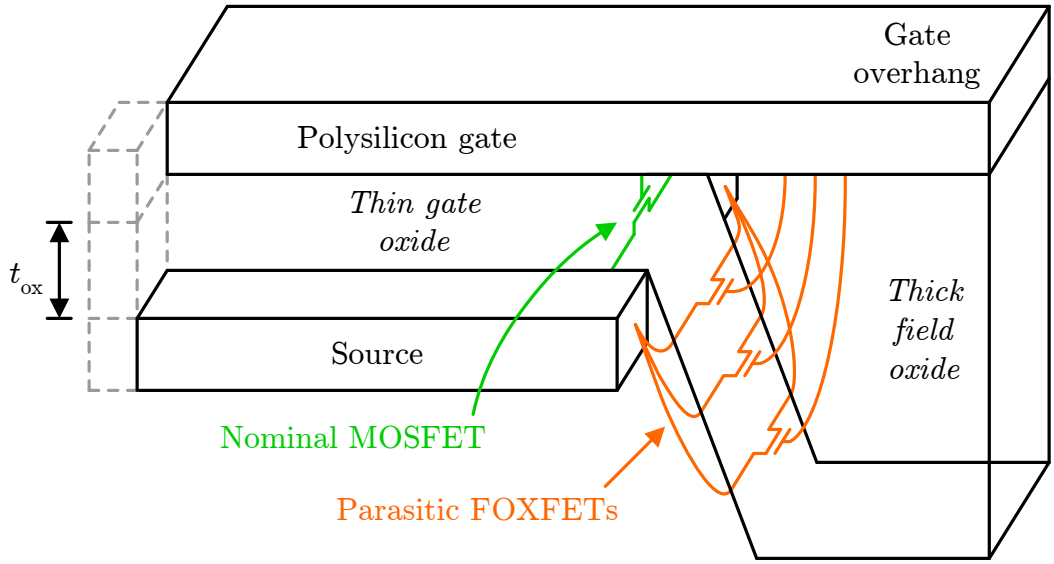

(b) STI field oxide.

Figure 2.13: Profile depicting nominal transistor (green) and parasitic FOXFETs (orange) for a MOSFET with either (a) LOCOS or (b) STI field oxide. Device source in foreground, drain in background (not labeled). Gate oxide thickness t_{ox} labeled to establish perspective: drawing not to scale.

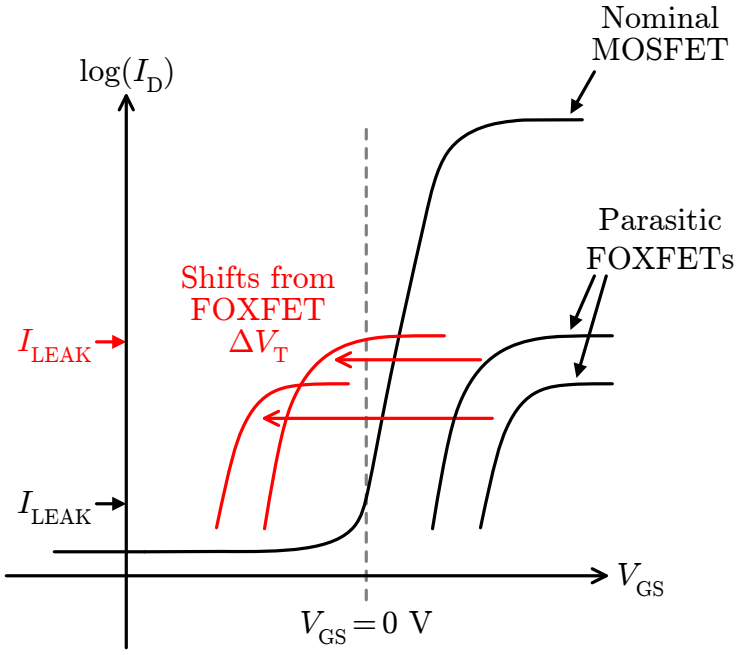

Figure 2.14: Impact of radiation-shifted FOXFET I_D-V_{GS} curves on leakage current. For FOXFETs, pre-irradiation I_D-V_{GS} curves shown in black, radiation-shifted curves shown in red.

thin and fairly radiation hard, translating to little I_{LEAK} contribution under the gate oxide. In contrast, the field oxide is thick and more susceptible to radiation damage. The FOXFET oxides can thus accumulate sufficient damage that significant leakage occurs under the field oxide. This effect is illustrated in Figure 2.14, where FOXFET V_T changes shift their I_D-V_{GS} curves, incurring ample I_{LEAK}.

Measured I_{LEAK} for multiple NMOS devices from the BiCMOS8iED process are shown in Figure 2.15. All devices display a sizable increase in I_{LEAK}—nearly 6 orders of magnitude to tens of microAmpere levels—beginning at ∼100 krad(Si). Given the relatively small $V_{T,N}$ shifts of Figure 2.8, it is unlikely that this I_{LEAK} flows under the thin gate oxide, suggesting it predominantly flows under the thick field oxide. Figure 2.16 shows the I_D-V_{GS} curve of one of these devices, indicative of all. The rise of a fairly flat current floor in the subthreshold regime further suggests translated FOXFET I_D-V_{GS} curves, corroborating a field oxide leakage interpretation.[40] Note

[40]Further evidence bolstering this explanation is presented in Section 5.2.2.2.

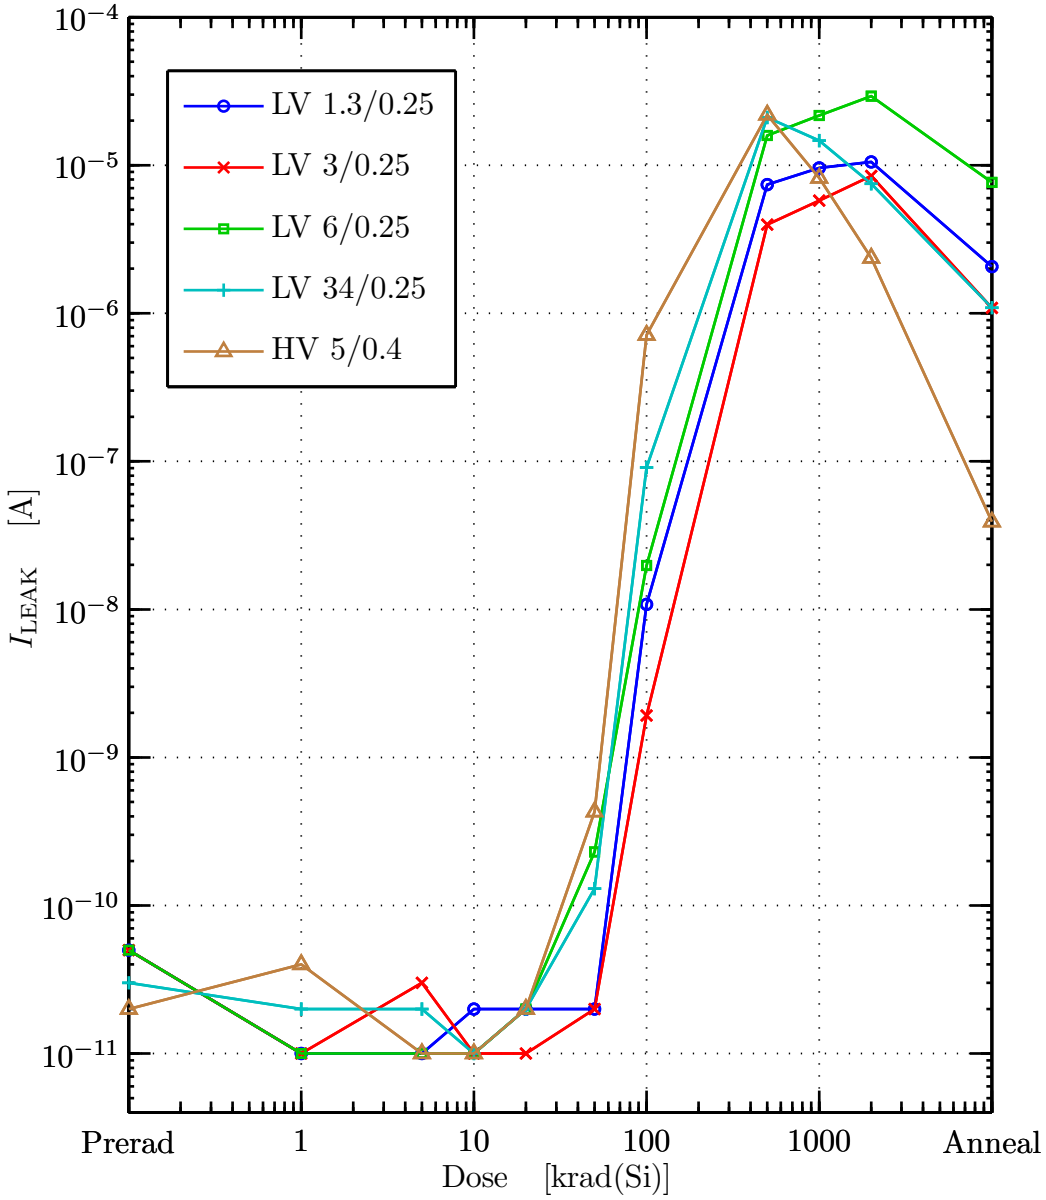

Figure 2.15: Measured leakage current I_{LEAK} versus total dose for multiple BiCMOS8iED NMOS devices. Device type W/L given in legend. I_{LEAK} measured as I_{D} when $V_{\mathrm{D}} = 2.5$ V and $V_{\mathrm{G}} = V_{\mathrm{S}} = V_{\mathrm{B}} = 0$ V. All devices irradiated under worst-case bias. Markers indicate measured data points.

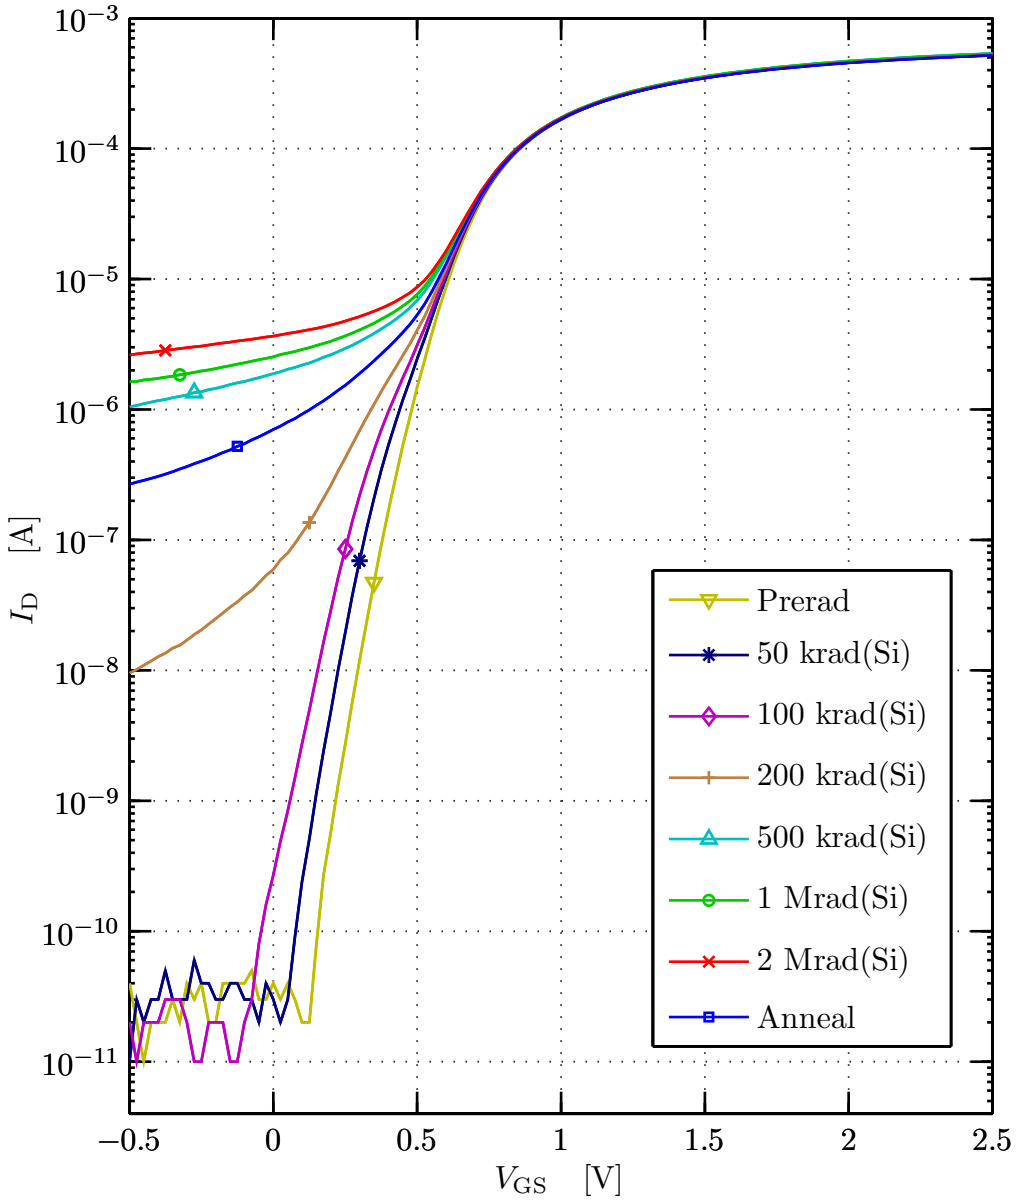

Figure 2.16: Measured radiation evolution of I_D-V_{GS} curve for LV $W/L = 6/0.25$ NMOS device from BiCMOS8iED. Select total doses shown. Measured with $V_S = V_B = 0$ V and $V_D = 0.1$ V; V_G swept in 25 mV increments. Device irradiated under worst-case bias.

that, as the leakage stems primarily from field oxide N_{ot} damage, it is not surprising that it decreases after annealing due to N_{ot} mitigation.

Finally, Figure 2.15 shows only NMOS devices. In modern technologies, I_{LEAK} is not an issue for PMOS devices: as $|V_{T,P}|$ only increases with radiation, I_D-V_{GS} curves only translate towards greater V_{GS}, in fact suppressing I_{LEAK}.

2.3.3.3 Transconductance

In general, the transconductance g_m of both NMOS and PMOS devices decreases with dose since radiation-induced N_{it} increases lattice and Coulomb scattering in the channel, degrading carrier mobility μ. A simple, empirical model for this degradation is [*Winokur*, 1989, p. 213]:

$$\mu = \frac{\mu_0}{1 + \alpha(\Delta N_{it})} \tag{2.7}$$

where μ_0 is the pre-irradiation mobility, ΔN_{it} is the change in the areal density of the interface traps (in particles/cm^2), and α is an empirically fitted parameter.[41] As g_m depends directly on μ, the degradation given in Equation (2.7) translates to transconductance loss. This effect is shown in Figure 2.17, which shows g_m decreases of nearly 10% for both an NMOS and PMOS device from a commercial 0.18-μm process [*Manghisoni et al.*, 2002].

In contrast, Figure 2.18 plots the radiation evolution of maximum transconductance $g_{m,max}$ for NMOS devices of varying size from BiCMOS8iED.[42] Bucking conventional wisdom, after an initial decrease $g_{m,max}$ actually increases with radiation. The change is modest, but the smoothness of the curves, as well as the universality of their character across all the devices (all show a $g_{m,max}$ peak near 500 krad(Si) before declining), suggest the variations are real effects. One explanation attributes the increase to interactions of radiation damage with the lightly doped drain (LDD)

[41] For example, *Sexton and Schwank* [1985] reported an α of $(7.0 \pm 1.3) \times 10^{-13}$ cm^2 for a 1.3-μm CMOS technology.

[42] The $g_{m,max}$ is found by sweeping V_G for fixed V_D and V_S: g_m is then the derivative dI_D/dV_G and $g_{m,max}$ the maximum value over V_G. The calculation of g_m is given formally in Section F.4.1. Given the experimental setup of these measurements (see Section F.3.2.1), $g_{m,max}$ occurs when the device is in the linear regime.

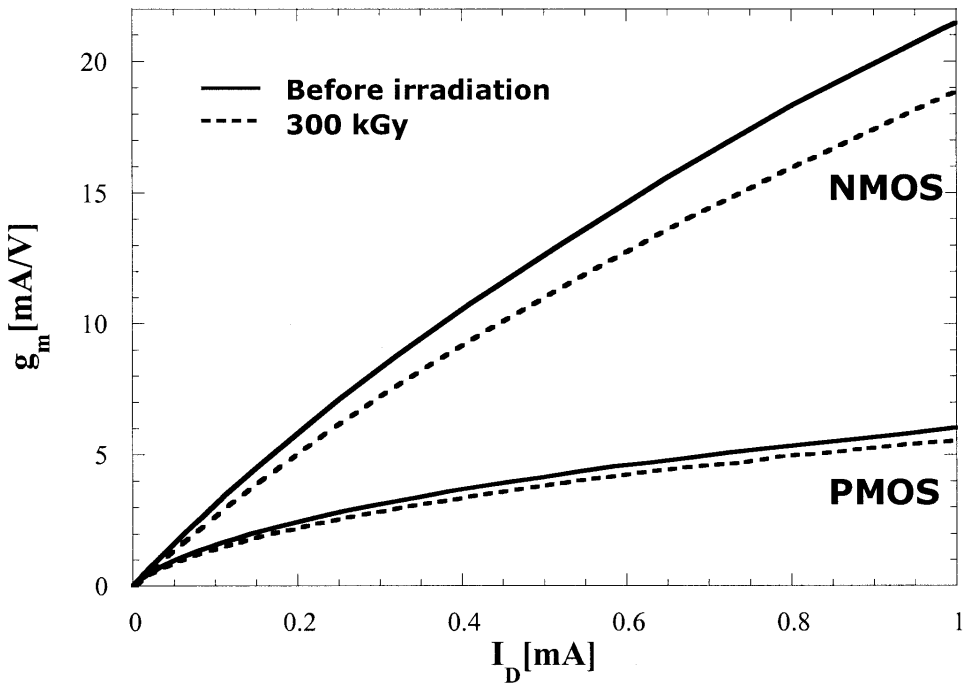

Figure 2.17: Reported transconductance g_{m} versus drain current I_{D} before (solid line) and after (dashed line) irradiation for NMOS ($W/L = 2000/0.7$) and PMOS ($W/L = 200/0.5$) devices from a commercial 0.18-μm technology (Group A). Measured with $|V_{\mathrm{DS}}| = 0.8$ V. Dose given assuming Si as incident material. Reproduced in total from *Manghisoni et al.* [2002].

regions of the device. This region is shown in Figure 2.19. Specifically, radiation-induced N_{ot} in the thicker sidewall oxide can drive the LDD regions more strongly into inversion, effectively shortening the channel and increasing the transconductance. The transconductance increase would be small and would be expected to eventually disappear as the interface trap population continues to accumulate, as occurs in Figure 2.18.[43] A complete confirmation of this explanation would require careful physical simulations of the device in a semiconductor simulator. From a design perspective, though, suffice to say that the result highlights the unpredictability of radiation

[43]This explanation was first suggested by Everett E. King of The Aerospace Corporation. Readers interested in the effects of sidewall oxide damage on device transconductance are recommended to *King et al.* [2000], which addresses the impact of electron damage on the sidewall oxide (in the radiation case, the damage is due to holes instead of electrons, but the principles are nonetheless similar).

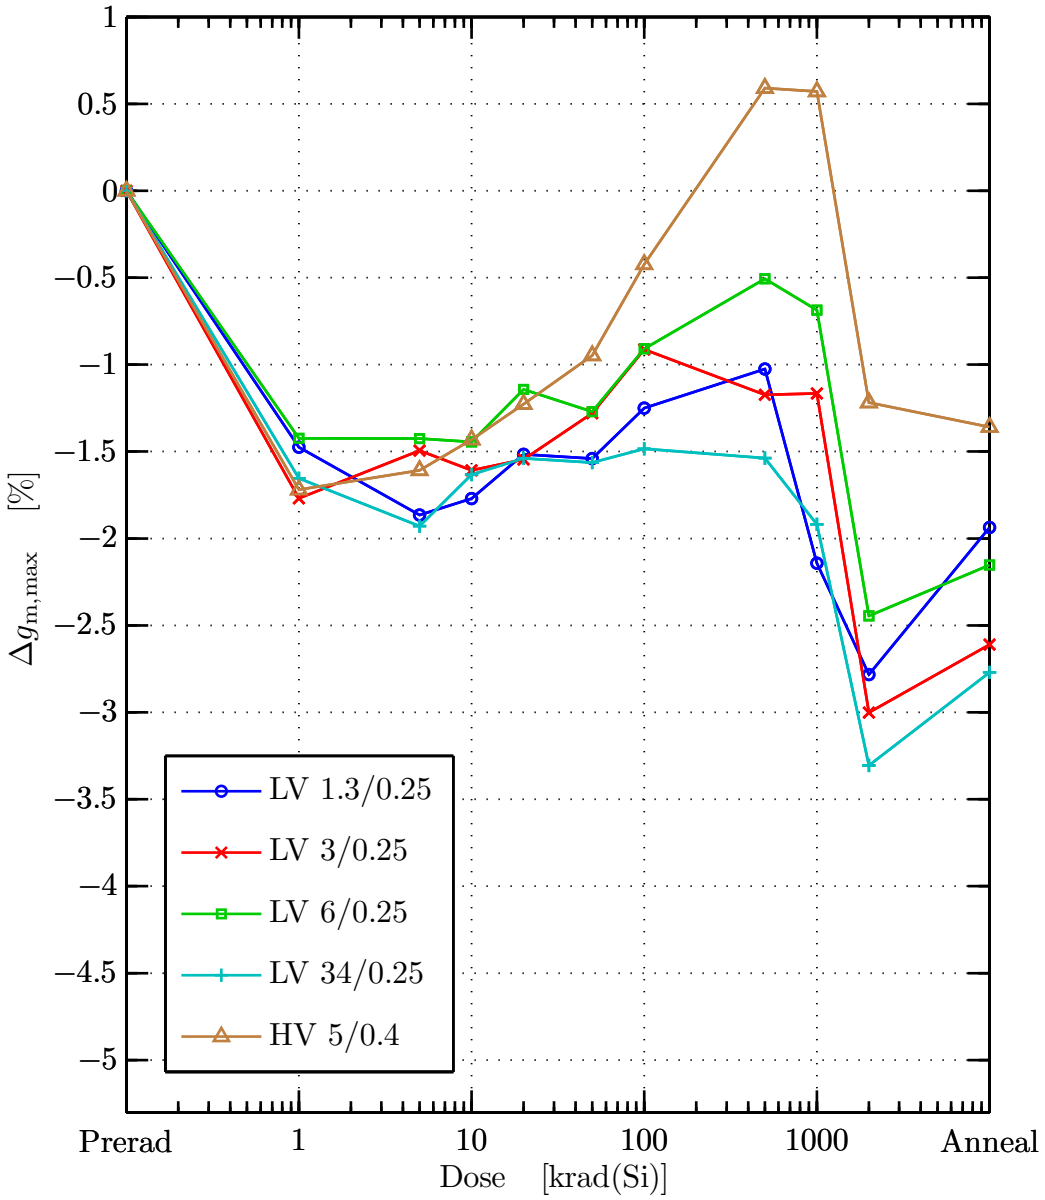

Figure 2.18: Measured change of maximum transconductance $\Delta g_{\mathrm{m,max}}$ versus total dose for multiple BiCMOS8iED NMOS devices. Device type and W/L given in legend. Percent change of $g_{\mathrm{m,max}}$ from pre-irradiation value shown. Measured with $V_{\mathrm{S}} = V_{\mathrm{B}} = 0$ V, $V_{\mathrm{D}} = 0.1$ V, and V_{G} swept from -0.5 V to V_{DD} in 50 mV increments. All devices irradiated under worst-case bias. Markers indicate measured data points.

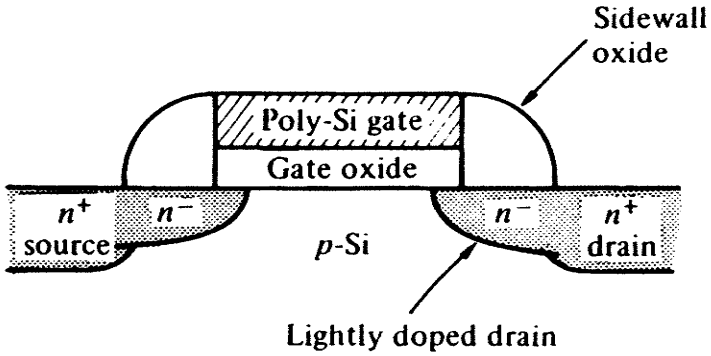

Figure 2.19: Cross-section depiction of the lightly doped drain (LDD) regions of an NMOS device. Reproduced in total from *Pierret* [1990, p. 139].

response.[44]

2.3.3.4 Noise

The noise of a MOSFET is often characterized in terms of its gate-referred noise voltage spectrum $S_e(f)$. This noise can be modeled with a white and flicker component as per [*Manghisoni et al.*, 2002]:

$$S_e^2(f) = \underbrace{S_w^2}_{\text{white noise}} + \underbrace{\frac{K_a}{C_{ox}^2 W L}\left(\frac{1}{f}\right)}_{\text{flicker noise}} \tag{2.8}$$

where K_a is an intrinsic process parameter and C_{ox} is the gate capacitance per unit area.

[44]The phenomenon of increasing $g_{m,max}$ with dose is not completely without precedent: *Nowlin et al.* [2005] reported $g_{m,max}$ increases of greater than a factor of four in devices from a 0.35-μm technology. However, these increases were seen in NMOS transistors drawn in nonstandard fashion, specifically, drawn in the enclosed terminal layout fashion discussed in Section 5.2.2. *Nowlin et al.* [2005] attributed the $g_{m,max}$ increase to the parallel current introduced by the series combination of the parasitic ring-MOSFET and edge-FOXFET associated with the enclosure (see footnote 15 of Chapter 5). In particular, physical device simulations affirmed that the latter functions as a current source that limits the former: this sudden saturation in the measured current results in a $g_m = dI_D/dV_{GS}$ spike. Unfortunately, such physics are not directly applicable here since Figure 2.18 depicts measurements of standard devices that do not have such a ring-MOSFET, edge-FOXFET structure.

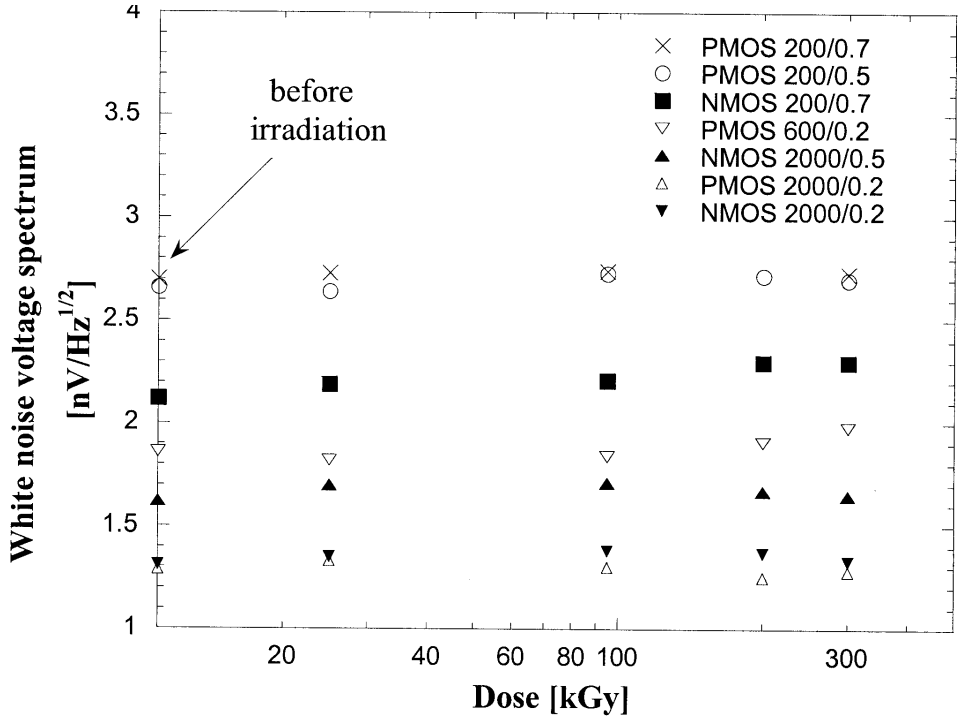

Figure 2.20: Reported white noise voltage spectrum $S_{\rm w}^2$ versus total dose for multiple NMOS and PMOS devices from a commercial 0.18-μm process (Group A). Measured with $|V_{\rm DS}| = 0.8$ V and $I_{\rm D} = 0.5$ mA for $W/L = 2000/0.2$ devices, $I_{\rm D} = 0.25$ mA for other devices. Dose given assuming Si as incident material. Reproduced in total from *Manghisoni et al.* [2002].

The radiation evolution of the white and flicker noise components of varying devices from commercial 0.18-μm and 0.35-μm processes is shown in Figures 2.20 and 2.21, respectively [*Manghisoni et al.*, 2002]. The white noise component shows little change with dose. However, the flicker noise component shows increases in $K_{\rm a}$ by a factor of 1.5 to 2 with dose. This result is expected: flicker noise is associated with interface traps (even though the exact mechanism is not yet well understood), and radiation-induced $N_{\rm it}$ should increase flicker noise. Flicker noise is especially important for plasma wave receivers, as it often dominates CMOS device noise at the low frequencies of the phenomena of interest.

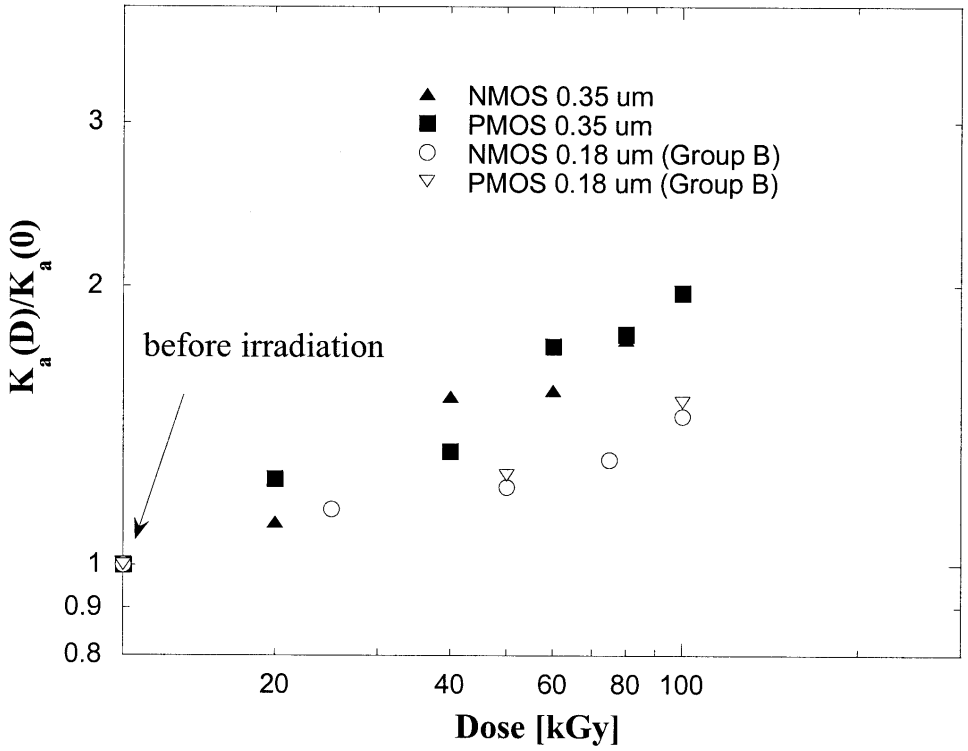

Figure 2.21: Reported flicker noise parameter K_{a} versus total dose for multiple NMOS and PMOS devices from commercial 0.18-μm and 0.35-μm processes. Ratio of K_{a} as a function of dose ($K_{\mathrm{a}}(D)$) to pre-irradiated value ($K_{\mathrm{a}}(0)$) shown. Dose given assuming Si as incident material. Reproduced in total from *Manghisoni et al.* [2002].

2.4 Conclusion

This chapter presented an introduction to the effects of radiation, specifically ionizing radiation, on MOS devices. These effects were divided into two categories: single-event effects (encompassing immediate effects from a single ionizing strike) and total-dose effects (encompassing cumulative long-term radiation damage). Single-event effects included latchup, a catastrophic, high-current failure condition. Total-dose effects included changes in MOSFET behavior due to accumulated populations of oxide trapped charge (N_{ot}) and interface traps (N_{it}) in the thin gate and thick field oxides of the device. In particular, the effects of N_{ot} and N_{it} on the key analog parameters of threshold voltage, leakage current, transconductance, and device noise, were investigated. These investigations included measured results of such

behavior shifts for the BiCMOS8iED process of the SVADC-1, supplemented by other, published results. Many of these results accorded with conventional wisdom, although some ran counter, underscoring the uniqueness of radiation response for each process.

In addition, the results presaged the challenges of designing circuits for radiation environments. The possibility of radiation-induced latchup, which can physically destroy a chip, is a prevailing concern. And the parameter shifts described—especially the increase in I_{LEAK}—can cause additional problems. Robust circuit design for radiation environments must thus take these effects into account. In particular, the SVADC-1 of this dissertation assumes a philosophy of radiation-hardness by design wherein circuits are designed to prevent or compensate for the radiation effects that most limit performance. This approach is expounded in detail in the descriptions of the circuit implementation of the SVADC-1 presented in Chapter 5.

Before then, though, recall that radiation tolerance is but one of the requirements of the SVADC-1. The other, naturally, is that the SVADC-1 perform the requisite analog-to-digital conversion. The next chapter thus investigates the conversion requirements in more detail. Appreciation of the radiation requirements and conversion requirements together enables a full specification of the SVADC-1.

Chapter 3

Analog-to-Digital Conversion

It is often said that real world signals such as plasma waves are analog, that is, continuous in both time and value. To better facilitate processing of these signals, it is often desirable to represent them as digital signals that are discrete in both time and value. This dual discretization is referred to as analog-to-digital conversion.

This chapter introduces analog-to-digital conversion and derives key conversion specifications for the SVADC-1. To begin, discretization in time (sampling) and discretization in value (quantization) are examined in more detail. While the former can be well described analytically, the latter is complicated by its inherent nonlinearity. Thus this chapter opts for a more practical approach, characterizing converter nonlinearity by metrics measurable in a real world context. These metrics are introduced and, given the heavy reliance on spectrograms in plasma wave analysis, the spurious-free dynamic range (SFDR) is identified as a key metric, in particular, the converter should achieve at least 90-dB SFDR (assuming a 100-Hz bin width). The consequences of this requirement on other converter properties—including sampling rate, quantizer resolution, and circuit noise—are assessed, and new results regarding the SFDR of noisy quantizers are introduced. An updated converter specification follows. The chapter concludes by considering the space of currently available, radiation-tolerant ADCs in light of these updated specifications: the lack of a suitable ADC motivates construction of a custom converter, the SVADC-1.

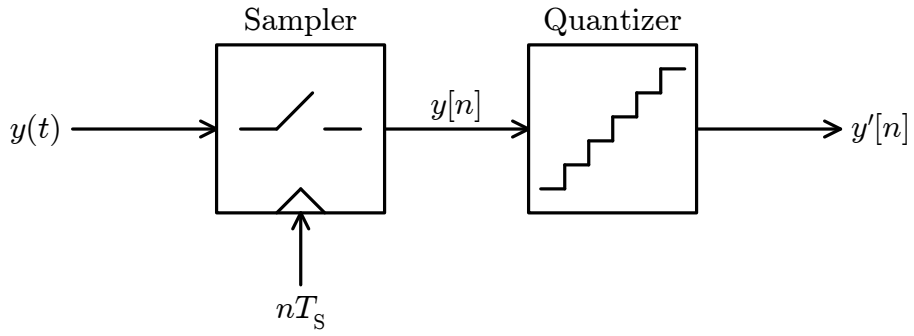

(a) Sampling then quantization.

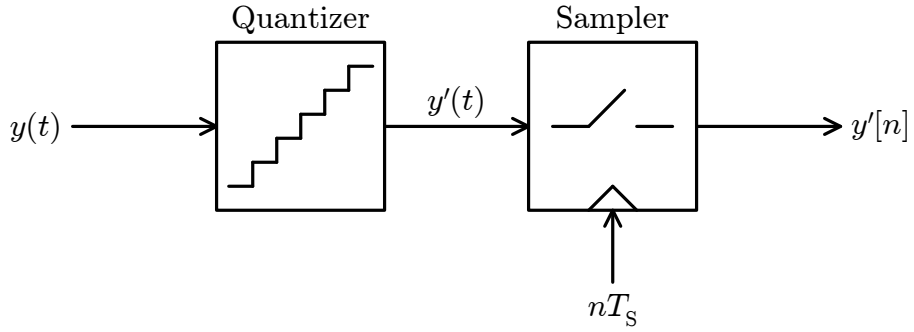

(b) Quantization then sampling.

Figure 3.1: Analog-to-digital conversion represented as sampling and quantization.

3.1 Fundamentals

Analog-to-digital conversion is typically analyzed by considering its two operations of sampling and quantization in series as shown in Figure 3.1. The representations are conceptual: while offering convenient ways to study the conversion process, they are not rigorously mathematically equivalent, and the order of operations chosen often depends on the conversion property under consideration. The following discussions adopt the model of a sampler followed by a quantizer as shown in Figure 3.1(a).

A note on nomenclature is in order. In this dissertation, a continuous argument is enclosed in parentheses $(\cdot)$ while a discrete argument is enclosed in square brackets $[\cdot]$. For example, $y(t)$ is continuous-time while $y[n]$ is discrete-time. In addition, a primed signal indicates the quantized version of the unprimed signal, so that $y'[n]$ is

the quantized version of $y[n]$.

3.1.1 Sampling

Sampling converts a continuous-time signal $y(t)$ into a discrete-time signal $y[n]$. The $y[n]$ are called the samples of $y(t)$. Mathematically, $y[n]$ is related to $y(t)$ as:

$$y[n] = y(nT_S) \tag{3.1}$$

where T_S is the sampling interval, or time between samples. For this dissertation, T_S is assumed constant, in which case Equation (3.1) represents uniform periodic sampling. The reciprocal of T_S is the sampling frequency f_S:

$$f_S = \frac{1}{T_S} \tag{3.2}$$

Sampling in the time domain corresponds to replication in the frequency domain. The frequency domain is defined by the continuous-time Fourier transform, where conversion between the continuous-time function $x(t)$ and its frequency domain representation $X(f)$ is through the Fourier transform pair:

$$X(f) = \int_{-\infty}^{\infty} x(t)e^{-i2\pi ft}\,\mathrm{d}t \quad , \quad x(t) = \int_{-\infty}^{\infty} X(f)e^{+i2\pi tf}\,\mathrm{d}f \tag{3.3}$$

$X(f)$ is called the spectrum of $x(t)$.[1] The sampling operation of Equation (3.1) can

[1] This Fourier transform definition is the one used by Bracewell, and readers interested in a deeper understanding of the Fourier transform are recommended to his classic text *The Fourier Transform and Its Applications* [*Bracewell*, 1986]. Another good resource is Oppenheim, especially his texts *Signals & Systems* [*Oppenheim et al.*, 1997] and *Discrete-Time Signal Processing* [*Oppenheim et al.*, 1999]. Oppenheim, though, adopts the ω-based definition [*Oppenheim et al.*, 1997, p. 288]:

$$X(j\omega) = \int_{-\infty}^{\infty} x(t)e^{-j\omega t}\,\mathrm{d}t \quad , \quad x(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} X(j\omega)e^{j\omega t}\,\mathrm{d}\omega$$

As the Oppenheim texts emphasize the interrelations between the various continuous-time and discrete-time transforms, they are cited extensively here. However, if only for its fearful symmetry (but also for its utility: see footnote 18 of this chapter, for example), this dissertation adopts the Bracewell definition. This choice does cause the peculiar situation wherein a result cites Oppenheim even though it technically cannot be found there as stated. Instead, for exact equivalence,

be modeled as multiplication of $y(t)$ by the impulse train $\psi_{\mathrm{s}}(t)$, where:

$$\psi_{\mathrm{s}}(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT_{\mathrm{S}}) \tag{3.4}$$

for $\delta(t)$ the Dirac delta function. This time domain operation is shown in Figure 3.2(a). The sampled signal $y_{\mathrm{s}}(t)$ is then:

$$y_{\mathrm{s}}(t) = y(t)\psi_{\mathrm{s}}(t) = \sum_{n=-\infty}^{\infty} y(nT_{\mathrm{S}})\delta(t - nT_{\mathrm{S}}) = \sum_{n=-\infty}^{\infty} y[n]\delta(t - nT_{\mathrm{S}}) \tag{3.5}$$

To derive $Y_{\mathrm{s}}(f)$, the spectrum of $y_{\mathrm{s}}(t)$, recall that multiplication in the time domain corresponds to convolution in the frequency domain [*Oppenheim et al.*, 1997, p. 322]. It can be shown that $\psi_{\mathrm{s}}(t)$ has Fourier transform [*Oppenheim et al.*, 1997, p. 299]:

$$\Psi_{\mathrm{s}}(f) = \frac{1}{T_{\mathrm{S}}} \sum_{m=-\infty}^{\infty} \delta\left(f - \frac{m}{T}\right) \tag{3.6}$$

Hence [*Oppenheim et al.*, 1997, p. 517]:

$$Y_{\mathrm{s}}(f) = Y(f) \otimes \Psi(f) = \frac{1}{T_{\mathrm{S}}} \sum_{m=-\infty}^{\infty} Y(f - mf_{\mathrm{S}}) \tag{3.7}$$

Sampling in the time domain thus creates replicas, or aliases, of the original signal spectrum $Y(f)$ in the frequency domain, spaced at integer multiples of the sampling frequency f_{S}, as shown in Figure 3.2(b).[2]

Oppenheim's derivations must be repeated under Bracewell's definition: such machinations are left as an exercise for the reader.

[2]The astute will note that this derivation has focused on the continuous-time signal $y_{\mathrm{s}}(t)$ and not the discrete-time signal $y[n]$. The two are different—the former is a function of a continuous variable, the latter a sequence of numbers—and it is not clear that the spectra of the two are the same. In fact, in an information sense, the spectra of impulse-train-sampled continuous-time signals and their discrete-time equivalents are indeed different. Recall that the transform pair for the continuous-time signal is (Equation (3.3)):

$$Y_{\mathrm{S}}(f) = \int_{-\infty}^{\infty} y_{\mathrm{s}}(t)e^{-i2\pi ft}\,\mathrm{d}t \quad , \quad y_{\mathrm{s}}(t) = \int_{-\infty}^{\infty} Y_{\mathrm{S}}(f)e^{+i2\pi tf}\,\mathrm{d}f$$

where f is the continuous-time frequency. The transform pair for the discrete-time equivalent is

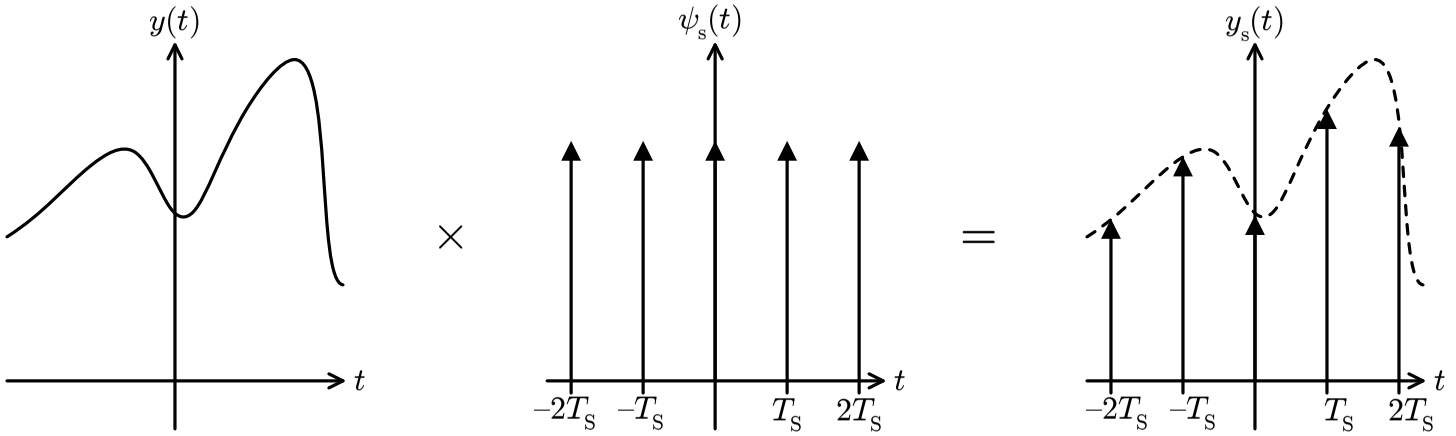

(a) Time domain representation.

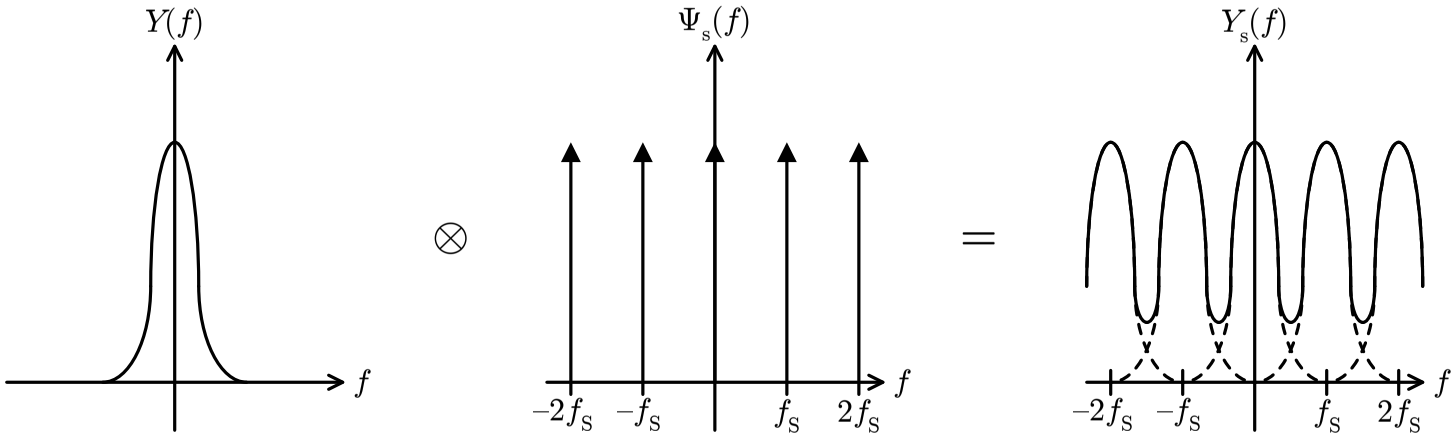

(b) Frequency domain representation.

Figure 3.2: Time domain sampling depicted in the (a) time and (b) frequency domains. The operation $\times$ denotes multiplication and the operation $\otimes$ denotes convolution.

Insofar as the replicas of $Y(f)$ overlap, they introduce aliasing distortion. However, if the original spectrum $Y(f)$ is bandlimited to a bandwidth f_{N}, that is:

$$Y(f) = 0 \quad \text{for} \quad |f| > f_{\mathrm{N}} \tag{3.8}$$

then sampling at a rate greater than $2f_{\mathrm{N}}$ prevents aliasing distortion. The argument is illustrated geometrically in Figure 3.3.[3] The frequency $2f_{\mathrm{N}}$ is called the Nyquist rate and sampling at $2f_{\mathrm{N}}$ is called Nyquist rate sampling. Sampling at frequencies greater than $2f_{\mathrm{N}}$ is called oversampling. A key result from sampling theory is that if $y(t)$ is bandlimited and sampled at the Nyquist rate or higher, then all of its information is captured in its samples $y[n]$. This result is typically stated mathematically by noting that a bandlimited signal $y(t)$ sampled at the Nyquist rate or higher can be perfectly reconstructed from its samples. For example, an oft-cited reconstruction formula is [*Oppenheim et al.*, 1999, p. 150]:

$$y(t) = \sum_{n=-\infty}^{\infty} y[n]\mathrm{sinc}(f_{\mathrm{S}}(t - nT_{\mathrm{S}})) \tag{3.9}$$

where:

$$\mathrm{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \tag{3.10}$$

given by the discrete-time Fourier transform (DTFT) [*Oppenheim et al.*, 1997, p.361] and is:

$$Y(\tilde{f}) = \sum_{n=-\infty}^{\infty} y[n]e^{-i2\pi\tilde{f}n} \quad , \quad y[n] = \int_{-1/2}^{+1/2} Y(\tilde{f})e^{+i2\pi n\tilde{f}}\,\mathrm{d}\tilde{f}$$

where $\tilde{f}$ is the discrete-time frequency. Comparing the two pairs, note that the entirety of the spectrum of f is required to transform $Y_{\mathrm{S}}(f)$ to $y(t)$, whereas only the spectrum of $\tilde{f}$ over $\tilde{f} \in [-1/2, 1/2]$ is required to transform $Y(\tilde{f})$ to $y[n]$. In a sense, then, the information of $Y_{\mathrm{S}}(f)$ extends over all frequencies, whereas the information of $Y(\tilde{f})$ is contained within a small range of frequencies. However, though their informational content is different, it can be shown that the values of $Y_{\mathrm{S}}(f)$ and $Y(\tilde{f})$ are the same to within a frequency scaling factor: see footnote 18 of this chapter. The fact that the frequency domain is the subject of this derivation and ensuing discussion justifies studying the spectrum of $y_{\mathrm{s}}(t)$ as a proxy for that of $y[n]$.

[3]Indeed, from the overlap geometry, a stronger bound can be set: if $Y(f)$ is bandlimited such that it is nonzero only in a closed, bounded interval of total frequency extent f_{N}, then Nyquist rate sampling still prevents aliasing distortion. This leads to techniques such as subsampling, where an intentionally low sampling rate effectively heterodynes down a narrowband high frequency signal by forcing a replica to fall into a lower frequency band.

The identification of the Nyquist rate as the minimal sampling rate required for perfect reconstruction is often referred to as the Nyquist-Shannon sampling theorem [*Nyquist*, 1928; *Shannon*, 1949].[4]

To prevent aliasing distortion, an anti-aliasing filter often precedes sampling to limit the pre-sampled signal bandwidth. Practically, this filter does not render out-of-band frequency components identically zero as required by theory, but instead suppresses them enough that aliasing distortion is sufficiently minimal for the application.[5]

3.1.2 Quantization

Quantization converts the continuous-valued samples $y[n]$ into discrete-valued samples $y'[n]$. The ideal input-output characteristic of a quantizer is shown in Figure 3.4. The quantizer divides the continuum of possible input values into a series of regions called bins; the bin boundaries are the transition levels or decision levels. Figure 3.4 shows a uniform quantizer wherein the width of each bin is the same: this width is the

[4] Attribution of the sampling theorem is actually rather more involved. In his paper (written in 1940 and published in 1949, though possibly selectively circulated before then [*Higgins*, 1985]), Shannon himself noted that the theorem had "been given previously in other forms by mathematicians but in spite of its evident importance seems not to have appeared explicitly in the literature of communication theory" [*Shannon*, 1949]. Shannon subsequently referenced the work of J. M. Whittaker on the cardinal series, establishing a mathematical basis for the theorem as far back as E. T. Whittaker (the father of J. M.) who examined the series as early as 1915 [*Higgins*, 1985]. However, the list of origins is wider than that. The Russian literature ascribes the theorem to Kotel'nikov, who published a version in 1933 (the Russians instead reference the Kotel'nikov sampling theorem). Other attributions include Someya in 1949 in the Japanese literature, and Weston in 1949 in the English. All are believed to be independent introductions of the theorem [*Higgins*, 1985]. Those further interested in the history of the sampling theorem are directed to *Higgins* [1985] to start. Higgins traces the theorem back to Borel in 1897, who (a rough familiarity with French is useful here) effectively noted the sufficiency of Nyquist rate sampling, though did not provide a consequent reconstruction formula. Those who find Higgins heightens more than slackens their thirst are further directed to *Jerri* [1977] (check [*Jerri*, 1979] for some minor corrections).

[5] This limitation is more fundamental than it would first appear. It can be shown that a (nontrivial) signal cannot be both bandlimited and time-limited. Hence, a bandlimited signal must have infinite time extent. While signals of infinite time extent may be mathematically possible, in practice signals are observed for only finite durations. Being thus time-limited, these practical signals cannot technically be bandlimited. It is possible, though, to refine the concepts of "bandlimited" and "time-limited" to better reflect the real world situation, and further to show the sampling theorem holds under the amended concepts. Readers intrigued by such subtleties are recommended to *Slepian* [1976] as a good starting point.

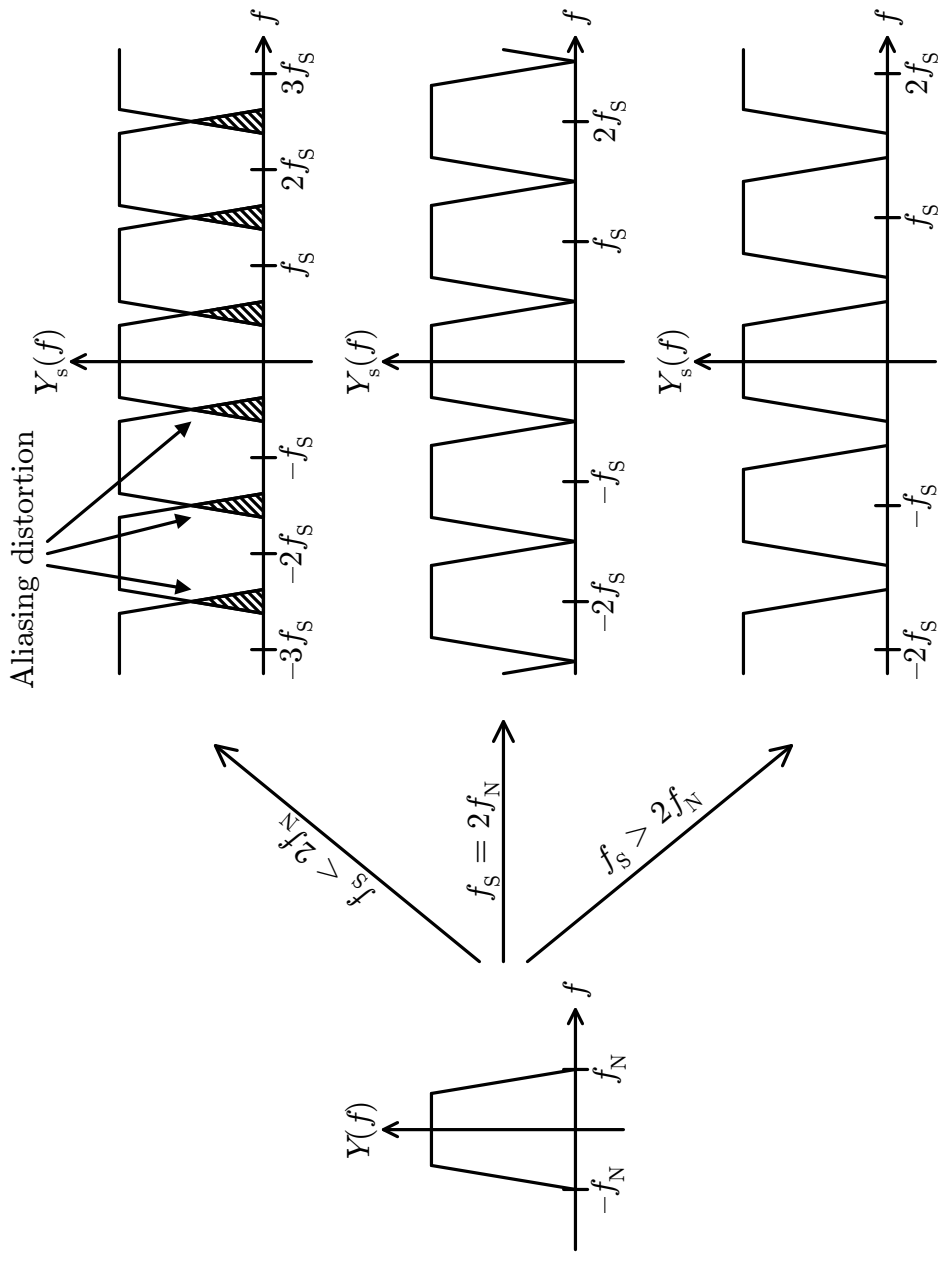

Figure 3.3: Frequency domain view of sampling a bandlimited function at different rates. From top to bottom: sampling at a rate below, at, and above the Nyquist rate. Aliasing distortion shaded in the topmost case.

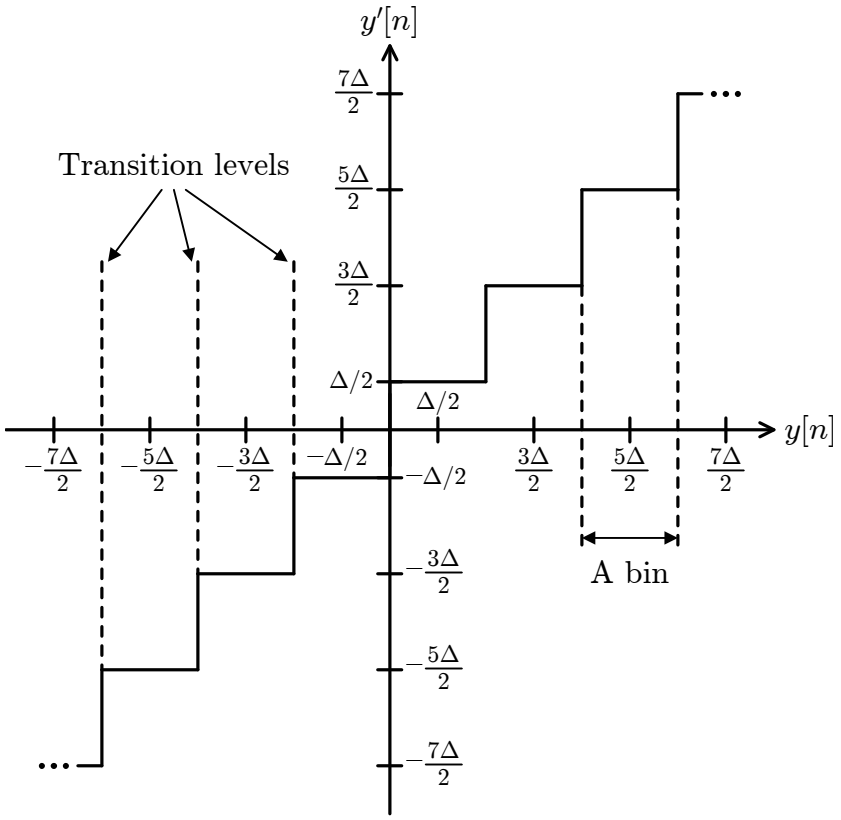

Figure 3.4: Input-output characteristic of a uniform, unbounded quantizer.

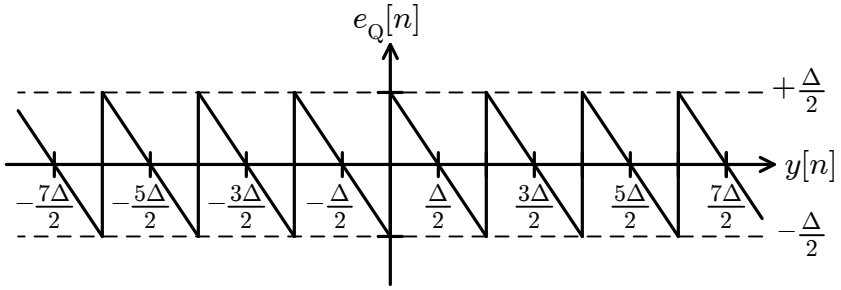

Figure 3.5: Quantizer error $e_{\mathrm{Q}}[n]$ of a uniform, unbounded quantizer.

quantizer step size Δ. When the input value falls within a bin, the quantizer outputs the center value of that bin, essentially converting the continuum to a discrete set by rounding. Naturally, quantization incurs inaccuracy. The quantizer error $e_{\mathrm{Q}}[n]$ is defined as:

$$e_{\mathrm{Q}}[n] = y'[n] - y[n] \tag{3.11}$$

The $e_{\mathrm{Q}}[n]$ of the ideal quantizer of Figure 3.4 is shown in Figure 3.5: as can be seen, $e_{\mathrm{Q}}[n]$ is regular and bounded between $\pm\Delta/2$.[6] Decreasing Δ thus decreases $e_{\mathrm{Q}}[n]$, improving input-output fidelity. Quantizers with small Δ (compared to the range spanned by the quantizer input signal) are said to be "high resolution" in that they determine (or resolve) the input value to a high degree of accuracy (or, equivalently, to a low degree of uncertainty). Quantizers with large comparative Δ are said to be "low resolution".

Practically implementable quantizers differ from the ideal one of Figure 3.4 in that they are bounded: quantization can be maintained only over a finite input range. Beyond this range, the quantizer overranges and saturates. The input-output characteristic of a bounded quantizer is depicted in Figure 3.6 and the resulting quantizer error is depicted in Figure 3.7. For the extremum bins, the quantizer error becomes unbounded. As labeled in Figure 3.7, the input value range over which the quantizer error remains bounded is called the full-scale range V_{FS}.[7] A finite number of output levels M_{level} is thus distributed over V_{FS}. A quantizer is said to be "coarse" if M_{level} is small, and "fine" if M_{level} is large.

The number of bits needed to represent all the output levels of a bounded quantizer is:

$$M_{\mathrm{min}} = \log_2(M_{\mathrm{level}}) \tag{3.12}$$

Quantizers are often identified by M_{min}: for example, a quantizer with 7 output

[6] Figure 3.4 depicts a midrise quantizer, where $y[n] = 0$ corresponds to a transition level. Alternately, quantizers may be midtread, where $y[n] = 0$ corresponds to the center of a bin with transitions levels at $\pm\Delta/2$. Conversion between the two is possible by addition of DC constants: for example, adding an offset of $+\Delta/2$ before, and $-\Delta/2$ after, a midrise quantizer results in a midtread quantizer [*Widrow and Kollár*, 2008, pp. 4–5]. This chapter thus just considers midrise quantizers.

[7] Technically, the particular bound chosen depends on the application, but typically, for uniform quantizers, the canonical bound of $\pm\Delta/2$ is used as depicted in Figure 3.7.

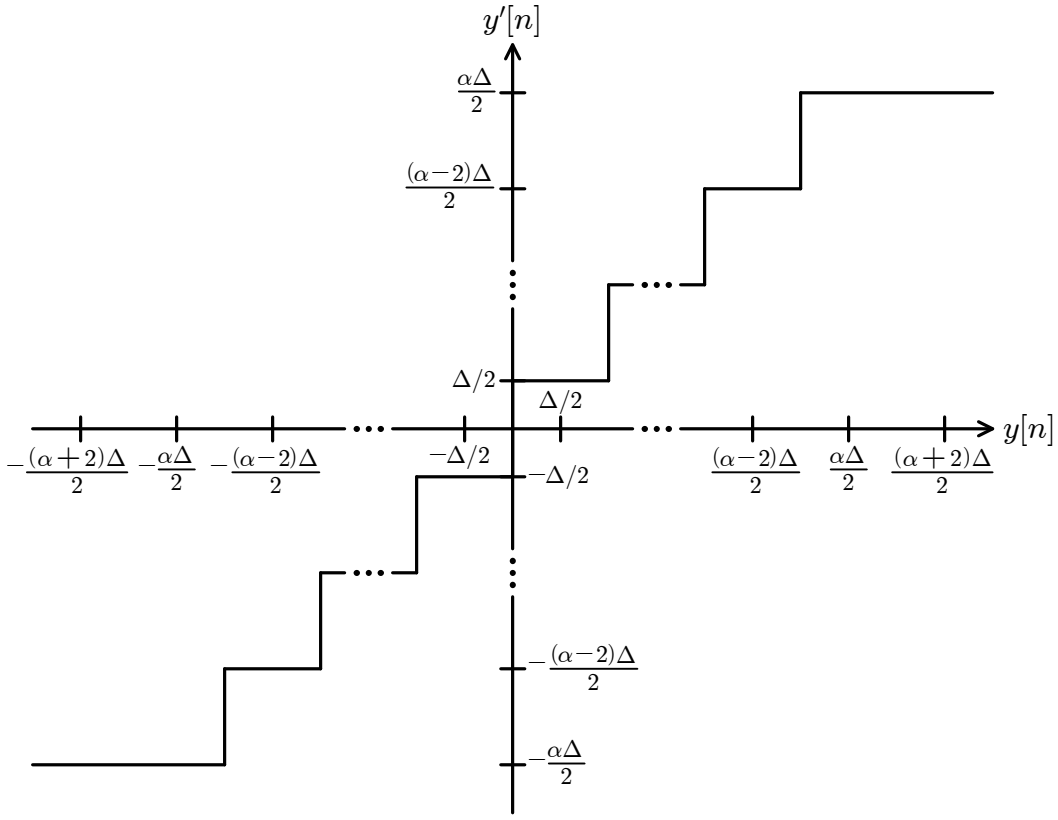

Figure 3.6: Input-output characteristic of a uniform, bounded quantizer.

levels is called a 2.8-bit quantizer.[8] The quantizer output is typically represented by

[8]Many other works use a more formal, entropy-based definition for the number of bits of a quantizer. Let X be a discrete random variable that can take on values $\{x_0, x_1, \ldots, x_{M-1}\}$. The entropy H of X is defined as [*Cover and Thomas*, 1991, p. 13]:

$$H(X) = -\sum_{k=0}^{M-1} p[x_k]\log_2(p[x_k])$$

where $p[x_k]$ is the probability of outcome x_k. It can be shown that the minimal number of bits required to encode X is given by $H(X)$ [*Cover and Thomas*, 1991, p. 86]. A quantizer, then, can be thought of as a discrete encoder with output X. $H(X)$ is then the number of bits of the quantizer. If it is assumed that the input is uniformly distributed over the quantizer input range (so that $p[x_k]$ reduces to the fraction of the input range occupied by the k-th bin), then it can be shown that, for a quantizer with M_{level} uniformly sized bins, $H(X) = M_{\text{min}}$. For a quantizer with nonuniformly sized bins, though, generally $H(X) \neq M_{\text{min}}$. However, for the nonuniform quantizers of this dissertation, it turns out that $H(X) \simeq M_{\text{min}}$ to enough precision that quantizer names in this dissertation are consistent with other works.

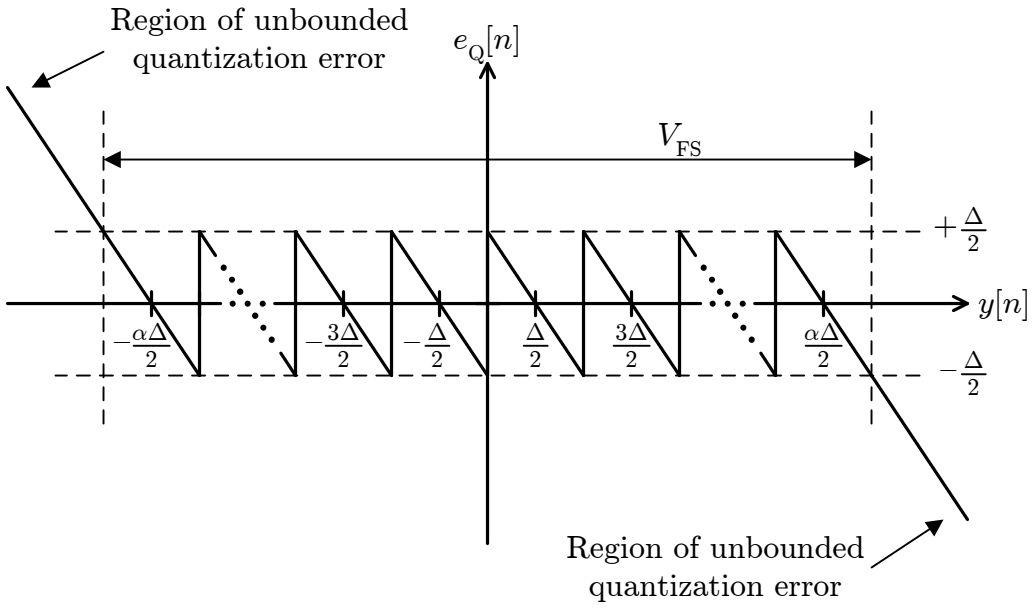

Figure 3.7: Quantizer error $e_{\mathrm{Q}}[n]$ of a uniform, bounded quantizer.

a binary word of an integer number of bits. The minimum integer number of bits needed is:

$$M_{\mathrm{out}} = \lceil M_{\mathrm{min}} \rceil \tag{3.13}$$

where $\lceil \cdot \rceil$ is the ceiling operation. The quantizer output, then, is represented by binary words of at least M_{out} bits. These output words are also interchangeably referred to as "codes".[9] While these output words may be represented by many formats, in this dissertation by convention they are assumed represented as unsigned integers, with 0 corresponding to the lowest output level. In describing the output words, it is customary to refer to the bit positions of highest and lowest value as the most significant bit (MSB) and least significant bit (LSB), respectively.

Often, M_{level} is chosen to be a power of 2 so that $M_{\mathrm{out}} = M_{\mathrm{min}}$. In this case, the LSB is equivalent to the quantizer step size Δ. However, while certainly efficient, the minimal M_{out} need not always be chosen. A 12-bit quantizer, for example, may be represented by 16-bit output words: though this quantizer has 2^{16} total possible codes,

[9]This nomenclature is in keeping with the interpretation of a quantizer as a discrete coder, as described in footnote 8 of this chapter.

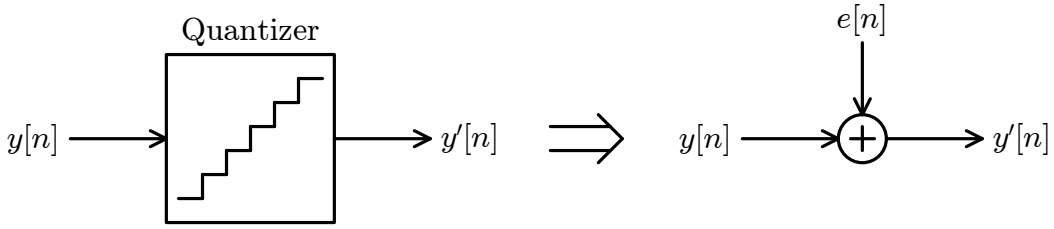

(a) Quantizer model.

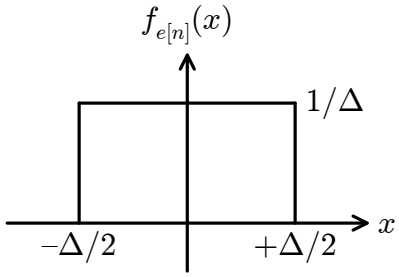

(b) PDF of $e[n]$.

Figure 3.8: The PQN model. (a) The quantizer replaced by an additive noise $e[n]$, (b) the probability density function (PDF) of $e[n]$.

only 2^{12} of the codes are active. Furthermore, these 2^{12} codes may be distributed such that all 16 bits are significant: the 2^{12} need not be restricted to the 2^{16}-space codes for which the three LSBs are 0. Naturally, in cases where M_{level} is not a power of 2, $M_{\text{out}} \neq M_{\text{min}}$, and the space of active codes is strictly less than the space of possible codes.

Quantization is a nonlinear operation on the input value.[10] As this nonlinearity complicates analysis, quantizers are often regarded statistically: it can be shown that quantizers act linearly on the input signal probability density function (PDF) [*Widrow and Kollár*, 2008, p. 61].[11] In statistical settings, the pseudo-quantization noise (PQN) model is often used. Shown in Figure 3.8, the PQN model approximates

[10]Nonlinearity is easily proved: given a DC $y[n]$ centered in a bin, adding another DC signal of as much as $\pm\Delta/2$ does not change the output. Hence the quantizer does not obey superposition. It can similarly be shown that it does not obey scaling. Therefore it is not linear.

[11]The statistical view of quantization theory is not covered in detail in the text proper. To give a brief introduction, though, assume that $y[n]$ is described by the PDF $f_{y[n]}(x)$. Let $y[n]$ be quantized by an unbounded, uniform quantizer of step size Δ. The operation of the quantizer on $f_{y[n]}(x)$

the quantizer error as an additive noise signal $e[n]$ that is white, independent of the quantizer input, and uniformly distributed with PDF [*Widrow and Kollár*, 2008, pp. 66–69]:

$$f_{e[n]}(x) = \begin{cases} \dfrac{1}{\Delta} & , \quad -\dfrac{\Delta}{2} < x < \dfrac{\Delta}{2} \\ 0 & , \quad \text{otherwise} \end{cases} \tag{3.14}$$

produces the output PDF $f_{y'[n]}(x)$ where [*Widrow and Kollár*, 2008, p. 62]:

$$f_{y'[n]}(x) = \sum_{r=-\infty}^{\infty} \delta(x - r\Delta) \int_{r\Delta-(\Delta/2)}^{r\Delta+(\Delta/2)} f_{y[n]}(x)\, \mathrm{d}x$$

that is, $f_{y'[n]}(x)$ is an impulse train, each $\delta(x)$ of which is weighted by the integral of $f_{y[n]}(x)$ over the appropriate bin. This area sampling can be modeled as a convolution of $f_{y[n]}(x)$ with a rectangle function, followed by multiplication with an impulse train. Recall that the characteristic function $\Phi_x(u)$ of a PDF $f_x(x)$ is defined as:

$$\Phi_x(u) = \int_{-\infty}^{\infty} f_x(x) e^{iux}\, \mathrm{d}x$$

The similarity to the Fourier transform is obvious. It can be shown that the characteristic function of the quantizer output $\Phi_{y'[n]}(u)$ is [*Widrow and Kollár*, 2008, p. 65]:

$$\Phi_{y'[n]}(u) = \sum_{l=-\infty}^{\infty} \Phi_{y[n]}\left(u + l\frac{2\pi}{\Delta}\right) \operatorname{sinc}\left(\frac{\Delta}{2}\left(u + l\frac{2\pi}{\Delta}\right)\right)$$

Similar to sampling theory, $\Phi_{y'[n]}(u)$ consists of a series of replicas at intervals of $2\pi/\Delta$, each replica being a sinc-weighted $\Phi_{y[n]}(u)$. Once again, replica overlap introduces distortion, preventing $\Phi_{y[n]}(u)$ recovery. The analogue to the Nyquist-Shannon sampling theorem, then, is called Quantizing Theorem I (QT-I) by Widrow. QT-I states that, if $\Phi_{y[n]}(u)$ is limited to a "bandwidth" less than $(2\pi/\Delta)/2 = \pi/\Delta$, then quantization incurs no aliasing distortion and the PDF of the quantizer input can be recovered from the PDF of the quantizer output [*Widrow and Kollár*, 2008, p. 69]. As a practical matter, though, in many applications only certain statistical properties of the quantizer input, instead of its entire PDF, are desired. In this case, aliasing distortion may be permitted. In particular, recall that the m-th moment of $f_{y[n]}(x)$ is [*Widrow and Kollár*, 2008, pp. 33–35]:

$$\mathrm{E}\{(y[n])^m\} = \int_{-\infty}^{\infty} x^m f_{y[n]}(x)\, \mathrm{d}x$$

The m-th moment can be derived from the characteristic function as [*Widrow and Kollár*, 2008, p. 35]:

$$\mathrm{E}\{(y[n])^m\} = \frac{1}{i^m} \frac{\mathrm{d}^m}{\mathrm{d}u^m} \Phi_{y[n]}(u)\bigg|_{u=0}$$

Thus, if $\Phi_{y'[n]}(u)$ at $u = 0$ is not distorted by aliasing, then the moments of $\Phi_{y[n]}(u)$ can be recovered. Hence for moment recovery, $\Phi_{y[n]}(u)$ only needs to be limited to a bandwidth of $2\pi/\Delta$. This result is called Widrow's Quantizing Theorem II (QT-II) [*Widrow and Kollár*, 2008, pp. 69–70].

Under this model, it can be shown that the second moment of the quantizer output $\mathrm{E}\{(y'[n])^2\}$ is:

$$\mathrm{E}\{(y'[n])^2\} = \mathrm{E}\{(y[n])^2\} + \frac{\Delta^2}{12} \tag{3.15}$$

Equation (3.15) leads to the popular statement that a quantizer adds a quantization noise power Q_{noise} of:

$$Q_{\mathrm{noise}} = \frac{\Delta^2}{12} \tag{3.16}$$

to the input signal.

Though the PQN model only holds under certain conditions of the input signal statistics and quantizer characteristic,[12] it is often applied in cases where these conditions are not strictly met, and often to surprising utility.[13] Unfortunately, in

[12]To continue the statistical view of quantization, it can be shown that the PQN model accurately predicts the PDF of $e_{\mathrm{Q}}[n]$ and various moments of the quantizer system (including cross moments, such as the correlation of $y[n]$ and $e_{\mathrm{Q}}[n]$, or of $y'[n]$ and $e_{\mathrm{Q}}[n]$) when QT-II is satisfied [*Widrow and Kollár*, 2008, pp. 162–163]. Furthermore, if QT-II holds, then Sheppard's corrections allow estimation of the input moments of $f_{y[n]}(x)$ from the output moments of $f_{y'[n]}(x)$ as [*Widrow and Kollár*, 2008, pp. 80–84]:

$$\begin{aligned}
\mathrm{E}\{y[n]\} &= \mathrm{E}\{y'[n]\} - 0 \\
\mathrm{E}\{(y[n])^2\} &= \mathrm{E}\{(y'[n])^2\} - \frac{\Delta^2}{12} \\
\mathrm{E}\{(y[n])^3\} &= \mathrm{E}\{(y'[n])^3\} - \frac{\Delta^2}{4}\mathrm{E}\{y'[n]\} \\
&\vdots
\end{aligned}$$

and so on. In this light, Equation (3.15) simply expresses Sheppard's correction to the second moment. Under less stringent conditions on $\Phi_{y[n]}(u)$, such as those required by Quantizing Theorems QT-III and QT-IV (not covered here), select portions, but not the whole, of the PQN model are accurate [*Widrow and Kollár*, 2008, pp. 163–164].

[13]For example, the PQN model is often applied to bounded quantizers. In general, the PQN model is not applicable to such quantizers: for one thing, if the quantizer saturates, $e_{\mathrm{Q}}[n]$ cannot be bounded to $\pm\Delta/2$ and Equation (3.14) is obviously invalid. Nonetheless, for bounded quantizers it can be shown that so long as 1) the quantizer input does not overrange, 2) M_{level} is asymptotically large, 3) Δ is asymptotically small, and 4) the joint PDF of the quantizer input at different times is smooth, then $e_{\mathrm{Q}}[n]$ is approximately 1) an independent and identically distributed (i.i.d.) process (and hence white) and 2) uniformly distributed as per the PQN model [*Gray*, 1997, p. 48]. It can further be shown that the quantizer input and $e_{\mathrm{Q}}[n]$ are approximately uncorrelated, but only under certain conditions of the input PDF near the quantizer bounds [*Gray*, 1997, p. 50]. Hence, while it does not strictly hold, under these conditions the PQN model is still approximately correct for bounded quantizers.

plasma wave receiver applications such a statistical view is of limited utility. Typically, statistics describe the long-term properties of signals: there is no guarantee that the signal displays such properties over the short term [*Widrow and Kollár*, 2008, App. E]. However, plasma wave analysis often investigates impulsive signals over short time intervals, often as little as 10 ms (see Section 3.2.3). In these cases, distortion from the quantizer nonlinearity is usually most limiting, a behavior not captured under the PQN model. Hence, the statistical view does not well characterize quantizer performance in these scenarios and other methods that explicitly describe quantizer nonlinearity are needed instead. To this end, the next section introduces a variety of converter metrics, many of which characterize the quantizer nonlinearity.

3.2 Converter Metrics

While many metrics are used to describe real world converter performance, the following subsections presents only those most pertinent for describing the SVADC-1. A more complete list of metrics can be found in IEEE Standard 1241-2000 [*IEEE Std 1241-2000*, 2000].[14]

A note on terminology: since a converter contains both sampling and quantizing functions, terms used in describing each are applied to the converter as a whole. Hence, the input-output characteristic of a converter means the input-output characteristic of the quantizing aspect of the converter.[15]

3.2.1 Transition Level Metrics

One strategy for assessing converter performance is to describe the input-output characteristic. Typically the characteristic is described by its transition levels $T[m]$, with $T[m]$ being the transition level between the $(m-1)$-th and m-th bins, for

[14] Many of the metrics described here accord with IEEE Standard 1241-2000 [*IEEE Std 1241-2000*, 2000]. However, some do differ: these differences are reported in footnote.

[15] To be sure, the sampling and quantizing functions do interact within a converter, for example, harmonic distortion can be incurred through the inherent nonlinearity of the quantizer, or the signal-dependent aspect of the sampling switch, or both. Hence addressing the sampling or quantizing function of the converter as a whole is justified by more than simply semantics.

$1 < m < M_{\text{level}}$.[16] An example of this indexing is shown in Figure 3.9 (at left) for a 3-bit converter. The $T[m]$ can be measured through a variety of techniques (as discussed in more detail in Section 6.2.2). The measured transition levels $T_{\text{meas}}[m]$ can then be compared to their ideal values $T_{\text{ideal}}[m]$. By convention, the $T_{\text{meas}}[m]$ are corrected for gain and offset (as these quantities constitute but linear deviations) so that the comparison has form:

$$aT_{\text{meas}}[m] + b + \epsilon[m] = T_{\text{ideal}}[m] \tag{3.17}$$

where a and b are the gain and offset correction, respectively, and $\epsilon[m]$ is the residual error after correction. The gain and offset can also be defined in a variety of ways (as discussed in more detail in Section 6.2.2). The transition levels after gain and offset correction are the corrected transition levels T_{corr}[m]:

$$T_{\text{corr}}[m] = aT_{\text{meas}}[m] + b \tag{3.18}$$

These corrected transition levels are used to define the metrics:

- *Differential nonlinearity (DNL)*
  The differential nonlinearity (DNL) is defined as the deviation of a bin width from its ideal width, normalized by the ideal width. For a uniform quantizer of M_{level} bins, the ideal width is Δ for all bins. The DNL for the m-th bin is then:

$$\text{DNL}[m] = \frac{T_{\text{corr}}[m+1] - T_{\text{corr}}[m] - \Delta}{\Delta} \quad , \quad 1 < m < M_{\text{level}} - 2 \tag{3.19}$$

  DNL is not defined for the first and last bins as these bins are unbounded [*IEEE Std 1241-2000*, 2000, p. 45]. For a perfectly uniform converter the DNL is 0.

- *Integral nonlinearity (INL)*
  The integral nonlinearity (INL) is defined as the deviation of the corrected

[16] While this $T[m]$ definition is mathematically accurate, in practice measurement noise complicates assessing the exact point between two bins. Hence, $T[m]$ is alternately defined as "the input value that causes 50% of the output codes to be greater than or equal to the upper code of the transition, and 50% to be less than the upper code of the transition" [*IEEE Std 1241-2000*, 2000, p. 8].

transition levels $T_{\text{corr}}[m]$ from their ideals $T_{\text{ideal}}[m]$, normalized by the ideal bin width. Again, assuming a uniform quantizer of M_{level} bins with ideal bin width Δ, the INL of the m-th transition level is:

$$\text{INL}[m] = \frac{\epsilon[m]}{\Delta} \quad , \quad 1 < m < M_{\text{level}} - 1 \tag{3.20}$$

For a perfectly uniform converter the INL is 0. While technically attributed to the transition levels, in common parlance INL is often attached to the output codes themselves: typically, INL$[m]$ is the INL of the m-th output code (that is, the m-th bin adopts the INL of its lower transition level).

The indexing of DNL and INL is also illustrated in Figure 3.9 (at right). Often, to summarize the metrics, only the maximum and minimum values of DNL and INL over all valid m are quoted. Also, although DNL and INL as defined above adopt units of fractions of a bin width, occasionally other units are used instead.

3.2.2 Spectral Metrics

Alternately, the converter can be described in terms of its spectral performance. The spectrum of discrete-time signals is defined by the discrete Fourier transform (DFT), which is typically computed by a fast Fourier transform (FFT) algorithm. Although an FFT is thus just a computational method for calculating the DFT, the association of DFT and FFT is so strong that the two terms are commonly used interchangeably. In this dissertation, then, the term FFT refers to both the transform and its computational algorithm.

3.2.2.1 Spectrum Definition

The quantizer output spectrum is computed by an N-point FFT:

$$Y'[k] = \frac{1}{N} \sum_{n=0}^{N-1} y'[n] e^{-i2\pi kn/N} \quad , \quad k = 0, 1, \ldots, N-1 \tag{3.21}$$

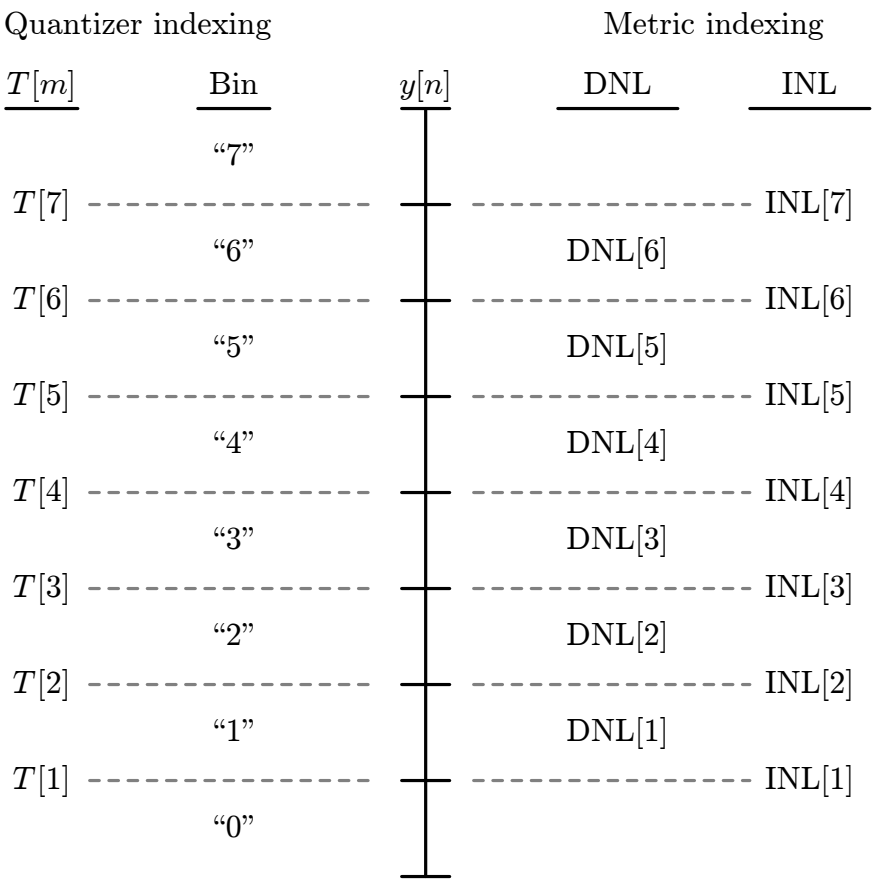

Figure 3.9: Example of indexing for both quantizer characteristics (left) and transition level metrics (right) for a 3-bit converter.

The FFT thus converts N samples of $y'[n]$ into the N-point spectrum $Y'[k]$.[17] N is called the FFT length. The $Y'[k]$ nominally represent the components of $y'[n]$ at the frequencies $f = kf_S/N$.[18] However, each $Y'[k]$ also contains nonzero contributions

[17] This FFT definition includes a $1/N$ scale factor so that coherent signal strengths (such as DC offsets and sinusoids) remain invariant to changes in N.

[18] The frequencies of the FFT, indexed by k, can be traced back to the continuous-time frequency f through the DTFT. As stated in footnote 2 of this chapter, the DTFT of a discrete-time signal $x[n]$ is defined as:

$$X(\tilde{f}) = \sum_{n=-\infty}^{\infty} x[n]e^{-i2\pi n\tilde{f}}$$

where $\tilde{f}$ is the discrete-time frequency. In value, $X(\tilde{f})$ is actually the same as the $X_s(f)$ of Equation (3.7) to within a frequency scaling factor. To derive this result, directly compute the continuous-time

from other frequency components. For this reason, the $Y'[k]$ are often referred to as bins:[19] the frequencies $f = kf_S/N$ are referred to as "on-bin" frequencies and the remainder of the frequencies are "off-bin" frequencies.

It is commonly assumed that the k-th bin represents frequency components in a bandwidth of f_S/N centered at kf_S/N: f_S/N is thus commonly identified as the "width" of the bin. In reality, though, each FFT bin represents a weighting of the entire underlying spectrum. For the k-th bin, this weighting function is given by [*Smith*, 2008, pp. 105–108]:

$$C_k(f) = \frac{1}{N} e^{-i\pi T_S[f-(k/N)f_S](N-1)} \frac{\sin(\pi N T_S[f-(k/N)f_S])}{\sin(\pi T_S[f-(k/N)f_S])} \tag{3.22}$$

As an example, Figure 3.10 plots $|C_3(f)|$ for a 16-point FFT. Figure 3.10 thus shows that the third FFT bin is in fact composed of many more off-bin components than

Fourier transform of Equation (3.5) [*Oppenheim et al.*, 1997, pp. 536–537]. Let $\mathcal{F}\{x(t)\}$ denote the continuous-time Fourier transform of $x(t)$. Then:

$$X_s(f) = \mathcal{F}\{x_s(t)\} = \mathcal{F}\left\{\sum_{n=-\infty}^{\infty} x[n]\delta(t-nT_S)\right\} = \sum_{n=-\infty}^{\infty} x[n]e^{-i2\pi nT_S f}$$

Comparing $X_s(f)$ with $X(\tilde{f})$:

$$X_s(f) = \sum_{n=-\infty}^{\infty} x[n]e^{-i2\pi nT_S f} \quad , \quad X(\tilde{f}) = \sum_{n=-\infty}^{\infty} x[n]e^{-i2\pi n\tilde{f}}$$

it is clear that, in value, these transforms describe the same spectrum if $\tilde{f} = T_S f = f/f_S$: the discrete-time frequency $\tilde{f}$ can be thought of as the continuous-time frequency f scaled such that f_S is equal to 1. The N-point FFT spectrum, then, is often conceived of as N samples of the DTFT spectrum, evenly spaced in frequency. Comparing:

$$X(\tilde{f}) = \sum_{n=-\infty}^{\infty} x[n]e^{-i2\pi n\tilde{f}} \quad , \quad X[k] = \frac{1}{N}\sum_{n=0}^{N-1} x[n]e^{-i2\pi nk/N}$$

it is clear that, if it is assumed that $x[n]$ is nonzero only over $n \in [0, N-1]$ (which is not entirely unrealistic in practice as only finite amounts of data can be computed at a time), then the $X[k]$ are indeed the samples $X(\tilde{f} = k/N)$ (to within a $1/N$ normalization of the spectrum), and $X[k] = X(\tilde{f} = k/N) = X_s(f = kf_S/N)$.

[19] The term "bin" here is slightly unfortunate, in that the divisions of a quantizer are also called bins. However, it should be obvious from the context as to which is meant. When confusion is nonetheless likely, the terms "frequency bin" or "FFT bin", and "quantizer bin", are used instead.

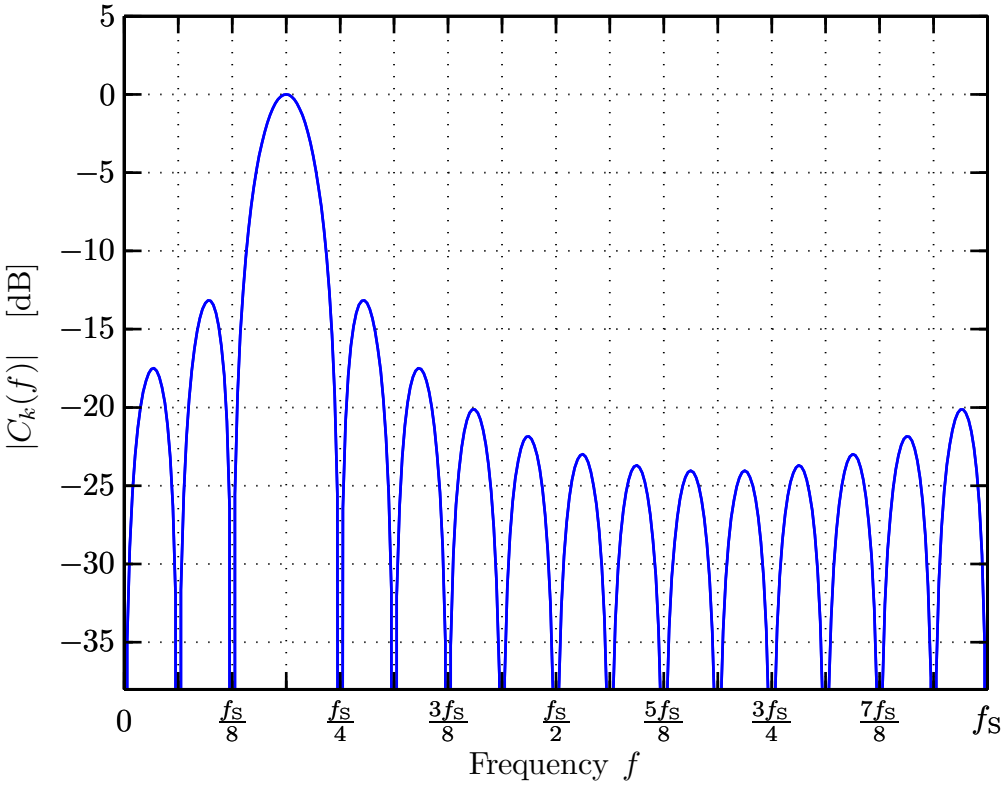

Figure 3.10: Magnitude of the frequency contribution $C_k(f)$ over continuous-time frequency f for the $k=3$ bin of a 16-point FFT. Magnitude plotted in dB.

just those components over a bandwidth of $f_S/16$ centered at $3f_S/16$. Nonetheless, as it is approximately true, the concept of a bin width of f_S/N is still used to describe FFT bin bandwidth in this dissertation.

Another consequence of Equation (3.22) is frequency smearing. It can be shown that, due to orthogonality of the $e^{-i2\pi kn/N}$ [*Smith*, 2008, p. 106]:

$$|C_k(f=mf_S/N)| = \begin{cases} 1 & , \quad m=k \\ 0 & , \quad \text{otherwise} \end{cases} \tag{3.23}$$

This orthogonality is visible in Figure 3.10 (though the nulls somewhat obscured by the dB-scaling). Equation (3.23) states that an on-bin frequency component contributes to only a single FFT bin. However, while on-bin frequency components

are thus confined, off-bin frequency components are not. In fact, since (as it can be shown) Equation (3.23) gives all of the $|C_k(f)|$ nulls, an off-bin frequency component contributes to all the bins. The result is a smearing of the power of the off-bin frequency component across the spectrum. In simulation settings this frequency smearing can be avoided by choosing on-bin frequencies. In laboratory settings, though, guaranteeing such f_S to input sinusoid frequency correspondence is more difficult and windowing techniques are used to contain the vast majority of the frequency smearing to only a few bins. And in real applications with arbitrary input signals (such as plasma wave analysis), smearing is unavoidable and becomes part of the uncertainty of definitions of the frequency-time disposition of the signal.

3.2.2.2 Metric Definitions

For the spectral metrics, a single tone sinusoid of frequency f_{in} is input to the converter and the output is collected. The output spectrum $Y'[k]$ is then computed by an N-point FFT (with possible windowing beforehand to reduce frequency smearing), and squared in magnitude to produce the power spectrum $|Y'[k]|^2$. Since $y'[n]$ is real valued, $Y'[k]$ is conjugate symmetric, and $|Y'[k]|^2$ is symmetric, about $f_S/2$ [*Oppenheim et al.*, 1999, p. 576]. Thus, only the spectral components of frequencies between 0 and $f_S/2$—that is, the components of the "positive frequencies"—are considered. Furthermore, in the case of oversampled converters, only those positive frequencies within the input signal bandwidth are considered.

As illustrated in Figure 3.11, the power spectrum is discriminated into three parts: the fundamental, the harmonics, and the nonharmonic portion. The fundamental is the signal found at f_{in}. The harmonics are the signals found at integer multiples of f_{in}. By convention $2f_{in}$ is called the second harmonic, $3f_{in}$ the third, and so on.[20] Harmonics arise from converter nonlinearities, such as quantization. Finally, what is neither fundamental nor harmonic is classified as the nonharmonic portion. The nonharmonic portion includes the noise from the circuitry implementing the converter (both white and flicker noise) and any nonharmonic tones. The latter may arise from aliasing of high frequency harmonics (as is explained in Section 3.3.2), for example.

[20] The term "first harmonic" is not used since $1f_{in}$ is already called the fundamental.

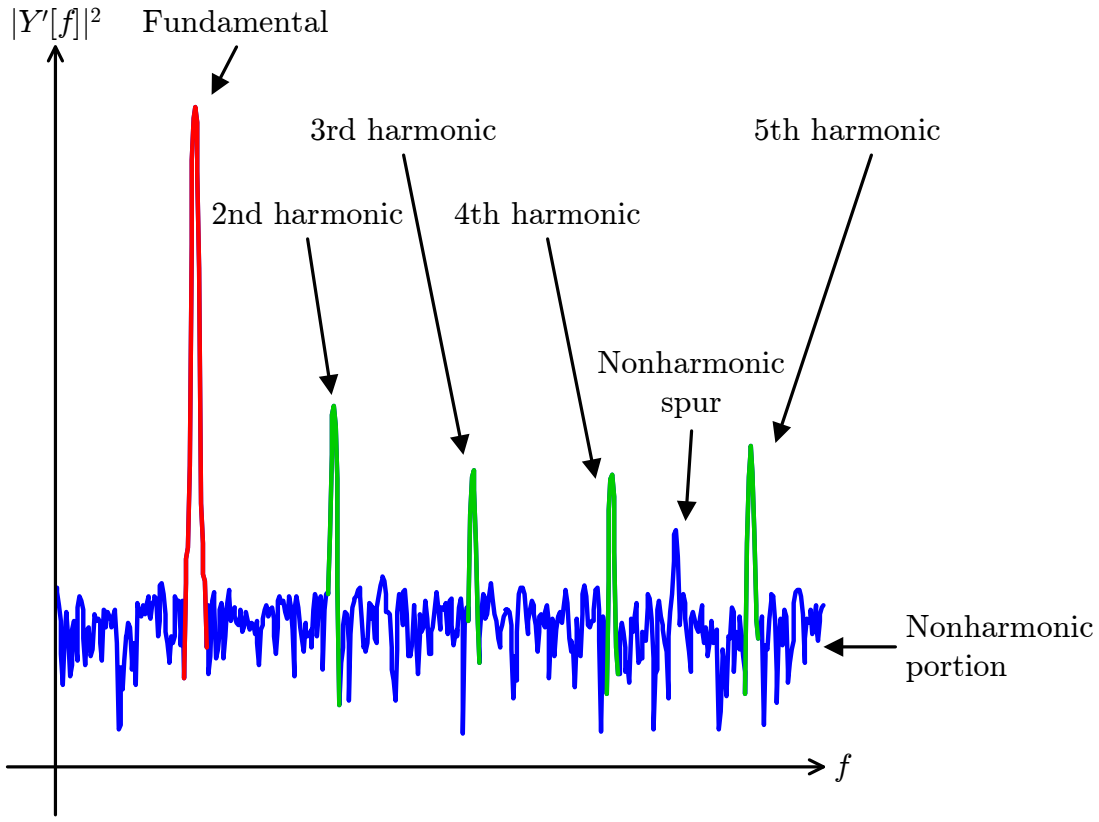

Figure 3.11: Decomposition of the power spectrum for spectral metrics. Fundamental shown in red, harmonics in green, and nonharmonic portion in blue.

Given this spectral decomposition, multiple metrics are defined.

- *Signal-to-noise ratio (SNR) and signal-to-noise-and-distortion ratio (SNDR)* The signal-to-noise ratio (SNR) and signal-to-noise-and-distortion ratio (SNDR) are the ratios of the power of the fundamental P_{fund} to various combinations of the non-fundamental portions of the signal. Specifically, defining P_{harm} as the sum of the powers of all the harmonics, and defining P_{nonharm} as the integral (over f) of the power of the nonharmonic portion, the SNR is:

$$\text{SNR} = \frac{P_{\text{fund}}}{P_{\text{nonharm}}} \tag{3.24}$$

and the SNDR is:

$$\text{SNDR} = \frac{P_{\text{fund}}}{P_{\text{nonharm}} + P_{\text{harm}}} \tag{3.25}$$

SNR and SNDR are typically given in dB ($10\log_{10}$).[21,22] Both quantities are functions of input power: typically, the SNR and SNDR initially rise with increasing input power before peaking and then falling as the input signal amplitude approaches full scale. In the case where a single value is given for SNR or SNDR, by convention the value refers to the peak SNR or peak SNDR, that is, the maximum of the quantity over input power. Finally, in this dissertation, of the two SNDR is often favored over SNR as a converter metric since it captures all of the non-fundamental portions of the signal.

- *Instantaneous dynamic range (IDR)*
  The instantaneous dynamic range (IDR) is defined as:

$$\text{IDR} = \frac{\text{Maximum signal power}}{\text{Minimum detectable signal}} \tag{3.26}$$

  where the minimum detectable signal is defined as the input power when the SNR = 0 dB. IDR is given in dB ($10\log_{10}$). In most ADC contexts, the IDR is typically called the dynamic range (DR). However, since the term "dynamic range" is easily overloaded in plasma wave receiver contexts (for example, see Section 1.2.1), in this dissertation dynamic range is used as a general descriptor of power range capability, while the formal performance metric is called IDR.

The relationship between SNDR and IDR is shown in Figure 3.12. Ideally the

[21] The SNR is equivalent to the signal-to-nonharmonic-noise-ratio (SNHR) of IEEE Standard 1241-2000 [*IEEE Std 1241-2000*, 2000, p. 13].

[22] The SNDR is conceptually equivalent to the signal-to-noise and distortion ratio (SINAD) of IEEE Standard 1241-2000 [*IEEE Std 1241-2000*, 2000, Sect. 4.5.1]. However, whereas the SNDR is defined in the frequency domain, the Standard recommends a time domain measurement of SINAD wherein the ADC output sequence $y'[n]$ is modeled as the sum of a single-frequency sinusoidal signal $s[n]$ and a non-signal, noise-and-distortion component $d[n]$. The signal $s[n]$ is extracted by a least-squares three-parameter (amplitude, phase, offset) or four-parameter (amplitude, phase, offset, frequency) fit of $y'[n]$ to a sinusoid [*IEEE Std 1241-2000*, 2000, Sect. 4.1.4]. The non-signal component $d[n]$ is then just $y'[n] - s[n]$. The ratio of the root-mean-square power of $s[n]$ to $d[n]$ is the SINAD. However, this dissertation favors the frequency-domain-based metric SNDR as SNDR more easily accommodates spectral bandlimiting in the case of oversampled conversion.

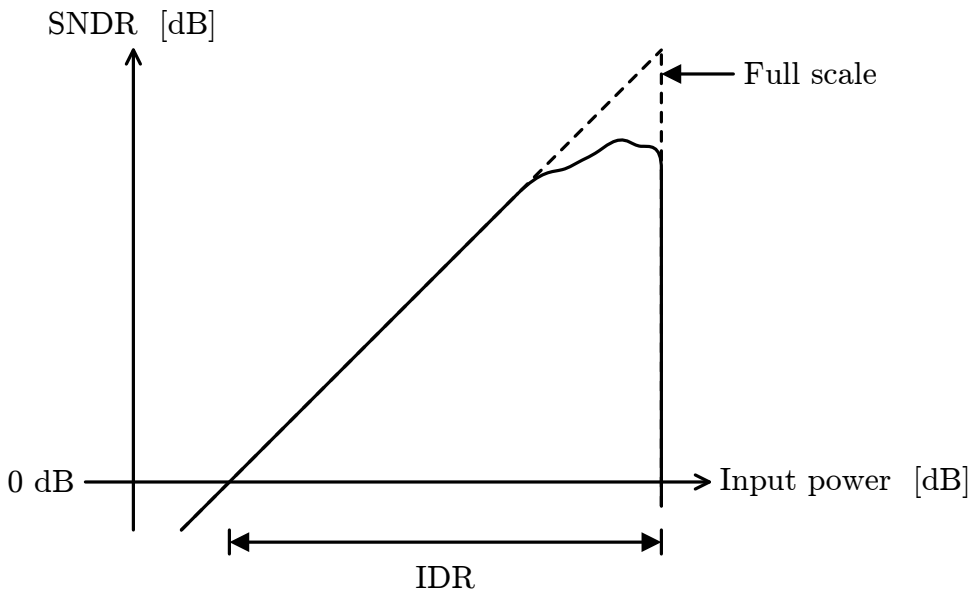

Figure 3.12: Illustration of the relationship between SNDR and IDR. Ideal SNDR shown as dashed line, typical measured SNDR shown as solid line.

SNDR rises with increasing signal power until full scale, as shown in the dashed line. However, in measured converters the SNDR often levels off and peaks before reaching full scale due to effects such as increasing harmonic distortion, as shown in the solid curve. Hence, in practice, the peak SNDR is less than the IDR. Generally, then, this dissertation favors the use of peak SNDR to characterize converter performance. Finally, in addition to being given in units of dB, the peak SNDR is often also given as an effective number of bits:

- *Effective number of bits (ENOB)*
  The effective number of bits (ENOB) is:

$$\text{ENOB} = \frac{[\text{Peak SNDR}] - 1.76}{6.02} \tag{3.27}$$

  where the peak SNDR is in dB and the ENOB is in bits.[23] Note that ENOB

[23]Equation (3.27) arises by considering the IDR of an ideal M-bit converter. The maximum input sinusoid has an amplitude of $V_{\text{FS}}/2$, or equivalently a power of $(V_{\text{FS}}/2)^2/2 = V_{\text{FS}}^2/8$. Under the PQN model, the converter introduces no distortion but, from Equation (3.16), adds a noise of power

is a general conversion: in many contexts, peak SNR and IDR are also given in ENOB.

The previous metrics consider the power of the fundamental in relation to the integrated power of the non-fundamental components of the spectrum. Another view regards the spectral performance on a frequency bin by frequency bin, or spot, basis.

- *Spurious-free dynamic range (SFDR)*
  The spurious-free dynamic range (SFDR) compares the power of the fundamental to the power of the largest non-fundamental spectral spur:

$$\text{SFDR} = \frac{[\text{power of fundamental}]}{[\text{power of largest non-fundamental spectral spur}]} \tag{3.28}$$

  The SFDR is typically given in dB ($10\log_{10}$). The largest non-fundamental spectral spur may be a harmonic, or may lie within the nonharmonic portion of the spectrum. Hence the SFDR may be decomposed into two components, the harmonic SFDR (hSFDR) and nonharmonic SFDR (nhSFDR):

$$\text{hSFDR} = \frac{[\text{power of fundamental}]}{[\text{power of largest harmonic}]} \tag{3.29}$$

  and:

$$\text{nhSFDR} = \frac{[\text{power of fundamental}]}{[\text{power of largest nonharmonic spur}]} \tag{3.30}$$

$\Delta^2/12$. Noting that $\Delta = V_{\text{FS}}/2^M$, the IDR is:

$$\text{IDR} = \frac{V_{\text{FS}}^2/8}{\Delta^2/12} = \frac{3}{2}2^{2M}$$

Converting to dB:

$$\text{IDR [dB]} = 10\log_{10}\left(\frac{3}{2}\right) + \frac{20}{\log_2(10)}M \simeq 1.76 + 6.02M$$

Hence an ideal M-bit converter achieves the IDR given above. Solving for M and substituting the peak SNDR for the IDR gives Equation (3.27). That is, based on its peak SNDR, the measured converter can be said to be equivalent to an ideal ENOB-bit converter.

The SFDR is then the worse of the two:

$$\text{SFDR} = \min\{\text{hSFDR}, \text{nhSFDR}\} \tag{3.31}$$

Similar to SNR and SNDR, the SFDRs are functions of the input signal power. In the case where a single value is given, by convention this value refers to the peak SFDR, that is, the maximum of the quantity across input power.

It should be noted that the term "nonharmonic spur" in Equation (3.30) is interpreted in a nonstandard fashion in this dissertation. Typically, nonharmonic spurs are constrained only to coherent signals [*MIL-STD-883G*, 2006, Sect. 4.4.5]. However, in this dissertation, nonharmonic spur is expanded to indicate the greatest power in a frequency bin within the nonharmonic portion of the spectrum. It thus includes not only coherent signals, but also incoherent noise. As is discussed in Section 3.3.3, the noise power in any frequency bin is dependent on the frequency width of the bin. Hence in this dissertation, the frequency bin width (or equivalently the spectral resolution in Hz/bin) is also given whenever stating an SFDR value.

3.2.3 Key Metrics

As described in Section 1.1.1, plasma waves evolve in both frequency and time and much of their analysis relies on spectrograms. To construct these spectrograms, the signal is first divided into a series of perhaps overlapping time intervals. FFTs are then taken of the signal over each time interval and the resulting spectra arranged as colored columns to create the spectrogram. The amount of overlap can be altered and additional signal processing (such as windowing) can be incorporated, but essentially spectrograms are built from a series of individual spectra.

For plasma wave analysis, then, it is of the utmost importance that each spectrum accurately depict the signal. But as the signal passes through the converter it accumulates non-signal distortions. To preserve spectral signal integrity, these distortions must be small, in particular, small enough to preserve visibility of the signals of Figure 1.10.

The metric that best captures this signal integrity requirement is SFDR. Consider the case of capturing an input composed of simultaneous strong and weak intensity phenomena. This case is often the most limiting in terms of spectral distortion in plasma wave receivers. Figure 3.13 shows such an input (in Figure 3.13(a)) converted by ADCs of both low (in Figure 3.13(b)) and high (in Figure 3.13(c)) SFDR. Both ADCs introduce distortion in the form of a noise floor and harmonics: without loss of generality, for both converters it is harmonic distortion of the strong intensity phenomenon that limits the SFDR. As depicted in Figures 3.13(b) and 3.13(c), the SFDR determines a window of signal discrimination. The top of the window is set by the strong intensity phenomenon, the bottom by the SFDR. Signals within this window are attributable to the input signal, whereas signals below it are indistinguishable from the distortion introduced by the ADC. Hence the weak intensity phenomenon is lost in the low SFDR case of Figure 3.13(b), whereas it is clearly visible in the high SFDR case of Figure 3.13(c). Thus, the SFDR characterizes signal visibility by determining the ability to see small signals in the presence of large signals.[24]

The frequency bin width must be specified when presenting SFDR (see Section 3.3.3). For plasma wave analysis, bin width is driven by the desire for good simultaneous resolution in both frequency and time. Assembling spectra, though, incurs an inherent tradeoff due to the uncertainty relation of Fourier transforms [*Bracewell*, 1986, p. 160]. Broadly, localization in one domain translates to dispersion in the other such that the product of the resolutions in both domains is constant. Hence good frequency resolution can be achieved, but the exact onset and duration of the frequency components would then not be well known. Similarly, good time resolution can be achieved by considering short time intervals, but the resulting frequency resolution over these intervals would then be coarse.

A good compromise for plasma waves is to set the FFT length to a time interval to 10 ms, corresponding to a frequency bin width of 100 Hz. As an example, consider chorus (an example of chorus was shown in the insert of Figure 1.4). Starting at a

[24]It is notable that the frequency domain representation is critical for signal visibility here: given the large difference in the strong and weak signal intensities as expected in plasma wave applications, the time domain signals of all three cases of Figure 3.13 are virtually indistinguishable.

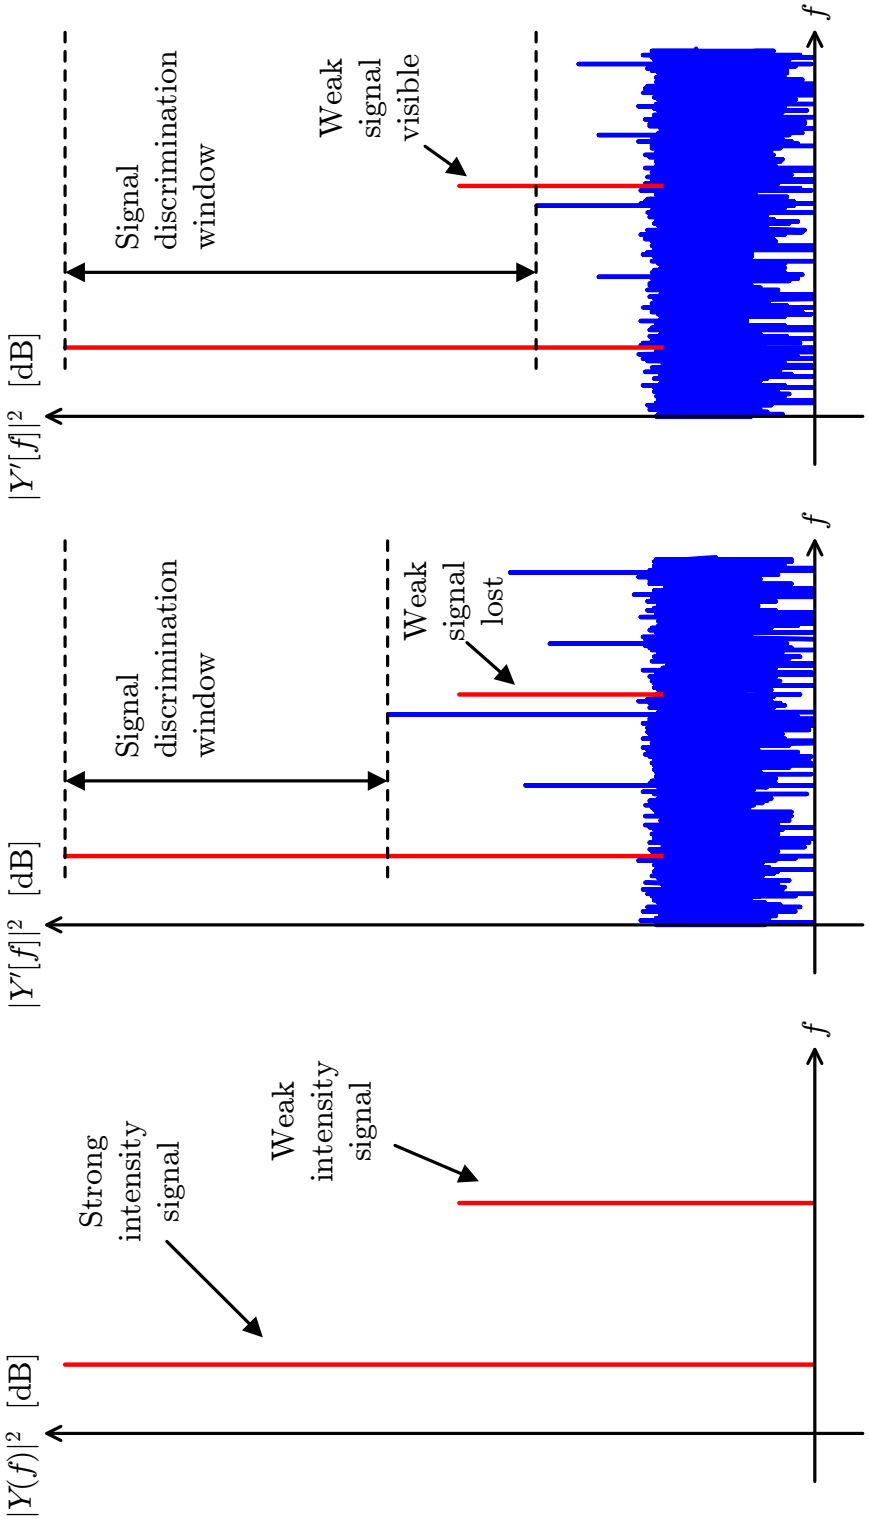

Figure 3.13: Illustration of SFDR and signal visibility, including (a) input signal, and input signal processed by converters of both (b) low SFDR and (c) high SFDR.

few kiloHertz, chorus evolves at a rate of about 1 kiloHertz per second and lasts for fractions of a second [*Inan et al.*, 1983]. To examine this phenomenon, a time interval of 1 ms is too short: as this interval captures only about a single cycle of the change, it is highly noise sensitive. On the other hand, a time interval of 100 ms is too long: as the entire phenomena only lasts on this order, its evolution with time is obscured. Hence a time interval of 10 ms strikes a good balance, providing sufficient signal cycles to belay noise while maintaining enough time resolution to clearly perceive frequency changes.

3.3 SFDR Impact

For the SVADC-1 the key metric is that it achieve at least 90-dB SFDR (as discussed in Section 1.3) assuming a frequency bin width of 100 Hz. Such an ADC can capture the entirety of the desired signal space identified in Figure 1.10. A 90-dB SFDR requirement, though, has ramifications for other aspects of the converter including sampling rate, quantizer resolution, and circuit noise.

3.3.1 Sampling Rate

An anti-aliasing filter precedes the ADC. The target plasma wave receiver of Figure 1.11 uses a 6-pole, Chebyshev Type I (ripple in the passband) low-pass filter that provides 96-dB signal suppression starting at 4 times the cutoff frequency. Factoring in this transition bandwidth and assuming a 1.08-MHz cutoff, the converter samples at 5 MS/s to reduce aliasing distortion by at least 90 dB over the 1-MHz signal bandwidth.

As 5 MS/s is 2.5 times the Nyquist rate needed to capture this 1-MHz bandwidth, the converter is oversampled. Thus, though the spectrum captures frequencies over a bandwidth of 0 Hz to 2.5 MHz, for spectral metrics only the frequency components over the 100-Hz to 1-MHz signal bandwidth are considered. To avoid confusion, presentations of converter spectral performance are often accompanied by explicit

statements of the spectral bandwidth considered in computing that performance.[25]

3.3.2 Quantizer Resolution

Being a nonlinear operation on the input signal value, quantization introduces harmonic distortion that (especially for wideband systems) diminishes the hSFDR and often limits the converter SFDR. Qualitatively, it is clear that the better the quantizer resolution—the smaller the Δ—the lesser this distortion. It is thus worthwhile to ask what Δ is necessary to achieve an SFDR specification.

Various methods have attempted to describe quantizer distortion quantitatively. Many of these methods are based on analytical techniques that rely on many assumptions and extensive mathematics, and often require heavily numerical computations to derive the ultimate hSFDR [*Martin and Secor*, 1981, Chap. 6].[26]

An alternative is to employ numerical methods. One such effort was pursued by *Pan and Abidi* [2004]. This method computes FFTs of uniformly quantized sinusoids and directly measures the SFDR. An example is shown in Figure 3.14 for a uniform 5-bit quantizer. To describe the quantizer itself, the quantizer error as a function of input value x is shown in Figure 3.14(a). The quantizer processes a full-scale sinusoid—that is, one that ranges over V_{FS}—whose frequency is chosen to lie on-bin within the first nonzero-frequency FFT bin. The resulting quantization error is shown in Figure 3.14(b). The quantized sinusoid is processed by a long FFT (typically at least (2^{M+4})-point for an M-bit quantizer): the resulting spectral magnitude for the signal of Figure 3.14(b) is shown in Figure 3.14(c) in dBc (that is, decibels-relative-to-carrier which, in this case, translates to decibels-relative-to-fundamental). Note

[25] For example, this bandwidth restriction is important in statements of SNDR since, assuming a 100-Hz to 1-MHz evaluation bandwidth, the nonharmonic portion is integrated only over 40% of the possible 2.5-MHz bandwidth.

[26] Briefly, *Martin and Secor* [1981, Chap. 6] detail four analyses (technically five, but two are the same to within the assumed quantizer input). Three rely on Fourier series expansions of the quantizer system: the first expands the quantizer output, the second the quantization error, and the third the input-output characteristic staircase itself. The fourth analysis relies on a Hermite polynomial expansion of the quantizer, based on a method by Blachman for describing the effects of arbitrary nonlinearities on signals consisting of sinusoids and Gaussian noise. All four analyses make assumptions on the quantizer and input signal to make the mathematics more tractable: this fact can limit their applicability.

that the input signal frequency is chosen to be on-bin to prevent frequency smearing, and chosen to be a relatively low frequency to (in combination with the long FFT lengths) reduce aliasing of high frequency harmonics.

Two characteristics of the output spectrum of Figure 3.14(c) are worth noting. First, unlike many nonlinear systems, harmonics do not strongly diminish with increasing frequency. Second, also unlike many nonlinear systems, the largest harmonic is not necessarily of low order. As labeled in Figure 3.14(b), the quantization error can be divided into three regions: the "bell" region (near the sinusoid's minima and maxima), the "sawtooth" region (near the sinusoid's zero crossings), and the "transition" region between the two. While the bell region contributes low frequency harmonics, it is usually the high frequency harmonics associated with the sawtooth region that limit the SFDR. Indeed, it can be shown that for an M-bit quantizer processing a sinusoid of frequency f_{in} it is typically the sawtooth fundamental at $f \simeq \pi 2^M f_{\text{in}}$ that is most limiting [*Pan and Abidi*, 2004; *Martin and Secor*, 1981, p. 6-23].

Pan and Abidi [2004] repeat this simulation for quantizers of various resolutions. The results are summarized in Figure 3.15, where the SFDR is shown as a function of the quantizer resolution. The measured SFDR is indicated by starred points and bounded by the solid and dashed lines. Given these boundaries, *Pan and Abidi* [2004] propose expressing the SFDR of a uniform M-bit quantizer as:

$$\text{SFDR} = 9.03M - \alpha(M) \tag{3.32}$$

where $\alpha(M)$ is an empirical quantity ranging from 0 dB to 6 dB for $M \in [1, 12]$ [*Pan and Abidi*, 2004].[27] Figure 3.15 shows that achieving 90-dB SFDR requires a converter resolution of at least 11 bits.[28]

[27] The $9M$ dependence of SFDR is explained intuitively by *Pan and Abidi* [2004]. If M increments by 1, Δ decreases by a factor of 2. Hence, under the PQN model, Q_{noise} decreases by a factor of 4, or 6 dB in power. This Q_{noise}, though, is distributed across some bandwidth. Consider: as M has incremented, the frequency of the sawtooth harmonics, $\pi 2^M$, has increased by a factor of 2. Intuitively, then, Q_{noise} is now distributed across twice the frequency bandwidth. Thus, in all the SFDR has improved by 6 dB from a direct decrease in Q_{noise}, and additionally by 3 dB from the remaining Q_{noise} being spread over twice the bandwidth, for a total decrease of 9 dB.

[28] While the focus is on SFDR, Figure 3.15 also shows the measured SNDR. As the input is full

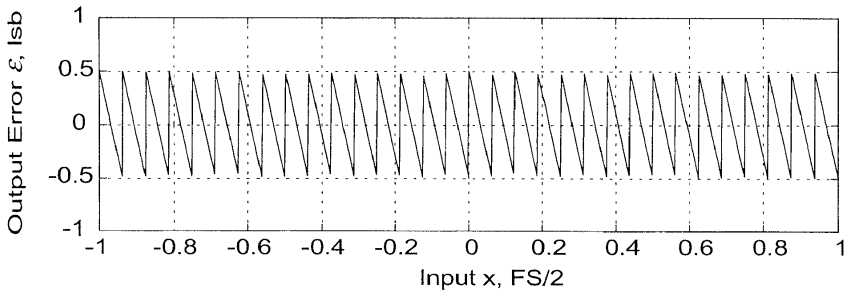

(a) Quantizer error.

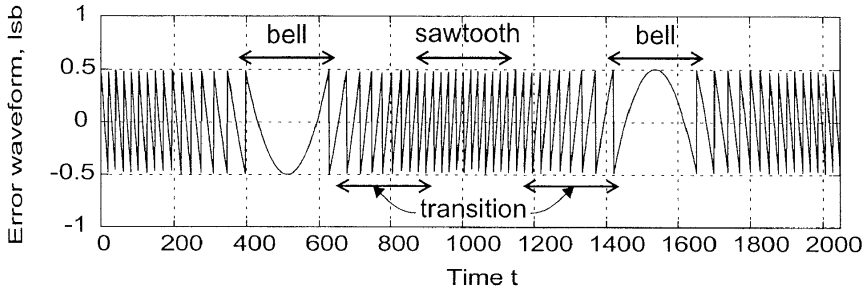

(b) Quantization error with sinusoid input.

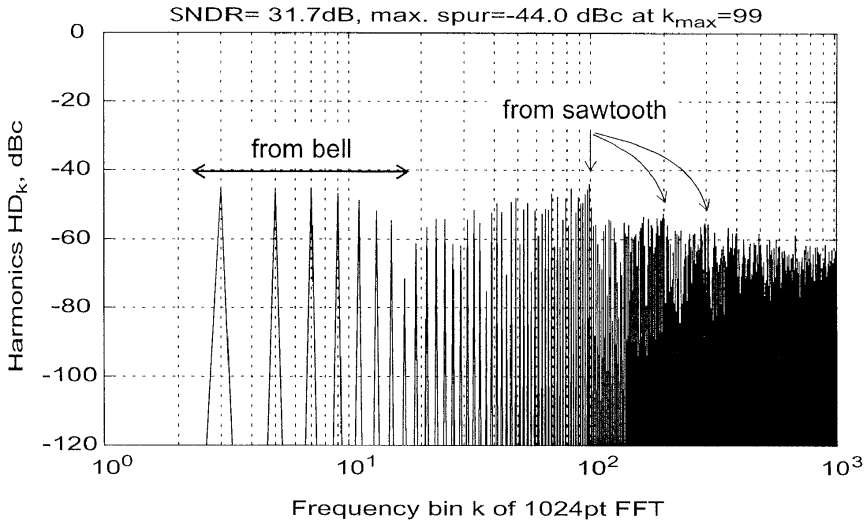

(c) Output spectrum from sinusoid input.

Figure 3.14: Demonstration of a uniform 5-bit quantizer processing a full-scale input sinusoid. (a) Quantizer error as a function of input. (b) Quantization error of quantizer processing a sinusoid. Sinusoid on-bin in first nonzero-frequency FFT bin. (c) Output spectrum of quantizer processing sinusoid of (b), sans fundamental. Reproduced in total from *Pan and Abidi* [2004].

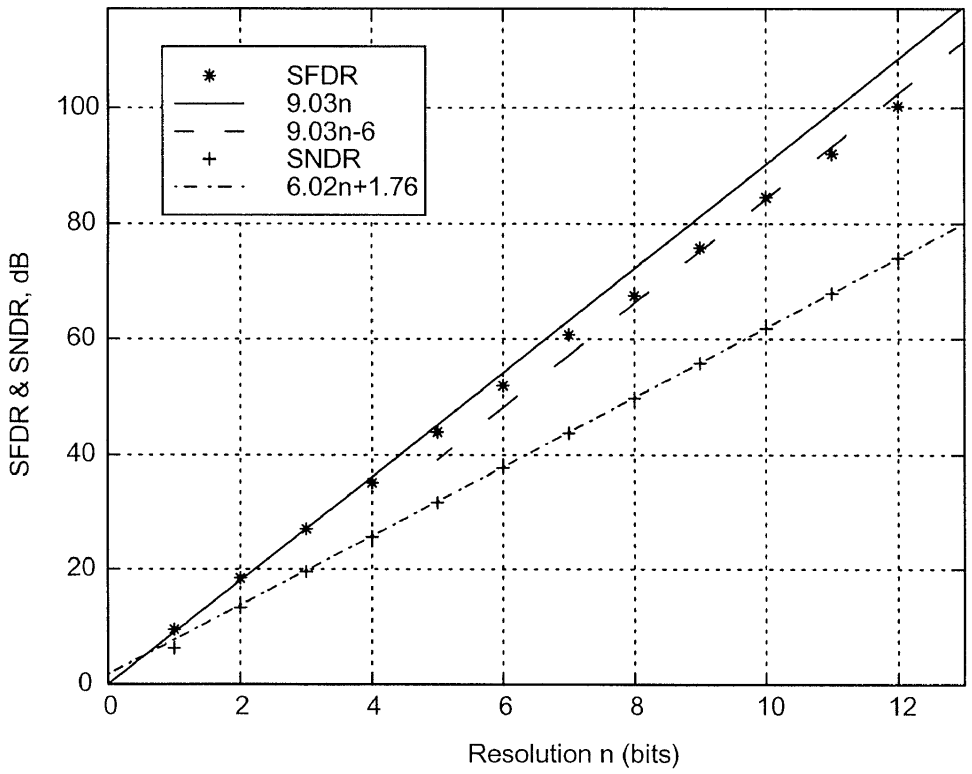

Figure 3.15: Simulated SFDR (and SNDR) achieved by a noiseless n-bit quantizer operating on a full-scale input sinusoid. Reproduced in total from *Pan and Abidi* [2004].

3.3.3 Circuit Noise

In addition to quantization noise, actual converters also contain other noise sources. In this dissertation, non-quantization noise is broadly referred to as circuit noise. Circuit noise is often input-referred and modeled as an input-independent signal $d[n]$ added prior to quantization as shown in Figure 3.16. Circuit noise can either degrade or improve the SFDR, depending on the balance of various effects.

Circuit noise aids SFDR through a dither effect. Intuitively, randomness imposed by $d[n]$ can be equivalently seen as a random disturbance of the converter transition levels. This disturbance in turn disturbs the coherence of periodic signals at the

scale, the SNDR gives the IDR of a noiseless n-bit converter. Notably, this IDR agrees with that predicted by ENOB in Equation (3.27).

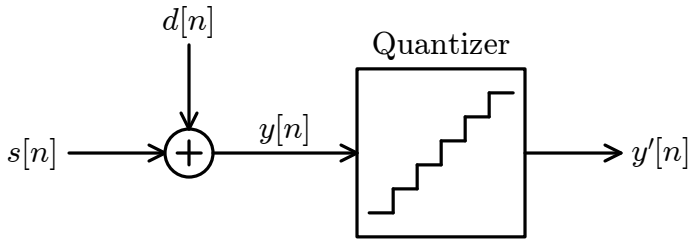

Figure 3.16: Quantizer model with input-referred circuit noise modeled by an input-independent, additive signal $d[n]$.

quantizer output: such signals—including both the fundamental and harmonics—demonstrate reduced power. Properly designed, the fundamental power reduction is fairly minimal, while the harmonic power reduction can be quite significant. Mathematically, it can be shown that, assuming the dither $d[n]$ is white, input-independent, and Gaussian-distributed with mean 0 and variance σ^2, and assuming the quantizer is unbounded and uniform, the power of the largest harmonic P_{IM} is approximately [*Martin and Secor*, 1981, p. 6-14]:

$$P_{\text{IM}}(\text{with dither}) \simeq e^{-(2\pi\sigma/\Delta)^2} P_{\text{IM}}(\text{without dither}) \tag{3.33}$$

Thus a dither σ of only a few Δ can achieve significant hSFDR enhancement.[29]

On the other hand, circuit noise degrades SFDR by raising the noise floor, reducing the nhSFDR. Assuming the noise is described by an independent, identically distributed (i.i.d.) Gaussian process with mean 0 and variance σ^2, it can be shown that the mean of the circuit noise in each FFT bin is (see Appendix H):

$$\mu_{\text{noise}} = \frac{\sigma^2}{N} \tag{3.34}$$

[29]Indeed, proper dither signal design can enforce satisfaction of QT-II, validating use of the PQN model; under the PQN model, the quantizer introduces no harmonic distortion. Assuming the model of Figure 3.16, since $s[n]$ and $d[n]$ are independent, the characteristic function of their sum is [*Leon-Garcia*, 1994, p. 271]:

$$\Phi_{y[n]}(u) = \Phi_{s[n]}(u)\Phi_{d[n]}(u)$$

where $\Phi_{s[n]}(u)$ and $\Phi_{d[n]}(u)$ are the characteristic functions of the signal and dither, respectively. In this way, dither can "anti-alias filter" the signal, bandlimiting the quantizer input $\Phi_{y[n]}(u)$ and better satisfying QT-II [*Widrow and Kollár*, 2008, Chap. 19].

where N is the FFT length.[30] The quantity μ_{noise}, though, is the average bin circuit noise: if several FFTs are taken and averaged, then the bin circuit noise tends to μ_{noise}. However, in plasma wave analysis, the rapid time evolution of the phenomena precludes such averaging. Such systems are therefore more concerned with the statistics of the bin circuit noise under a single FFT. The derivation of the statistics in the single FFT case is the subject of Appendix H, where it is shown that the circuit noise power in each FFT bin has the distribution:

$$f_{|X[k]|^2}(x) = \begin{cases} \sqrt{\dfrac{N}{2\pi\sigma^2 x}}\, e^{-Nx/(2\sigma^2)} & , \quad k = 0 \text{ or } k = N/2 \\ \dfrac{N}{\sigma^2} e^{-Nx/\sigma^2} & , \quad \text{otherwise} \end{cases} \tag{3.35}$$

assuming N is even and for $x > 0$ (naturally, $f_{|X[k]|^2}(x) = 0$ for $x < 0$ as $|X[k]|^2$ must be positive). It is further shown in Appendix H that Equation (3.35) can be extended to describe the circuit noise floor and, subsequently, the nhSFDR. Assuming the circuit noise floor is considered over η bins (encompassing only positive frequency bins and excluding the DC bin), the probability that the nhSFDR is at least x dB is then:

$$\mathrm{P}\{\text{nhSFDR} > x\} = \left[1 - e^{-\left(N/\sigma^2\right)\left(A_{\text{in}}^2/4\right)10^{-x/10}}\right]^{\eta} \tag{3.36}$$

where A_{in} is the amplitude of an on-bin input sinusoid. As the nhSFDR is probabilistic, it is impossible to completely bound it. However, from a design perspective it is sufficient that the nhSFDR be guaranteed to be greater than a bound with some (high) certainty. Expressing that certainty as a percentage $p\%$ the resulting bound is:

$$\text{nhSFDR}_p = 10\log_{10}\left(-\left(\frac{N}{\sigma^2}\right)\frac{A_{\text{in}}^2/4}{\ln(1-(p/100)^{1/\eta})}\right) \tag{3.37}$$

[30] Note that increasing N decreases μ_{noise}. This fact gives rise to the intuitive interpretation that, at least on average, the circuit noise power σ^2 is evenly spread amongst the FFT bins. Increasing N thus spreads σ^2 over more bins, resulting in less noise power per bin and improving the nhSFDR. Thus, intuitively at least, bin width must be given when specifying SFDR (see Section 3.3.3.2 for a more complete reason).

That is, the system nhSFDR is at least nhSFDR_p dB $p\%$ of the time.[31]

Equation (3.33) (combined with Figure 3.15) for dither and Equation (3.37) for the circuit noise floor together estimate the impact of circuit noise on the SFDR of a quantizer. The results are plotted in Figure 3.17. The quantizer input here is a full-scale sinusoid with additive i.i.d. Gaussian noise so that the combination displays an SNR of SNR_{FS}.[32] The first five curves are the SFDR predicted by considering dither as per Equation (3.33) in combination with the results of *Pan and Abidi* [2004] as shown in Figure 3.15 for varying quantizer resolutions. These curves show an improvement in hSFDR with increasing noise (i.e., with decreasing SNR_{FS}). The last curve is the nhSFDR predicted by considering the circuit noise floor as per Equation (3.37) (and assuming $\eta = (N/2) - 1$ and 95% certainty) and shows a loss in nhSFDR with increasing noise. The total SFDR is the lesser of that predicted by the two effects. Thus, dither effects dominate the SFDR at high SNR_{FS}, whereas the circuit noise floor dominates at low SNR_{FS}.

From a practical perspective, two results are readily apparent in Figure 3.17. First, for all quantizer resolutions shown a net SFDR gain from dither is possible. Second, for quantizers of 10 bits or more, better than 90-dB SFDR is achievable with as little as 58-dB SNR_{FS} (corresponding to an ENOB of only 9.34 bits).

3.3.3.1 Simulation of Quantizer with Circuit Noise

Fundamentally, Equations (3.33) and (3.35) are merely estimates: Equation (3.33) is such explicitly, while Equation (3.35) ignores quantization effects. Furthermore, they implicitly assume that the impact of circuit noise is separable into dither and noise

[31] As an aside, the nhSFDR_p can be connected to IDR: if A_{in} is full scale and saturation effects are ignored, then:

$$\text{IDR} = \frac{A_{\text{in}}^2/2}{\sigma^2}$$

and:

$$\text{nhSFDR}_p\ [\text{dB}] = \text{IDR}\ [\text{dB}] + 10\log_{10}\left(-\left(\frac{N}{2}\right)\frac{1}{\ln\left(1-(p/100)^{1/\eta}\right)}\right)$$

[32] That is, in the system of Figure 3.16, $s[n]$ is full scale and $d[n]$ chosen such that $s[n]+d[n]$ has an SNR of SNR_{FS}. Note that the SNR_{FS} is thus the SNR evaluated over the full bandwidth of the spectrum, that is, over $f \in [0, f_S/2]$.

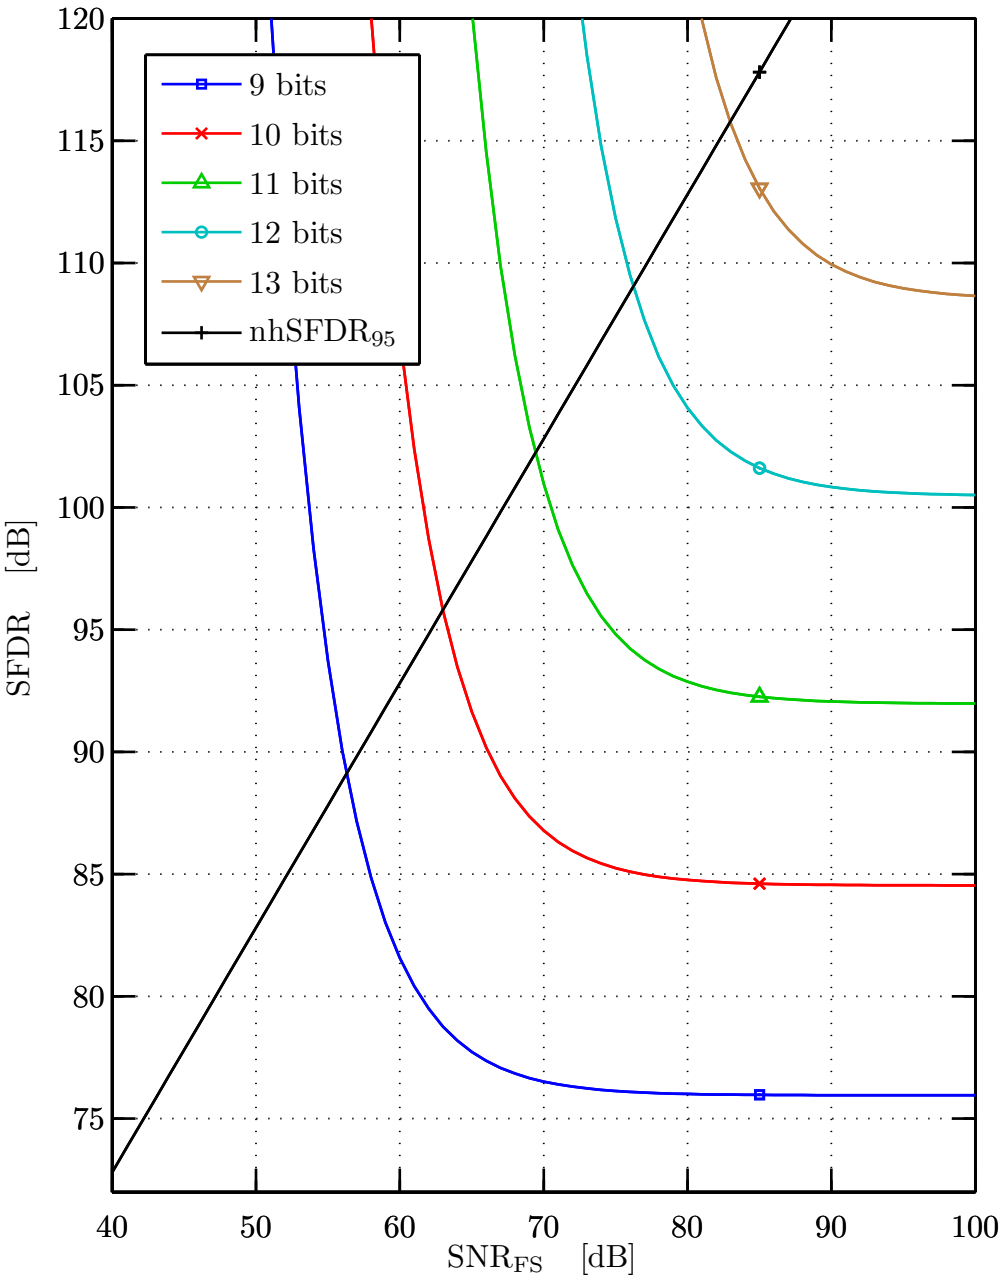

Figure 3.17: Theoretical SFDR predicted by Equation (3.33) and Figure 3.15 (describing dither), and Equation (3.37) (describing the circuit noise floor). Input signal composed of full-scale sinusoid and input-independent, i.i.d. Gaussian noise of power such that input signal achieves SNR of SNR_{FS}. For SFDR_{95}, a 50,000-point FFT is assumed (corresponding to 100-Hz bin width at 5 MS/s).

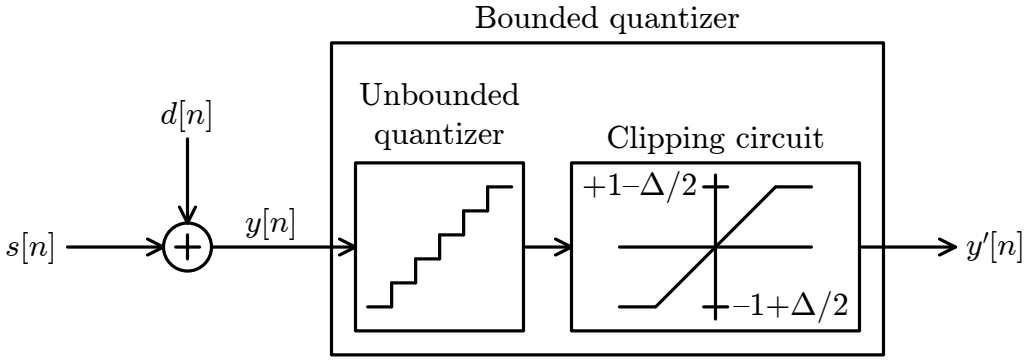

Figure 3.18: Model of bounded quantizer with circuit noise used in numerical SFDR simulations.

floor contributions. To more fully assess the impact of circuit noise on SFDR, and to better consider possible interactions between the two effects, this dissertation thus introduces new numerical simulations in the spirit of *Pan and Abidi* [2004].

The simulated system is shown in Figure 3.18. The input signal $s[n]$ is a sinusoid whose frequency is placed on-bin in the first nonzero-frequency FFT bin. To this signal is added an input-independent circuit noise $d[n]$, which is modeled as an i.i.d. Gaussian process of zero mean and preset variance. The combined signal $y[n]$ is then quantized by a bounded quantizer and the resulting $y'[n]$ passed through an N-point FFT and assessed for SFDR.

In contrast to the simulations of *Pan and Abidi* [2004], the presence of $d[n]$ means the quantizer input can now overrange. To model this effect, the quantizer is modeled as an unbounded quantizer followed by a clipping circuit. The clipping levels are set to $\pm(1-\Delta/2)$ so that the combination of quantizer and clipping circuit achieves the input-output characteristic of Figure 3.6. Clipping introduces saturation effects. In particular, given clipping the maximum SFDR is typically achieved before $s[n]$ is full scale. In all, the system thus operates under three open parameters:

1. *Quantizer resolution*

   The resolution is given in number of bits M. The resulting quantizer step size $\Delta = V_{\mathrm{FS}}/2^M = 2/2^M$.

2. *Input signal amplitude*

The amplitude is indicated in volts. Note that the input sinusoid $s[n]$ has no DC offset. A full-scale sinusoid has 1-V amplitude.

3. *Noise power*

   The noise power is specified by SNR_{FS} in dB. SNR_{FS} is the SNR of $y[n] = s[n] + d[n]$ if $s[n]$ were full scale. Put another way, the power of the circuit noise $d[n]$ is SNR_{FS} dB below the power of a full scale $s[n]$, regardless of the actual $s[n]$ amplitude.

For each triad of parameter values, 1000 Monte Carlo runs of the system of Figure 3.18 are simulated and the SFDR assessed for each run.[33]

3.3.3.2 Simulation Results

Before presenting the full results from the Monte Carlo simulations, to build intuition some individual results are shown in Figure 3.19. These output spectra, collected from an 11-bit quantizer processing the same input sinusoid under different SNR_{FS} conditions, demonstrate the impact of increasing circuit noise. In the low circuit noise power case—that is, the high SNR_{FS} case at left—the circuit noise is insignificant and the system essentially functions as a noiseless quantizer. The sawtooth fundamental, at approximately $\pi 2^{11}(100 \text{ Hz}) \simeq 643$ kHz, limits the SFDR. Increasing the circuit noise power improves SFDR through dither, as shown at center. Notably, the center spectrum is also much flatter, indicative of a whitening effect. However, continued circuit noise increase eventually raises the circuit noise floor enough to not only eliminate the dither benefit but also diminish the overall SFDR, as depicted at right.

The complete simulation results, acquired over multiple M, SNR_{FS}, and $s[n]$ amplitude values, are summarized in Figure 3.20. For each M and SNR_{FS} pair, the $s[n]$ yielding the maximum mean SFDR is determined. The mean of the Monte Carlo runs of that $s[n]$ amplitude, accompanied by error bars for the SFDR_5 and SFDR_{95} values, is then plotted. The estimates from Figure 3.17 are also included as dashed lines. Furthermore, for Figure 3.20 the SFDR is evaluated by a 50,000-point FFT

[33] Given a particular triad, though the $d[n]$ power remains the same through all 1000 runs—as dictated by the SNR_{FS}—the actual $d[n]$ sequence is different, introducing the element of randomness.

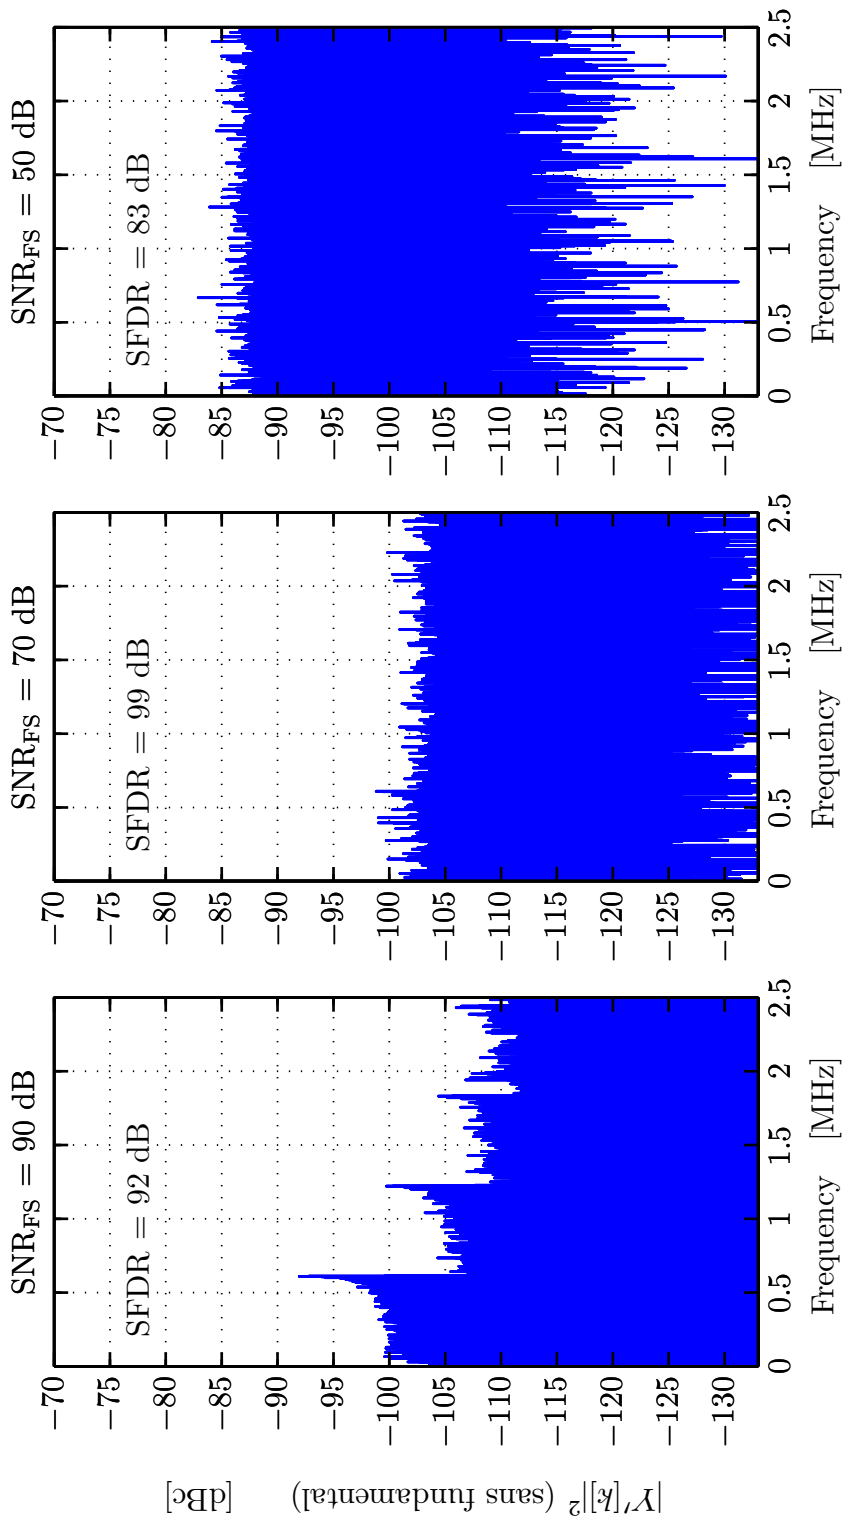

Figure 3.19: Sample output spectra of an 11-bit quantizer processing an input sinusoid under different $\mathrm{SNR}_{\mathrm{FS}}$ conditions. Input sinusoid is the same in all cases, with frequency set on-bin in first nonzero-frequency FFT bin; the first bin component is omitted to better depict the non-signal components. Spectra captured with a 50,000-point FFT, equivalent to a 100-Hz bin width at 5 MS/s.

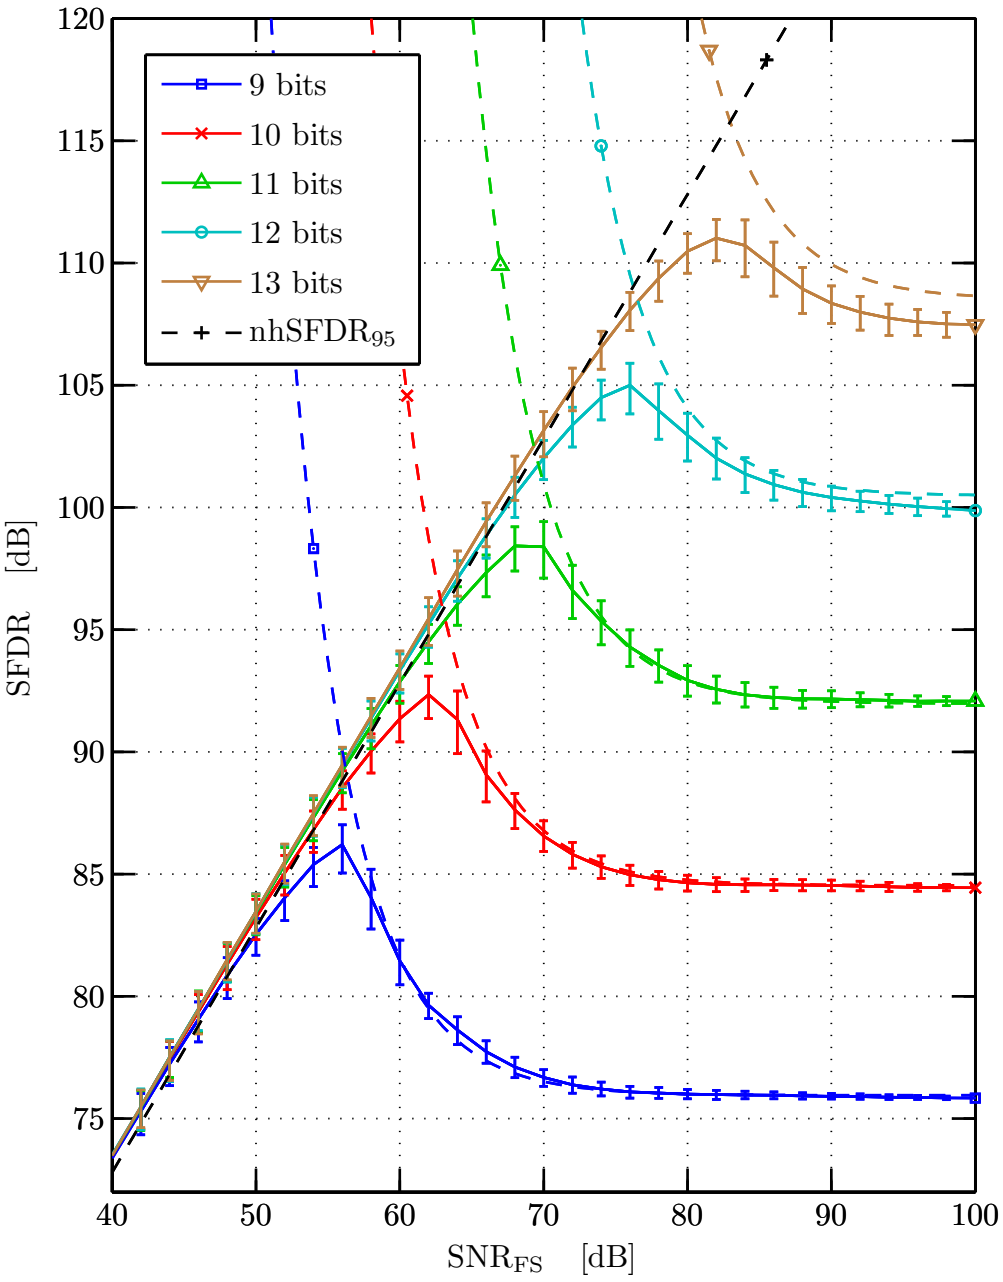

Figure 3.20: SFDR simulation results of a quantizer with circuit noise for quantizer resolutions of 9 bits to 13 bits. Numerical simulation results shown in solid curves with error bars corresponding to $SFDR_5$ and $SFDR_{95}$ values. Theoretical results (reproduced from Figure 3.17) shown as dashed curves. All results computed with a 50,000-point FFT, equivalent to 100 Hz/bin at 5 MS/s.

at 5 MS/s for a 100-Hz bin width in accordance with the SVADC-1 specification of Table 3.1. Notably, the results are strongly FFT length dependent as demonstrated in Figure 3.21, which shows results from 2,000-point, 50,000-point, and 200,000-point FFTs from left to right (the 50,000-point case of Figure 3.21 is the same as that of Figure 3.20).

Both Figures 3.20 and 3.21 show good accord between theory and simulation in the low $\mathrm{SNR}_{\mathrm{FS}}$ regime, where the SFDR is circuit noise floor dominated. Furthermore, in this regime nhSFDR_p is a strong function of the FFT length N, as depicted in Figure 3.22, which plots the theoretical nhSFDR_p of Equation (3.37) versus N: the relationship between nhSFDR_{95} and N is monotonic with a slope of roughly 9 dB/decade. This result is expected: intuitively, given constant noise power longer FFTs should spread that noise power over more bins, lowering the noise floor and improving SFDR.

On the other hand, in the high $\mathrm{SNR}_{\mathrm{FS}}$ regime, accordance between theory and simulation holds well for low resolution quantizers, but decreases for high resolution quantizers. This effect is exacerbated at low N, as seen in Figure 3.21. The lower simulated SFDR versus higher predicted SFDR is due to increased aliasing of high frequency quantization harmonics (such as the sawtooth fundamental) adding coherently to baseband signals. To understand why the SFDR decreases with increased quantizer resolution M for fixed FFT length N, recall that the sawtooth fundamental is at $f \simeq \pi 2^M f_{\mathrm{in}}$. Hence as the quantizer resolution M increases, the quantization distortion extends to higher frequencies and, if the FFT length N is not sufficient, coherent aliasing of these high-frequency distortions decreases the SFDR. This effect is not captured in the theory because this aliasing is ignored in the simulations of *Pan and Abidi* [2004]; indeed, *Pan and Abidi* [2004] intentionally increase the FFT length to prevent its occurrence. To understand why the SFDR increases with FFT length N for fixed quantizer resolution M, note that since the input signal frequency is chosen to always lie in the first nonzero-frequency FFT bin, the sawtooth fundamental frequency can also be written as $f \simeq \pi 2^M f_{\mathrm{S}}/N$. Hence as the FFT length N increases, the input signal frequency decreases, and the quantization distortion remains confined to lower frequencies. The reduced aliasing

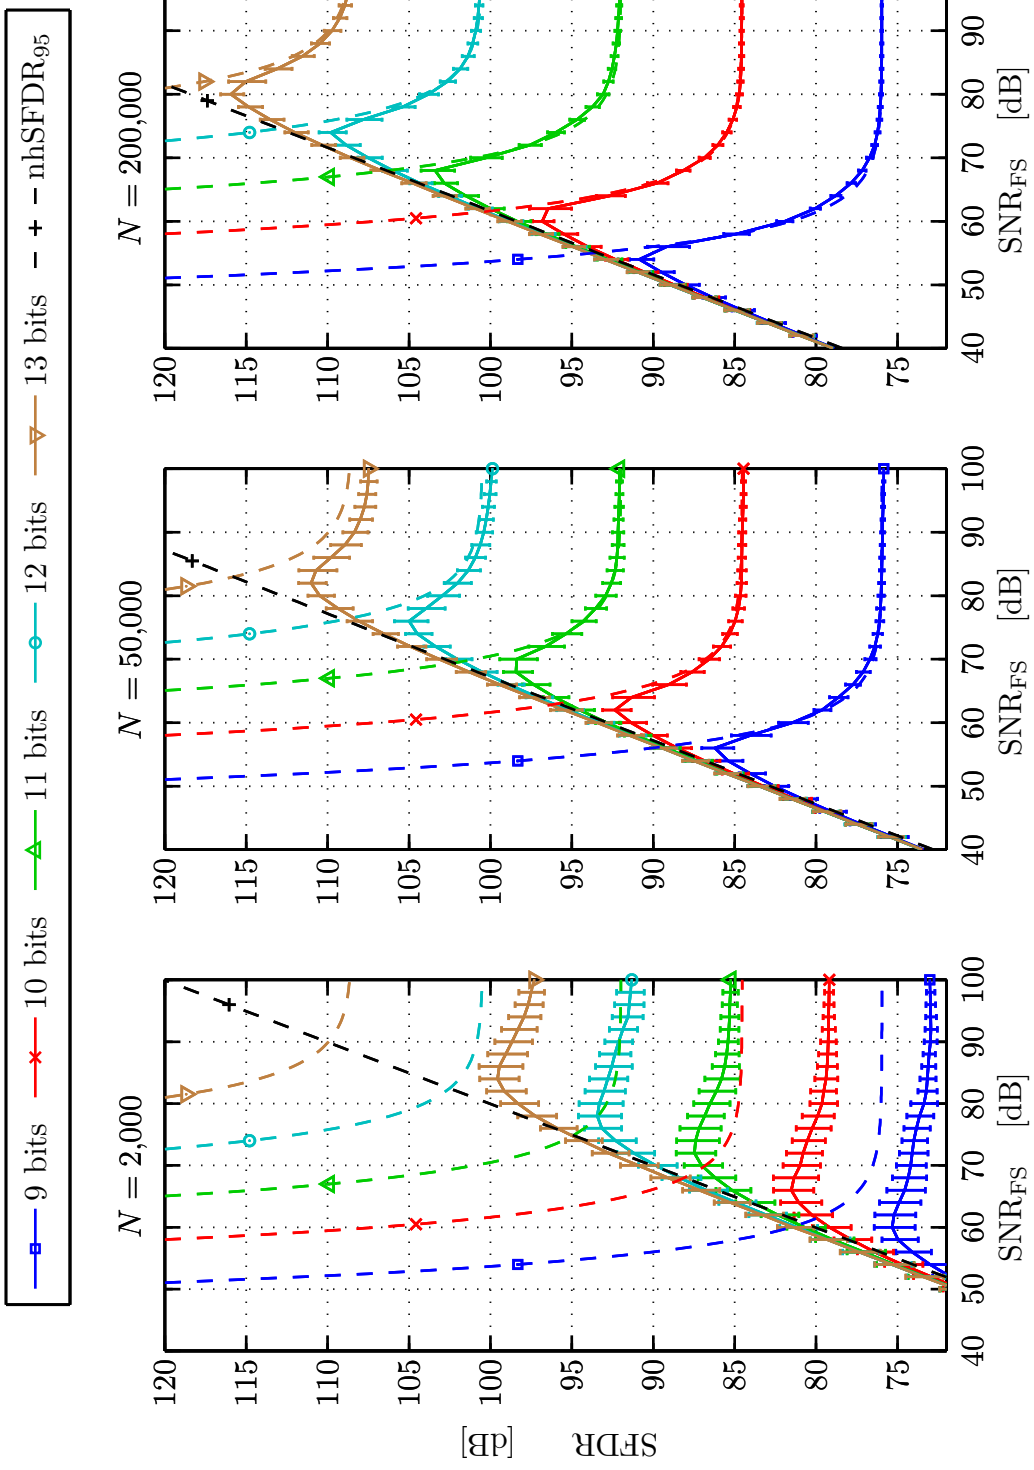

Figure 3.21: SFDR simulation results of a quantizer with circuit noise for quantizer resolutions of 9 bits to 13 bits at different FFT lengths N (equivalent to 2.5 kHz/bin, 100 Hz/bin, and 25 Hz/bin at 5 MS/s from left to right). Numerical simulation results shown in solid curves with error bars corresponding to SFDR_5 and SFDR_{95} values. Theoretical results (reproduced from Figure 3.17) shown as dashed curves. Note that the $N = 50,000$ case is the same as Figure 3.20.

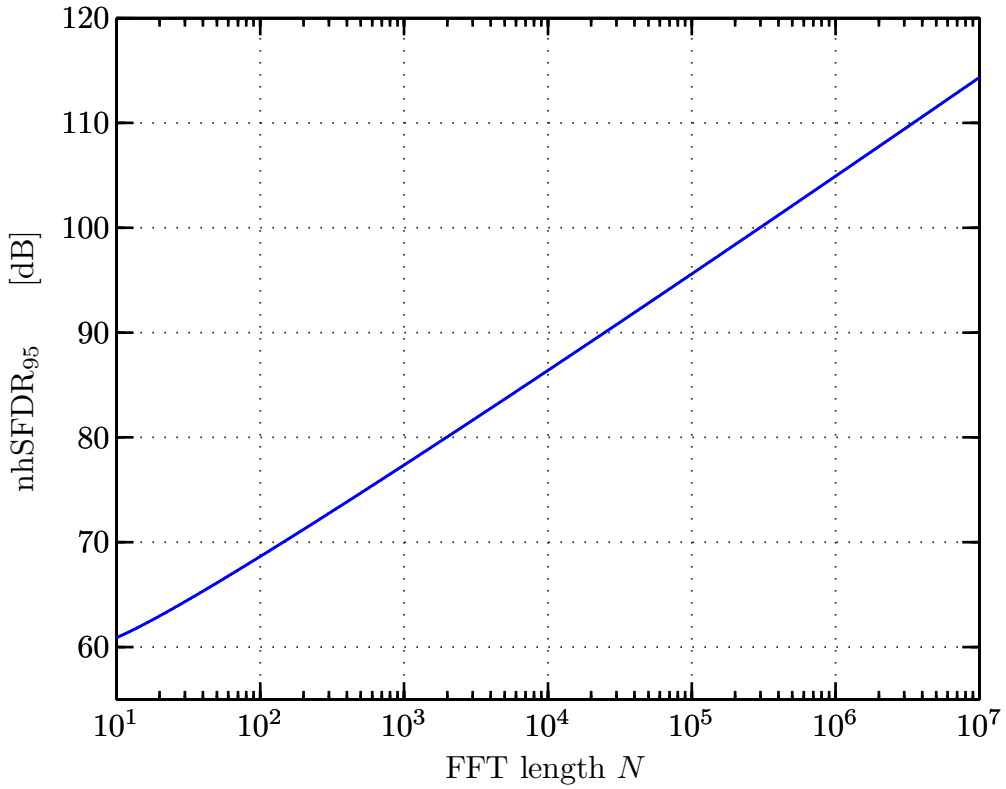

Figure 3.22: Demonstration of the relationship of nhSFDR_p to FFT length N. Parameters: $A_{\text{in}} = 1$ and $\sigma = 7.07 \times 10^{-4}$ (for $\text{SNR}_{\text{FS}} = 60$ dB), and $p = 95$ and $\eta = (N/2) - 2$.

both increases the SFDR and promotes accord with the simulations of *Pan and Abidi* [2004]. In total, then, these arguments demonstrate that in the high SNR_{FS} regime, the subsequent interrelation between quantizer resolution, FFT length, and quantization distortion is input frequency dependent and can be difficult to predict.

3.3.4 Converter Specification

Ultimately, Figure 3.20 with its 100-Hz bin width enables determination of the quantizer resolution of the SVADC-1.

In theory, with proper circuit noise levels, 90-dB SFDR is achievable with just a 10-bit converter. In practice, though, a higher resolution is desirable. First, it is

desirable to operate in the low $\mathrm{SNR_{FS}}$, circuit-noise-floor-limited regime to avoid the FFT-length-dependent distortion effects of the high $\mathrm{SNR_{FS}}$ regime. Second, the numerical simulations heretofore presented all assume uniform quantization. However, in reality, quantization is not uniform due to implementation nonidealities. The result is a reduction in SFDR, although the amount of reduction is difficult to predict.[34]

Given these concerns, the SVADC-1 targets a 12-bit resolution. This resolution provides a wide range of $\mathrm{SNR_{FS}}$—roughly between 58 dB and 66 dB (accounting for error bars)—wherein the SVADC-1 should maintain at least 90-dB SFDR while remaining noise floor dominated. Designing for this portion of Figure 3.20 also provides a healthy margin for SFDR loss from conversion nonuniformity: for example, if the implemented resolution drops to 11 bits, the SVADC-1 can still achieve greater than 90-dB SFDR. Notably, the $\mathrm{SNR_{FS}}$ over this range is significantly lower than the 74 dB typically required of a 12-bit ENOB conversion. This seeming incongruity stems from the fact that the $\mathrm{SNR_{FS}}$ requirement ultimately derives from the SFDR: the latter accommodates higher circuit noise since this noise is spread amongst all the frequency bins by the FFT.

An updated statement of the target specifications for the SVADC-1 is given in Table 3.1. The Table includes the conversion requirements derived in this chapter as well as specifications imposed by the target plasma wave receiver described in Section 1.3. The $\mathrm{SNR_{FS}}$ requirement is presented as the minimum converter SNR in the Table. Thus the SNR specification assumes that the converter is circuit noise, rather than quantization noise, limited. Also note that the SNR specification is over the full 0 Hz to 2.5 MHz bandwidth, and not just the 100 Hz to 1 MHz signal bandwidth. Finally, the Table includes radiation requirements as well: to reiterate, the converter should display no latchup, and should maintain all performance specifications through a total ionizing dose of at least 100 krad(Si).

[34]For example, systematic deviations of transition levels from their ideal positions often produce distortion worse than that anticipated if deviations of the same size had been randomly assigned [*Widrow and Kollár*, 2008, pp. 628–631]. Hence two converters of the same minimum and maximum INL, say, can nonetheless display much different distortion, depending on the exact shape of the entire INL curve.

Characteristic	Limit	Specification	Units
Receiver requirements			
Input signal range	Up to	1	V_{PP}
Input bandwidth		10^2–10^6	Hz
Sampling rate		5	MS/s
Power consumption	At most	60	mW
SFDR	At least	90	dB
SFDR requirements (100-Hz bin width)			
Resolution		12	bits
SNR (0 Hz to 2.5 MHz)		58–66	dB
Radiation requirements			
Single-event latchup (SEL)		None	
Total ionizing dose (TID)	At least	100	krad(Si)

Table 3.1: Target specifications for the SVADC-1.

3.4 Availability

With the converter required by our application and specifications well-specified by Table 3.1, the space of currently available, radiation-tolerant converters can now be assessed. To this end, Table 3.2 culls converters from the most recent JPL A/D Selection Guide [*JPL*, 2005], which is composed of converters currently approved for spaceborne applications. The Guide converters are trimmed to those that best meet the target specifications of Table 3.1: the key specifications are reiterated at the top of Table 3.2 for convenience. Of necessity, the Guide converters are also trimmed to those for which the necessary performance data is publicly available.[35]

While there are many converters that meet most of the target specifications, there are none that meet all. Perhaps notably, two converters meet all the requirements

[35]The Guide lists both custom-built converters and radiation-tolerant COTS parts, but states only limited performance data for both. For the custom-built converters, finding additional performance data can be difficult as datasheets are rare (many of these converters are not intended for sale to the public). For the COTS converters, searches are often more fruitful: if available, datasheets for the radiation-tolerant part version are used, but otherwise datasheets for the terrestrial equivalent are used instead.

Bits	Manufacturer	Part	Power	Rate	SFDR	SEL	TID
			[mW]	[MS/s]	[dB]	[MeV-cm^2/mg]	[krad(Si)]
	Required specifications		60	5	90	No latchup	100
12	Honeywell	RH9225	240	20	85	No latchup	300
14	Analog Devices	AD6644	1300	65	92	No latchup	>100
14	Analog Devices	AD6645	1500	105	93	No latchup	>100
14	Analog Devices/Maxwell	7871RP	50	0.083	88	No latchup	30–100
16	Linear Technology	LTC1604	220	0.333	93	55–70	100
16	Texas Instruments/Maxwell	7809LPTRP	150	0.1	100	No latchup	100

Table 3.2: Currently available, radiation-tolerant ADCs. Required converter key specifications included: value highlighted in red if it does not satisfy specification. Culled from the most recent JPL A/D Selection Guide [*JPL*, 2005].

but power consumption. As they also sample much faster than needed, it may be thought that these converters can be down-clocked to reduce power consumption. However, this idea does not necessarily work in practice. First, converter power consumption is usually dominated by the analog (and not the digital) circuitry and analog power consumption often does not scale with conversion rate. Second, running a converter at a reduced rate can compromise its radiation performance. For example, many converters use switched-capacitor circuitry. However (see Section 5.2.2), such circuits can be suspect to increased radiation-induced leakage in their switches, an effect that is magnified at lower clock rates. Hence, without the benefit of additional radiation testing, or a detailed understanding of the underlying circuit implementation (especially unlikely for commercial parts), stated performance is evaluated "as is". Table 3.2, then, shows that to obtain an ADC that meets the specifications of Table 3.1, the development of a custom application-specific integrated-circuit (ASIC)—namely, the SVADC-1—is required.

3.5 Architecture Choice

Analog-to-digital conversion can be achieved through a variety of different architectures. This section briefly reviews these architectures and explains the choice of a pipeline converter for the SVADC-1. It should be noted that this review is by no means exhaustive, but instead confines itself to architectures most suitable for the high-fidelity, megasample-per-second converter requirements of Table 3.1.[36]

Perhaps the most direct conversion architecture is the flash ADC, which is composed of a collection of parallel comparators each of which compares the input against a particular transition level. As all the comparators operate on the input signal simultaneously, flash ADCs can be very fast. However, as an M-bit

[36] A more complete review of converter architectures can be found in the *Data Conversion Handbook* [*Kester and Bryant*, 2005]. In addition to explaining the many architectures, the *Data Conversion Handbook* also assumes a historical bent, tracing the implementations and evolutions of each architecture through time. Each architecture discussion thus often includes an overview of the currently available commercial products that utilize that architecture (albeit with bias towards the product line of Analog Devices, Incorporated, the authors' employer), which is useful for getting a feel for the state-of-the-art performance achieved by the architecture.

conversion requires $2^M - 1$ comparators, the power and area of flash ADCs grows exponentially with resolution. Furthermore, flash ADCs are sensitive to circuit nonidealities including comparator offset and comparator cross-talk (along shared input or reference lines), limiting their resolution. It is possible to reduce the number of flash ADC comparators—and hence alleviate the rapid area and power growth with resolution—by folding and interpolation techniques. A good review of these techniques can be found in *Limotyrakis* [2004, pp. 40–45]. However, these techniques have their own limitations and hence, even with them employed, flash ADCs are typically limited to resolutions of ∼7 bits [*Limotyrakis*, 2004, p. 40].

Flash ADCs resolve all the bits of the conversion in a single cycle. In contrast, successive approximation ADCs resolve only a single bit each cycle [*Kester and Bryant*, 2005, pp. 185–190]. Essentially, successive approximation ADCs implement a binary search. The input is first compared against a reference set to the midpoint of the input range. This comparison identifies whether the input lies in the upper or lower bisection. On the next cycle the reference is set to the midpoint of the occupied bisection and the input compared to this new reference. This process of ever-finer bisection repeats, with the first cycle producing the MSB, and the last the LSB. To achieve an M-bit conversion requires M cycles. Successive approximation ADCs can be made very accurate—given the small amount of hardware required, they are readily amenable to factory calibration—but they are much slower than flash ADCs, especially at higher resolutions.

A compromise between the strict parallel nature of the flash ADC and the strict serial nature of the successive approximation ADC is a class of converter architectures called subranging ADCs [*Kester and Bryant*, 2005, pp. 190–202]. These ADCs rely on the generation of residues to perform quantization. Briefly, these ADCs perform relatively coarse quantizations of the input signal. They then determine the error of this coarse quantization. The error of the coarse quantization is subsequently itself quantized. By properly combining the coarse quantization with this quantized error the input can then be estimated to a resolution higher than each individual quantization. In practice, to relieve downstream accuracy requirements, a gain is applied to the error of the coarse quantization before it is quantized: the gained

error is called the residue. Examples of subranging ADCs include two-step ADCs, cyclic ADCs, and pipeline ADCs.[37] Of these architectures, the pipeline architecture achieves the highest throughput and hence is selected for the SVADC-1: not only can the pipeline ADC architecture achieve the 5 MS/s sampling rate required of the SVADC-1, but it can do so with margin, offsetting possible conversion speed loss due to radiation.

3.5.1 ΣΔ-Modulators

Another candidate architecture worth considering is the ΣΔ-modulator. ΣΔ-modulators use oversampling and noise shaping to achieve high SNR from very coarse quantizers. Canonically, ΣΔ-modulators embed a single-bit quantizer in a integrating feedback loop. If the input and output are taken at the right points in the loop, it can be shown that the quantization noise of the quantizer is noise-shaped by the loop with a high-pass response, whereas the input passes through the loop largely unchanged. Thus, if the input is oversampled and hence confined to low frequencies, then over the input bandwidth the quantization noise is attenuated and high SNR achieved. Appendix I addresses ΣΔ-modulators in more detail.

In general, ΣΔ-modulators are well known for their high tolerance to component mismatch and circuit nonidealities [*Brandt et al.*, 1997]. Still, historically such modulators have been mostly confined to audio applications, where the low input signal bandwidth (tens of kiloHertz) allows high oversampling factors (64 to 256) before encountering process technology limitations. These modulators not only achieve resolutions of 16–18 bits [*Vleugels et al.*, 2001], but also maintain high linearity (a key requirement in audio applications where fidelity is of exceptional concern).

[37] While all of these ADCs rely on residue generation, they differ in implementation. The two-step ADC achieves the aforementioned algorithm in two stages: one stage quantizes the input and the other quantizes the residue. Various circuit elements of the two stages can be time-shared to reduce overall hardware. The cyclic ADC collapses these two stages into a single stage: the residue is simply cycled back to the input of the stage. As cyclic ADCs consume so little hardware, they can be readily calibrated, enabling high resolution. However, they can only process one input every several cycles. In contrast, then, is the pipeline ADC, which cascades a series of residue-generating stages and then operates them in assembly line fashion. Pipeline ADCs thus achieve high throughputs—processing a new input each cycle—and hence high effective conversion rates, albeit at the cost of a small latency.

Improvements in technology, though, coupled with much active research into advanced modulator topologies and optimal circuit design, have recently produced ΣΔ-modulators with input bandwidths in the megaHertz range (e.g., *Vleugels et al.* [2001]; *Nam et al.* [2005]; *Kulchycki et al.* [2008]). Indeed, some of these newer modulators demonstrate performance very similar to that required of the SVADC-1: for example, *Nam et al.* [2005] achieved 97-dB SFDR over a 1.25-MHz input bandwidth while consuming just 44 mW analog power.

The SVADC-1 nonetheless opts for a pipeline architecture for fundamentally practical reasons. First, at the time of design, the radiation response of the BiCMOS8iED manufacturing process was not known. Hence, it was desirable that any radiation-induced performance changes be easily traceable to changes in the underlying circuitry. Given its more direct approach to conversion, such attributions were felt more straightforward in the feedforward-based pipeline architecture than the feedback-based ΣΔ-modulator. Indeed, concerns of ΣΔ-modulator loop stability (an issue in even terrestrial applications [*Adams and Schreier*, 1997]) under both total-dose parameter changes and single-event overranging events would further complicate the radiation response. Second, the digital backend of a ΣΔ-modulator, while well understood [*Norsworthy and Crochiere*, 1997; *Brandt and Wooley*, 1994], is much more involved than that of a pipeline converter, requiring additional hardware. These concerns favored adoption of a pipeline converter.

Notably, it is possible to overcome these challenges: *Edwards et al.* [1999] constructed a ΣΔ-modulator for extreme radiation environments that maintained 9.1-bit ENOB and 40-dB SFDR over a 63-kHz input bandwidth up to 23 Mrad(Si). However, this modulator was manufactured in a space-qualified, radiation-hard silicon-on-sapphire CMOS process that had been radiation-characterized.[38] In the end, then, as the SVADC-1 is the first design of a radiation-hard converter in its process, and as that process had not been radiation-characterized at design time, a more conservative approach was adopted and a pipeline architecture chosen. In the

[38]For more information on silicon-on-sapphire processes and their use in radiation applications see footnote 4 of Chapter 5.

future, though, given the device-level radiation results of Chapter 2 and the converter-level results of Chapter 6, $\Sigma\Delta$-modulators should be given strong consideration as a candidate architecture for spaceborne converters.

3.6 Conclusion

This chapter introduced analog-to-digital conversion and considered the conversion requirements of the SVADC-1.

Terms and concepts for the both the sampling and quantization aspects of conversion were established. Sampling is well described mathematically. Quantization, being a nonlinear operation on the input signal value, is not. To describe the inherent nonlinearity of conversion, then, various performance metrics were considered and SFDR found to best describe the performance required of the SVADC-1. A study of SFDR and the impact of various converter properties on its value resulted in the more complete converter description of Table 3.1. This study included new theoretical and numerical results investigating the impact of circuit noise on SFDR. The specifications of Table 3.1 were then used to consider the field of currently available, radiation-tolerant ADCs: the results, summarized in Table 3.2, show that a custom ASIC—the SVADC-1—is needed to achieve the performance requirements of the target plasma wave receiver described in Section 1.3.

After a consideration of converter architectures, a pipeline ADC architecture was chosen for the SVADC-1. The next chapter addresses the pipeline architecture in more detail. In particular, while the pipeline architecture achieves good resolution at fast rates, it is subject to SFDR loss from various circuit implementation nonidealities. To overcome these limitations, a novel self-calibration technique based on DAC differencing is introduced.

Chapter 4

Pipeline Converters

The first published CMOS pipeline converter is widely attributed to *Lewis and Gray* [1987]. The architecture itself, though, is much older, and can be traced back at least as far back as *Smith* [1956], where a converter composed of a cascaded series of binary coders was considered and implemented using operational amplifiers and vacuum tubes.[1] Nowadays, pipeline converters are known for their ability to achieve high sampling rates at medium to high resolutions by processing the analog input signal in an assembly line fashion.

This chapter discusses pipeline converters, with a view toward designing the architecture of the SVADC-1. As SFDR is a key specification for the SVADC-1, this chapter concentrates on the quantizing aspect of the architecture. It begins by introducing a single pipeline stage, then cascades a series of stages to form a complete pipeline converter. The effects of various implementation nonidealities—including overrange, offset, analog-digital mismatch, nonlinearity, and noise—are then assessed, and analog-digital mismatch found to most limit the converter SFDR. However, the analog-digital mismatch can be overcome by a variety of correction techniques. After reviewing such techniques, a set of novel DAC-differencing self-calibration techniques are proposed. One of these techniques—increased sub-ADC transition levels—is chosen and a complete converter architecture is presented.

[1] In fact, *Smith* [1956] mentions that similar methods had already been partially described in a thesis by R. P. Sallen at the Massachusetts Institute of Technology in 1949.

4.1 Architecture

A pipeline converter is formed by a series of analog stages as shown in Figure 4.1. As the analog input traverses the pipeline, each stage produces a number of bits that represent the input with the more significant bits extracted earlier in the stream. The full ensemble of bits is then digitally reconstructed to form a final digital output word that represents the input signal.

The pipeline operates entirely in sampled time. Furthermore, the stages are timed in assembly line fashion: at set times each stage outputs its result to the succeeding, downstream stage and inputs the result from the preceding, upstream stage. The stage then operates on this input (extracting bits, for example) and completes its work by the next transfer time. In this way, multiple analog input samples are processed simultaneously throughout the converter. Such timing is called "pipelining" in digital design parlance and lends the converter its name. The timing is indicated in Figure 4.1 by the sample-and-hold (S/H) circuits preceding each stage: the S/Hs establish the interstage transfer times. Pipelining enables high throughput through the converter, though at the cost of a slight delay.

With the timing understood, the rest of this chapter focuses on the quantizing aspect of the architecture. Henceforth, all signals are assumed to be discrete time, with "analog" and "digital" signals referring to signals of continuous and discrete value, respectively. This convention extends to "analog-to-digital converters" (ADCs) and "digital-to-analog converters" (DACs), which thus convert discrete-time signals between continuous-value and discrete-value representations. Also, the Chapter 3 convention wherein a primed signal indicates the digital version of the unprimed signal is maintained. Finally, it is worth reiterating the convention adopted throughout this dissertation of using parentheses $(\cdot)$ for continuous arguments, and square brackets $[\cdot]$ for discrete arguments.

To develop the pipeline architecture, this chapter first describes the operation of a single stage from both an equation-centric (i.e., algebraic) and graph-centric (i.e., geometric) perspective. Multiple stages are then cascaded to form the pipeline and the resulting converter is again considered both algebraically and geometrically.

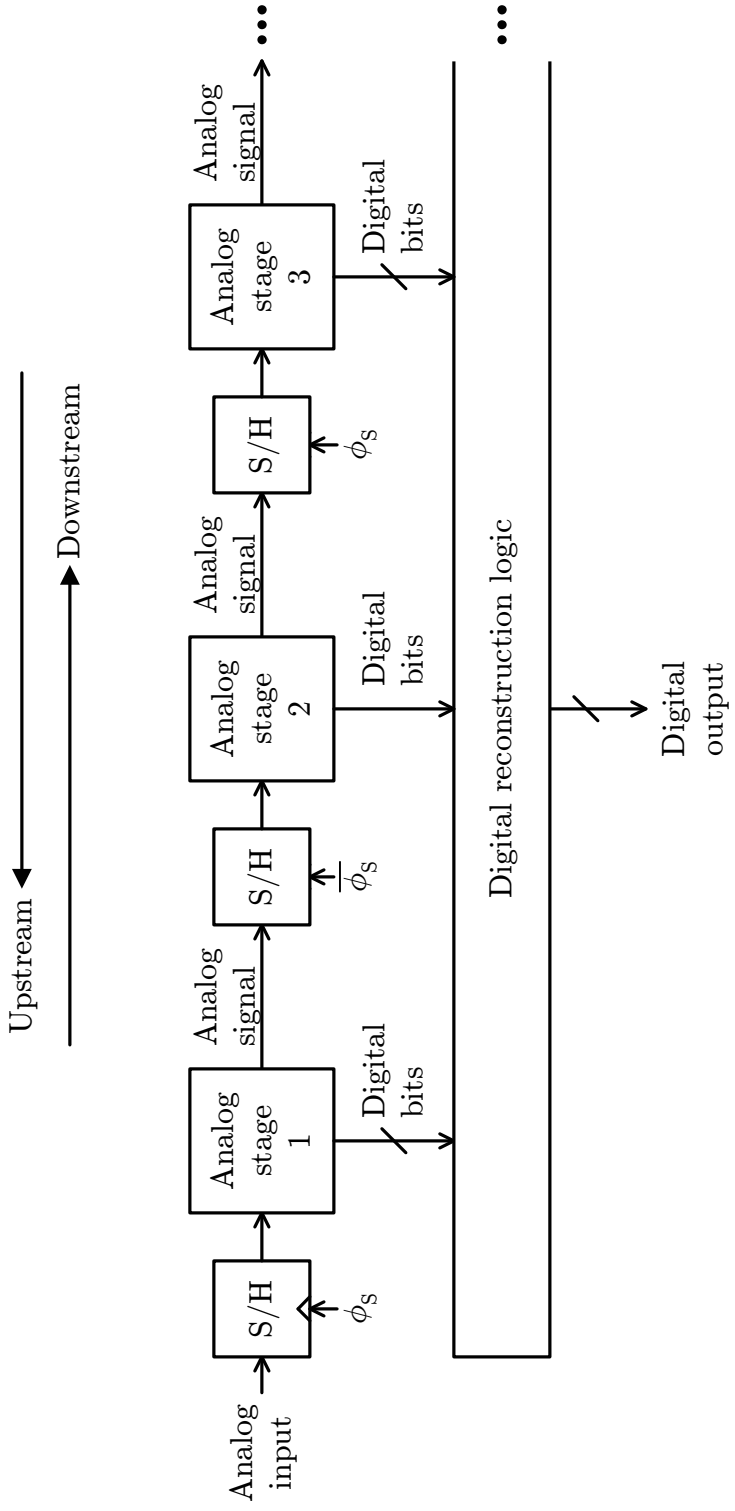

Figure 4.1: The architecture of a pipeline converter. Sample-and-hold circuits denoted as S/H.

4.1.1 Stage Operation, Algebraically

Each analog stage implements the same essential algorithm, shown towards the top of Figure 4.2. To perform quantization, the stage passes its analog input V_{in} through a coarse sub-ADC to produce the digital output B, essentially extracting the most significant bits of V_{in}. In this dissertation, B is assumed to be in unsigned integer form starting from 0.

B is taken as the DAC code input to a sub-DAC to create $V_{\text{DAC}}[B]$, an analog analogue of B. $V_{\text{DAC}}[B]$ is subtracted from V_{in} to form V_{err}:

$$V_{\text{err}} = V_{\text{in}} - V_{\text{DAC}}[B] \tag{4.1}$$

V_{err} can be thought of as the quantization error of the coarse sub-ADC. This quantization error is linearly gained by G to produce the analog residue V_{res}:

$$V_{\text{res}} = G\,V_{\text{err}} = G\,(V_{\text{in}} - V_{\text{DAC}}[B]) \tag{4.2}$$

and V_{res} is passed to a backend ADC for further conversion.

Construction of V'_{in}, a digital estimate of V_{in}, requires only B, V'_{res} (a digital estimate of V_{res} provided by the backend ADC), and knowledge of the analog stage. In fact, to compute V'_{in}, simply reverse the operation of the analog stage in the digital domain, as shown towards the bottom of Figure 4.2. Attenuation of V'_{res} by G' (a digital version of the analog gain G) produces V'_{err}, a digital estimate of V_{err}. Next, $V'_{\text{DAC}}[B]$ (a digital version of sub-DAC output voltages $V_{\text{DAC}}[B]$) is added to V'_{err} to produce:

$$V'_{\text{in}} = V'_{\text{err}} + V'_{\text{DAC}}[B] = \frac{V'_{\text{res}}}{G'} + V'_{\text{DAC}}[B] \tag{4.3}$$

V'_{in} is thus the digitization of the input analog signal V_{in}. Note that $V'_{\text{DAC}}[B]$ can be simply implemented as a lookup table on B.[2] The digital logic that performs these manipulations is called the digital reconstruction.

[2] In fact, in some cases an even simpler implementation is possible: see footnote 4 of this chapter.

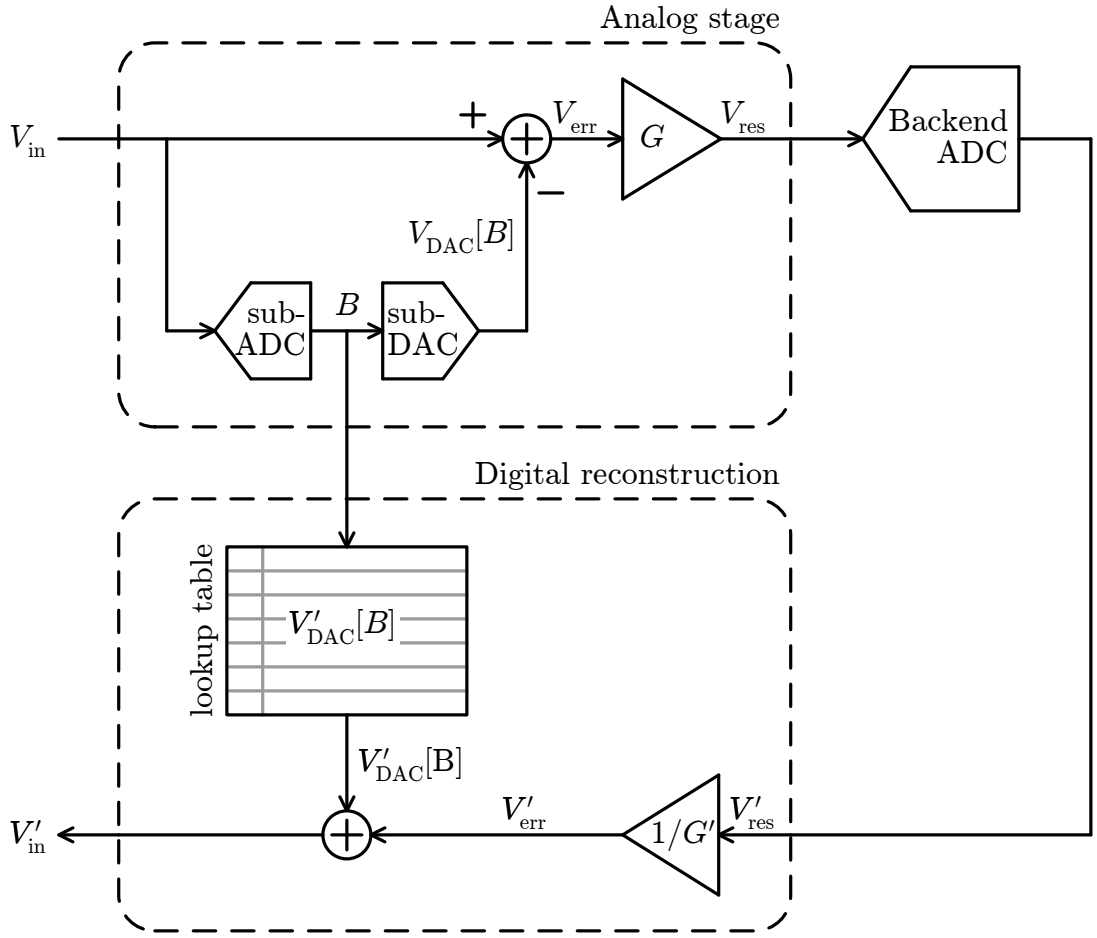

Figure 4.2: A single pipeline stage including both an analog stage and its corresponding digital reconstruction. Analog-to-digital converter denoted as ADC, digital-to-analog converter denoted as DAC.

The quantization error of this stage-and-backend-ADC system is:

$$Q_{\text{sys}} = V_{\text{in}} - V'_{\text{in}} = \left(\frac{V_{\text{res}}}{G} + V_{\text{DAC}}[B]\right) - \left(\frac{V'_{\text{res}}}{G'} + V'_{\text{DAC}}[B]\right) \tag{4.4}$$

If the analog stage and digital reconstruction match—that is, if $G = G'$ and $V_{\text{DAC}}[B] = V'_{\text{DAC}}[B]$—then Q_{sys} reduces to:

$$Q_{\text{sys}} = \frac{V_{\text{res}} - V'_{\text{res}}}{G} = \frac{Q_{\text{backend}}}{G} \tag{4.5}$$

where Q_{backend} is the quantization error of the backend ADC. Thus under perfect matching the system quantization error is the backend ADC quantization error improved by the stage gain.

Equation (4.5) explains how the system can accomplish high resolution even if the stage sub-ADC is low resolution. Intuitively, the principle is that, even though the quantization of V_{in} in a given stage is coarse, if the error of that quantization (i.e., V_{err}) is known to high accuracy, then V_{in} itself can also be known to high accuracy. This accurate estimate of the error is accomplished by the backend ADC. Hence the stage effectively defers the task of fine quantization to the backend ADC. Notably, the stage does apply a gain G to V_{err} before passing it to the backend ADC, relieving the backend ADC resolution by a factor of G. This gain contribution proves important when multiple stages are cascaded, as discussed in Section 4.1.3.

4.1.2 Stage Operation, Geometrically

An analog stage is completely specified by its input-output transfer function. The transfer function is usually presented graphically: examples are shown in Figures 4.3 and 4.4. The analog output V_{res} of the stage is plotted as a function of the analog input V_{in} and the digital output B of the stage is indicated below the axes. As expected, the transfer functions are piecewise linear functions whose segments produce a sawtooth characteristic reminiscent of the quantizer error (compare Figure 3.7).

Stages are typically named by the number of bits in their sub-ADCs:[3] Figures 4.3 and 4.4 are thus 1-bit and 2-bit stages, respectively. Furthermore, Figures 4.3 and 4.4 are examples of classical M-bit stages composed of a uniform M-bit sub-ADC, a corresponding M-bit sub-DAC, and a gain of 2^M, for $M \in \mathbb{N}$. That is, assuming the stage input range spans from $-V_{\text{REF}}$ to V_{REF} (a convention stemming from common stage implementations: see Section 4.2.3.1), in a classical M-bit stage the sub-ADC

[3] It is recognized that this nomenclature is not unique: for example, two stages employing the same sub-ADC characteristic, but sporting different gains, would share the same moniker. Nonetheless, this convention is followed throughout the literature.

Figure 4.3: Transfer function of a classical 1-bit pipeline stage. sub-ADC transition levels: 0. sub-DAC output voltages: $-\frac{1}{2}V_{\text{REF}}$ and $\frac{1}{2}V_{\text{REF}}$. Gain: 2.

Figure 4.4: Transfer function of a classical 2-bit pipeline stage. sub-ADC transition levels: $-\frac{1}{2}V_{\text{REF}}$, 0, and $\frac{1}{2}V_{\text{REF}}$. sub-DAC output voltages: $-\frac{3}{4}V_{\text{REF}}$, $-\frac{1}{4}V_{\text{REF}}$, $\frac{1}{4}V_{\text{REF}}$, and $\frac{3}{4}V_{\text{REF}}$. Gain: 4.

uniformly divides this range into 2^M partitions as per:

$$B = \begin{cases} 0 & , & V_{\text{in}} < -V_{\text{REF}} + \Delta \\ 1 & , & -V_{\text{REF}} + \Delta < V_{\text{in}} < -V_{\text{REF}} + 2\Delta \\ & \vdots & \\ 2^M - 2 & , & -V_{\text{REF}} + \big(2^M - 2\big)\Delta < V_{\text{in}} < -V_{\text{REF}} + \big(2^M - 1\big)\Delta \\ 2^M - 1 & , & -V_{\text{REF}} + \big(2^M - 1\big)\Delta < V_{\text{in}} \end{cases} \tag{4.6}$$

where:

$$\Delta = \frac{2V_{\text{REF}}}{2^M} \tag{4.7}$$

and the sub-DAC outputs the partition midpoints:

$$V_{\text{DAC}}[B] = \left[\frac{1}{2^M}(2B - 1) - 1\right] V_{\text{REF}} \tag{4.8}$$

When repeated throughout a pipeline, these types of M-bit stages implement M-ary searches (see Section 4.1.4). Of course, many other kinds of transfer functions are also possible.

Correspondences can be drawn between the transfer function geometry of Figures 4.3 and 4.4 and the stage block equations of Equations (4.6) through (4.8). These correspondences are illustrated in Figure 4.5 for an arbitrary stage transfer function. Specifically: the transition levels of the sub-ADC translate to the V_{in}-axis locations of the discontinuities between segments, the sub-DAC output voltages translate to the zero crossings of each segment, and the gain is the slope of all the segments. Given the clear correspondences, graphical transfer functions are often favored over algebraic equations when describing stages.

Graphical transfer functions also provide insight into digital reconstruction. The backend ADC estimates V_{res} but, as can be clearly seen in the transfer functions shown so far, a given V_{res} may correspond to multiple V_{in}. B, then, selects the proper V_{in} by indicating which transfer function segment is occupied by V_{in}.

In fact, there is a fully graphical view of digital reconstruction. Rearranging

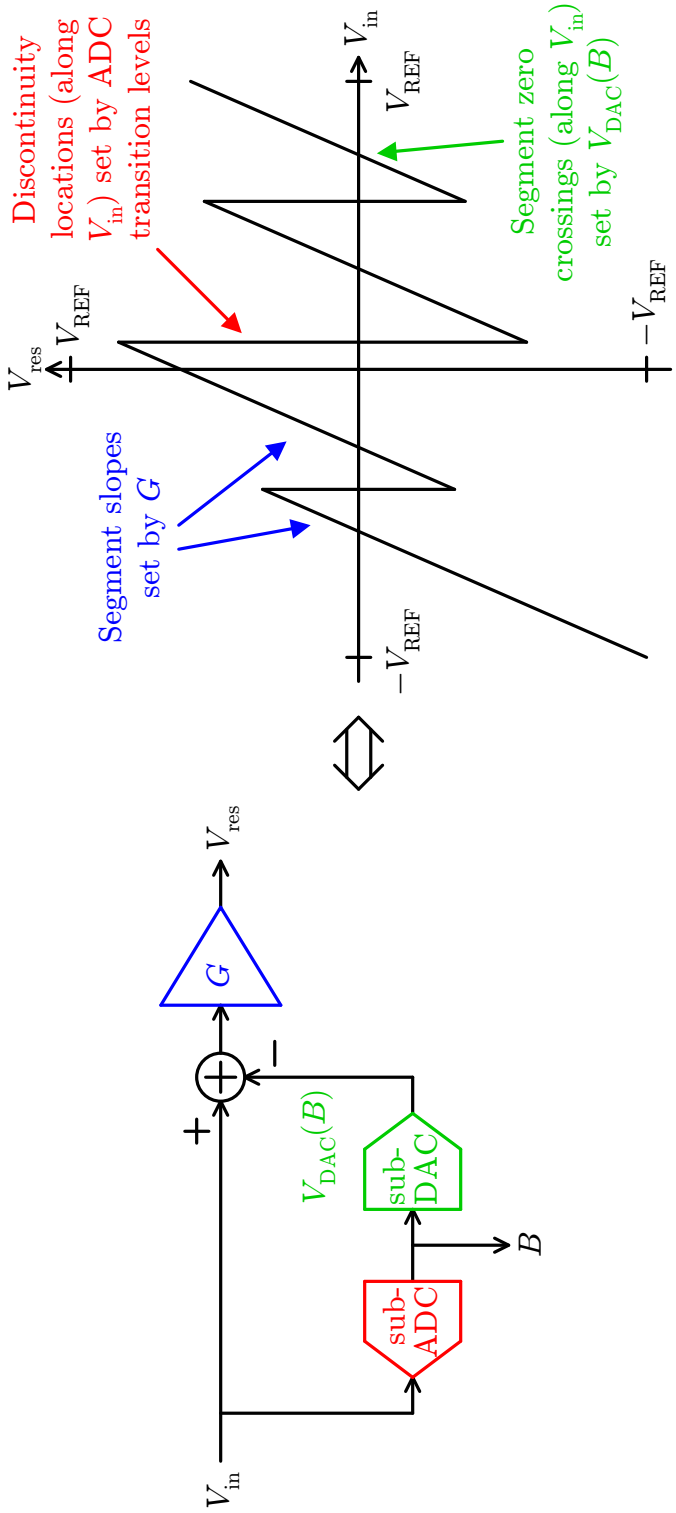

Figure 4.5: Equivalences between an analog stage and its transfer function.

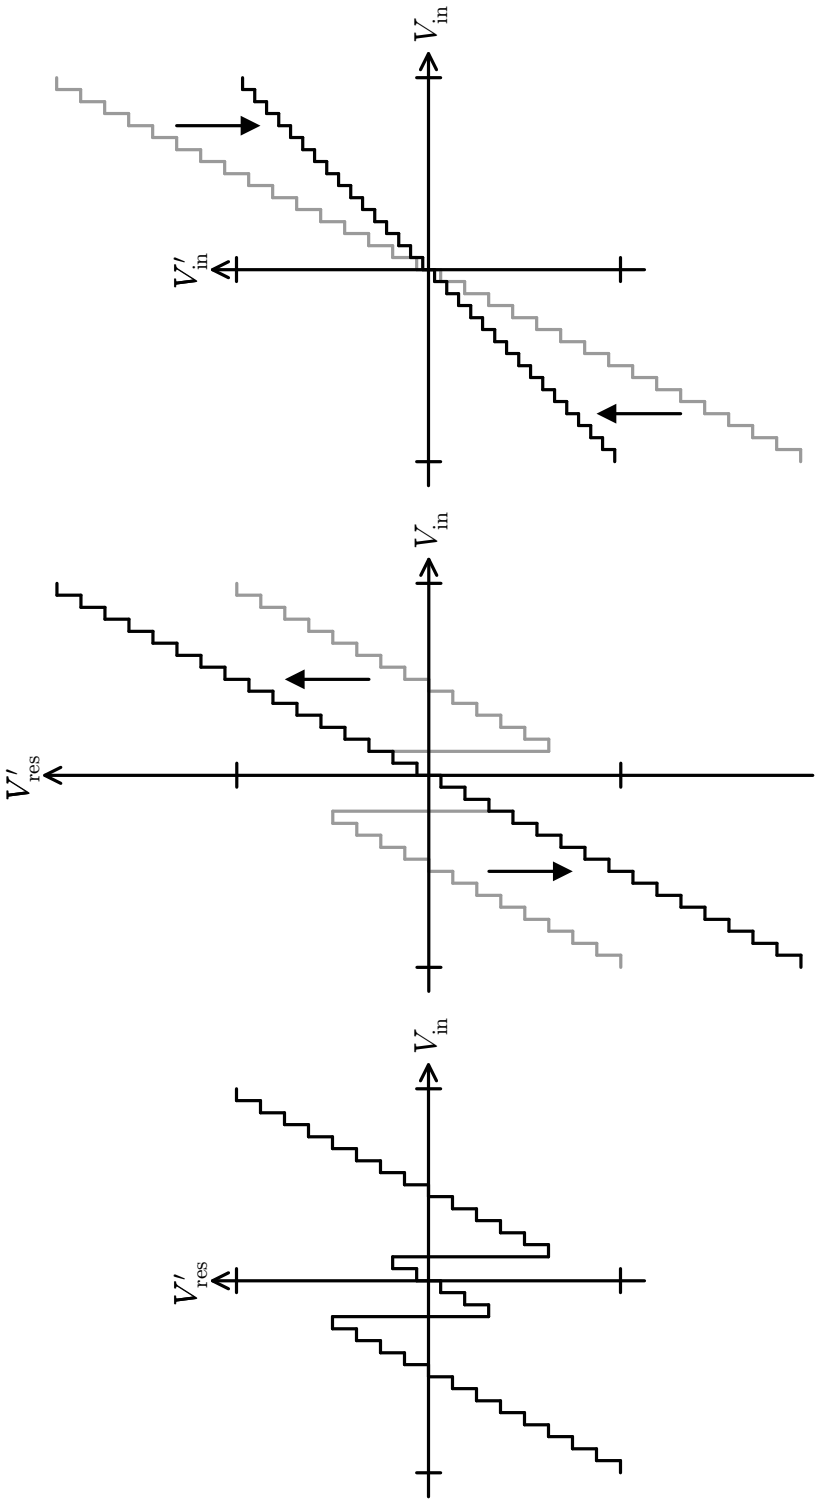

(a) V'_{res}-V_{in} characteristic. (b) Segment realignment. (c) V'_{in}-V_{in} characteristic.

Figure 4.6: Geometric interpretation of digital reconstruction as segment realignment. (a) V'_{res}-V_{in} characteristic achieved by stage. (b) Offsetting segments by $G'V'_{\text{DAC}}[B]$. (c) Applying $1/G'$ to achieve final V'_{in}-V_{in} characteristic.

Equation (4.3) as:

$$V'_{\text{in}} = \frac{1}{G'}\left(G'V'_{\text{DAC}}[B] + V'_{\text{res}}\right) \tag{4.9}$$

yields an interpretation of digital reconstruction as segment realignment. This interpretation is shown in Figure 4.6. In Equation (4.9), V'_{res} can be seen as providing digitized representations of each transfer function segment (depicted in Figure 4.6(a)). These segments are then displaced by $G'V'_{\text{DAC}}[B]$ to realign the segments end to end (depicted in Figure 4.6(b)) and the result scaled by $1/G'$ to produce V'_{in} (depicted in Figure 4.6(c)). Note that the amount of displacement between segments is equivalent to the height of the discontinuities in the stage transfer function. These heights prove important in the converter linearization techniques of Section 4.3.2.

4.1.3 Converter Operation, Algebraically

In a pipeline converter, the backend ADC of Figure 4.2 is implemented by an analog stage followed by another backend ADC. This recursion expands into a series of analog stages terminated by a final, typically coarse, ADC, as depicted in Figure 4.7 for a P-stage pipeline with the stages indexed from 1 (most upstream) to P (most downstream). Note that the P-th stage is actually a terminating ADC: when implemented, it contains neither the sub-DAC nor subtraction blocks shown in Figure 4.7. Rather, these blocks are included to make explicit the terminating ADC quantization error V_{err}. Assuming that the p-th stage, for $1 < p < P-1$, acts as:

$$V_{\text{res},p} = G_p\left(V_{\text{in},p} - V_{\text{DAC},p}[B_p]\right) \tag{4.10}$$

and noting that the terminating ADC quantization error is:

$$V_{\text{err}} = V_{\text{in},P} - V_{\text{DAC},P}[B_P] \tag{4.11}$$

it can be shown that:

$$y = V_{\text{DAC},1}[B_1] + \frac{V_{\text{DAC},2}[B_2]}{G_1} + \cdots + \frac{V_{\text{DAC},P}[B_P]}{G_1G_2\cdots G_{P-1}} + \frac{V_{\text{err}}}{G_1G_2\cdots G_{P-1}} \tag{4.12}$$

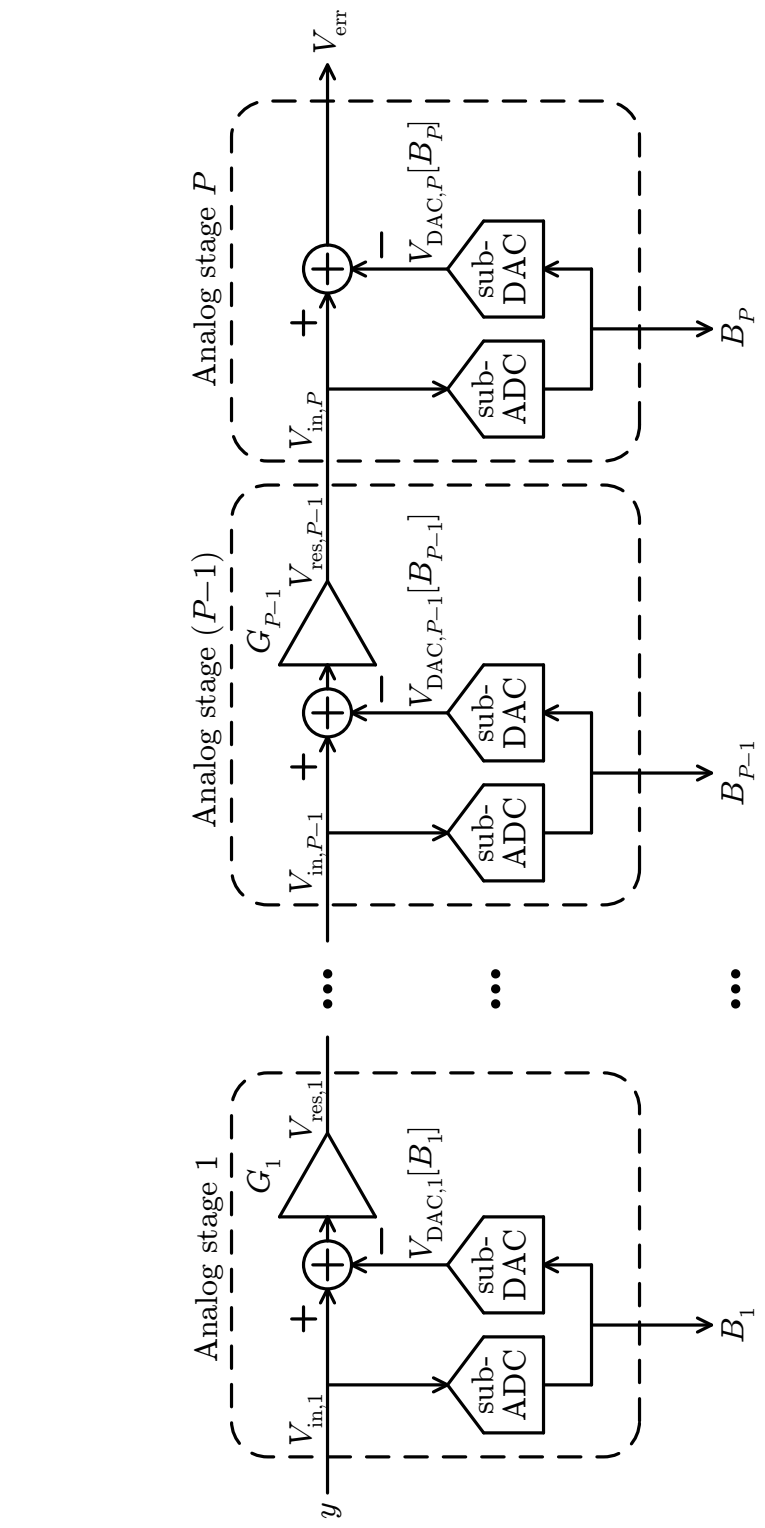

Figure 4.7: Analog stage portion of a pipeline converter: a series of $P-1$ analog stages terminated by a final ADC (labeled analog stage P).

for y the pipeline analog input. The first P terms of Equation (4.12) represent a decomposition of the input analog value into a sum of weighted sub-DAC voltages. The digital reconstruction effectively reconstitutes these terms as per:

$$y' = V'_{\mathrm{DAC},1}[B_1] + \frac{V'_{\mathrm{DAC},2}[B_2]}{G'_1} + \cdots + \frac{V'_{\mathrm{DAC},P}[B_P]}{G'_1 G'_2 \cdots G'_{P-1}} \tag{4.13}$$

This digital reconstruction can be implemented by extending the digital reconstruction of Figure 4.2 as illustrated in Figure 4.8.

On the other hand, the digital reconstruction of Figure 4.8 requires multiple multiplications, which can be expensive to implement. It can be simplified by noting that:

$$\begin{aligned}(G'_1 G'_2 \cdots G'_{P-1})y' = {} & (G'_2 \cdots G'_{P-1})G'_1 V'_{\mathrm{DAC},1}[B_1] \\ & + (G'_3 \cdots G'_{P-1})G'_2 V'_{\mathrm{DAC},2}[B_2] \\ & + \cdots \\ & + (1)G'_{P-1} V'_{\mathrm{DAC},P-1}[B_{P-1}] \\ & + V'_{\mathrm{DAC},P}[B_P]\end{aligned} \tag{4.14}$$

This digital reconstruction is shown in Figure 4.9. As the augmented table entries can be precomputed, it reduces digital reconstruction to only additions and table lookups and hence is often used in practice. Also, although the output of Figure 4.9 is a gained version of y', this gain is typically absorbed into the converter output formatter (which may convert y' to unsigned integer or two's complement formats, for example) and subsequently ignored.[4]

Assuming perfect matching between the analog stages and the digital reconstruction, the pipeline quantization error is:

$$Q_{\mathrm{pipeline}} = V_{\mathrm{in}} - V'_{\mathrm{in}} = \frac{V_{\mathrm{err}}}{G_1 G_2 \cdots G_{P-1}} \tag{4.15}$$

[4]In fact there is an even simpler form of digital reconstruction. Equation (4.14) is essentially a weighted sum of the B_p, where the weight is composed of the combination of the G'_p and $V'_{\mathrm{DAC},p}$. If the G'_p and $V'_{\mathrm{DAC},p}$ are properly chosen, this weighting can reduce to powers of 2 and the implementation reduce to a simple addition of bit-shifted B_p. For example, a pipeline of two 2-bit stages followed by a 2-bit terminating ADC—as depicted in Figure 4.11 and described by Equation (4.16)—can be reconstructed as per:

$$32y' = 16B_1 + 4B_2 + B_3$$

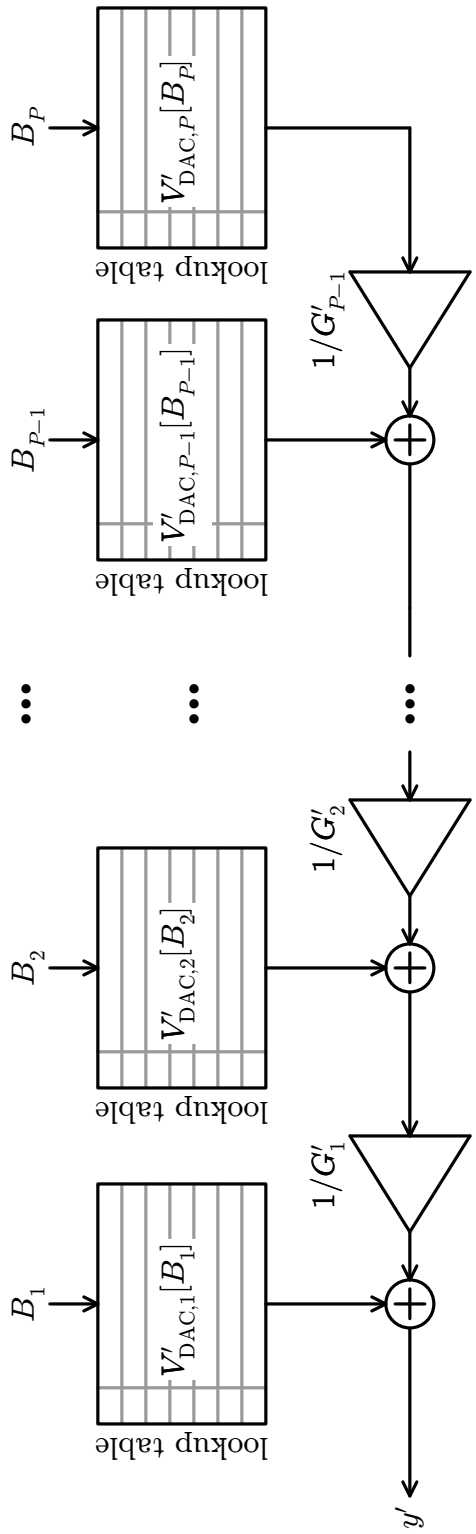

Figure 4.8: Digital reconstruction portion of a pipeline converter.

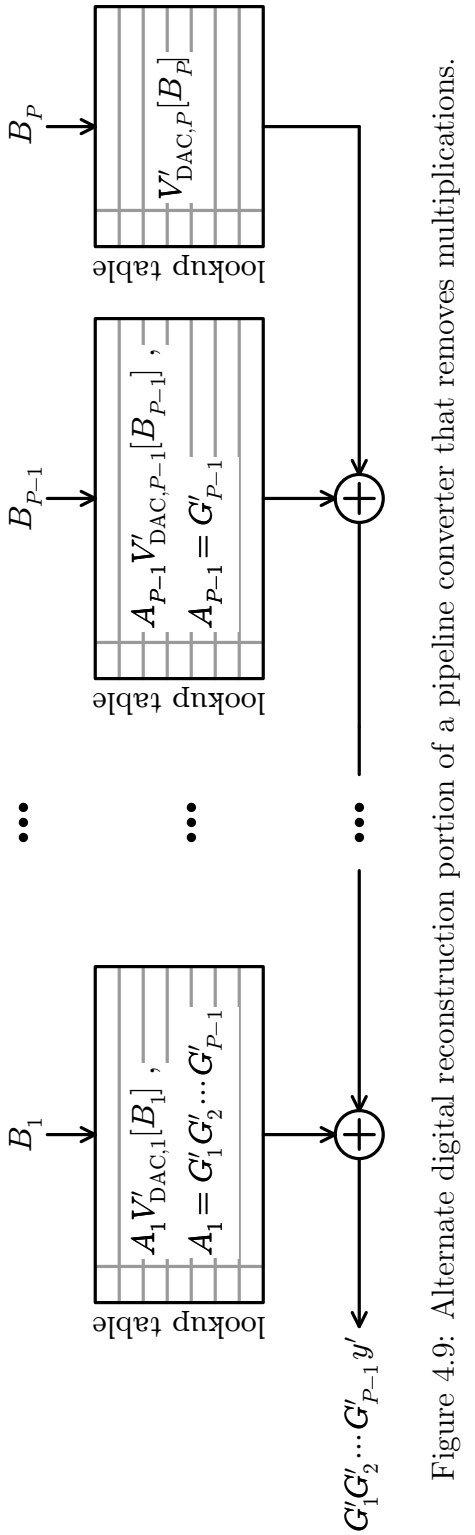

Figure 4.9: Alternate digital reconstruction portion of a pipeline converter that removes multiplications.

which is simply the last term of Equation (4.12). The similarity to the Q_{sys} of Equation (4.5) is clear. Equation (4.15), then, states that the quantization error of the terminating ADC, mitigated by the overall pipeline gain, determines the overall pipeline quantization error. Importantly, it highlights the counterintuitive result that it is ultimately the stage gains, and not the stage sub-ADC resolutions, that determine the overall pipeline resolution. In this way, the pipeline can achieve high resolution even though coarse sub-ADCs are used in each stage.

4.1.4 Converter Operation, Geometrically

A useful way to visualize the operation of a stage in a pipeline context is in terms of signal ranges, as illustrated in Figure 4.10. The leftmost column represents the stage's analog input range. Its sub-ADC divides this range and identifies the partition containing the input (here the $B = 2$ partition). The subtraction and sub-DAC blocks

which is the sum of B_3 plus B_2 bit-shifted two times plus B_1 bit-shifted four times, that is:

$$\begin{array}{cccccccc} & B_1\langle 1\rangle & B_1\langle 0\rangle & & & & & \\ & & & B_2\langle 1\rangle & B_2\langle 0\rangle & & & \\ + & & & & & B_3\langle 1\rangle & B_3\langle 0\rangle \\ \hline & y''\langle 5\rangle & y''\langle 4\rangle & y''\langle 3\rangle & y''\langle 2\rangle & y''\langle 1\rangle & y''\langle 0\rangle \end{array}$$

where $y'' = 32y'$. (By convention, for the multibit digital signal x, $x\langle 0\rangle$ represents the LSB.) Indeed, note that in this case even the summation itself is unnecessary: every y'' bit is directly equal to one of the analog stage digital output bits. Perhaps a more interesting example, then, is to consider a pipeline of two 2.8-bit stages (see Figure 4.14) followed by a 2-bit terminating ADC. It can be shown that this pipeline has the same Q_{pipeline} as the previous, but is instead reconstructed as per:

$$32y' = 8B_1 + 2B_2 + B_3$$

This pipeline can thus be reconstructed as:

$$\begin{array}{cccccccc} & B_1\langle 2\rangle & B_1\langle 1\rangle & B_1\langle 0\rangle & & & & \\ & & & B_2\langle 2\rangle & B_2\langle 1\rangle & B_2\langle 0\rangle & & \\ + & & & & & B_3\langle 1\rangle & B_3\langle 0\rangle \\ \hline & y''\langle 5\rangle & y''\langle 4\rangle & y''\langle 3\rangle & y''\langle 2\rangle & y''\langle 1\rangle & y''\langle 0\rangle \end{array}$$

where B_1 and B_2 are represented by 3 bits since they range over $[0, 6]$. Such shift-and-add reconstructions obviate even table lookups, replacing them with (often free) bit-shifts. They are thus typically used in practice. However, as the self-calibration techniques of Section 4.3.2 use amended lookup table entries to linearize the pipeline, the lookup table implementation is nonetheless assumed in these discussions.

then center the occupied partition (as shown in the center column) and the gain block amplifies the result to produce an output signal range (as shown in the rightmost column). The stage thus essentially isolates the input-containing partition and then zooms in on it.

A cascade of three such stages is shown in Figure 4.11, where the center column of Figure 4.10 is omitted for clarity. Again, the first stage zooms in on the input-containing partition. The next stage continues this process, partitioning the zoomed view, determining the input-containing partition, and zooming in still further. The analog input is thus determined to successively lesser degrees of uncertainty until the desired resolution is achieved.

The zooming-in process of Figure 4.11 is reminiscent of a quaternary search. This intuition can be formalized by applying the sub-DAC expression of Equation (4.8) to the analog decomposition of Equation (4.12) to derive the operation of the pipeline as:

$$y = \frac{V_{\text{REF}}}{32}\left(16B_1 + 4B_2 + B_3\right) + V_{\text{REF}}\left(-1 + \frac{1}{64}\right) + \frac{V_{\text{err}}}{32} \tag{4.16}$$

The first three terms (involving the B_p) are the decomposition of y achieved by this pipeline. These terms denote a third-fold decomposition by powers of 4: the pipeline does indeed essentially implement a 3-step quaternary search.[5] In fact, it can be shown that in general a cascade of P M-bit stages of the classical design described in Section 4.1.2 operates as an M-ary search of depth P.[6]

4.2 Errors

When actually implemented, the analog stage transfer function deviates from the ideals hitherto presented. If nothing else, slight variations in fabricated component values due to manufacturing tolerances readily affect the accuracy of the analog stage

[5]The following two terms represent an offset arising from the unsigned integer convention used to express the B_p, and (assuming perfect matching between the pipeline stages and the digital reconstruction) the last term represents the quantization error of the pipeline.

[6]The pipeline operation of Equation (4.16) is also indicative of an interpretation of a pipeline converter as a long division operation. Readers interested in this interpretation (and still cognizant of grade school arithmetic) are recommended to *Cline* [1995, Chap. 3].

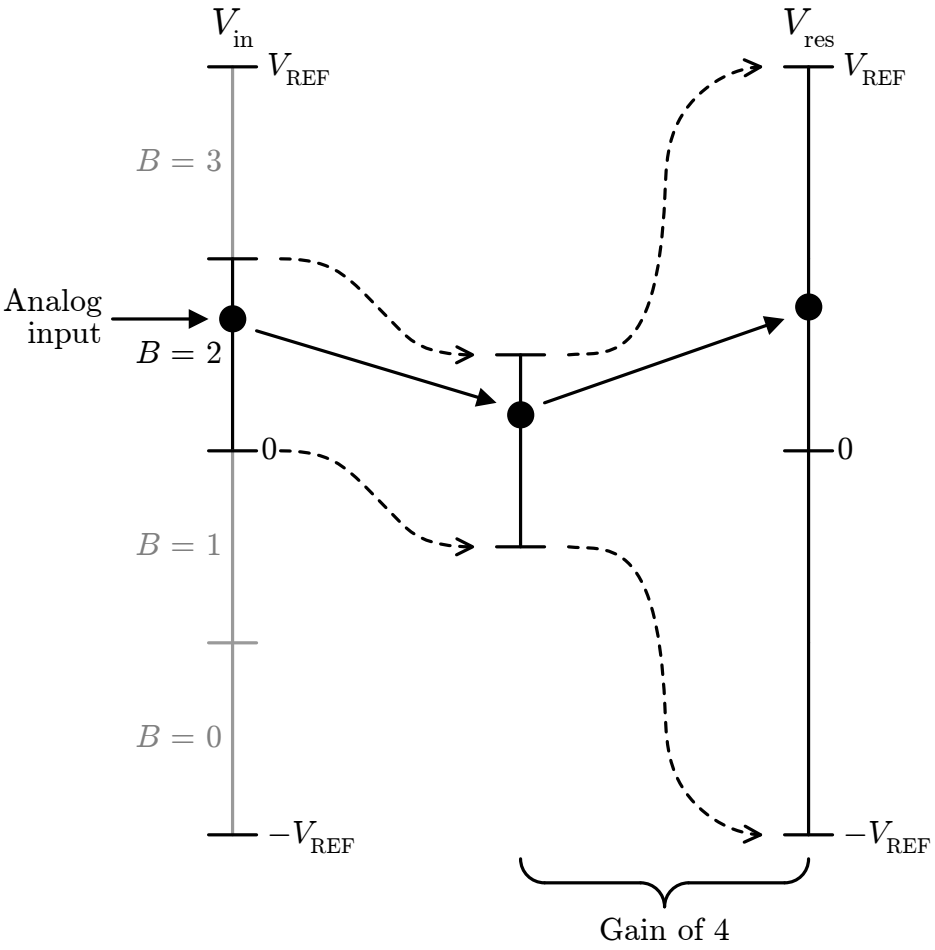

Figure 4.10: Operation of an analog stage in terms of signal ranges for a 2-bit stage as per Figure 4.4. Location of analog input indicated by black dot. In leftmost column (representing the stage's analog input signal range), partition occupied by analog input depicted in black, other partitions in gray.

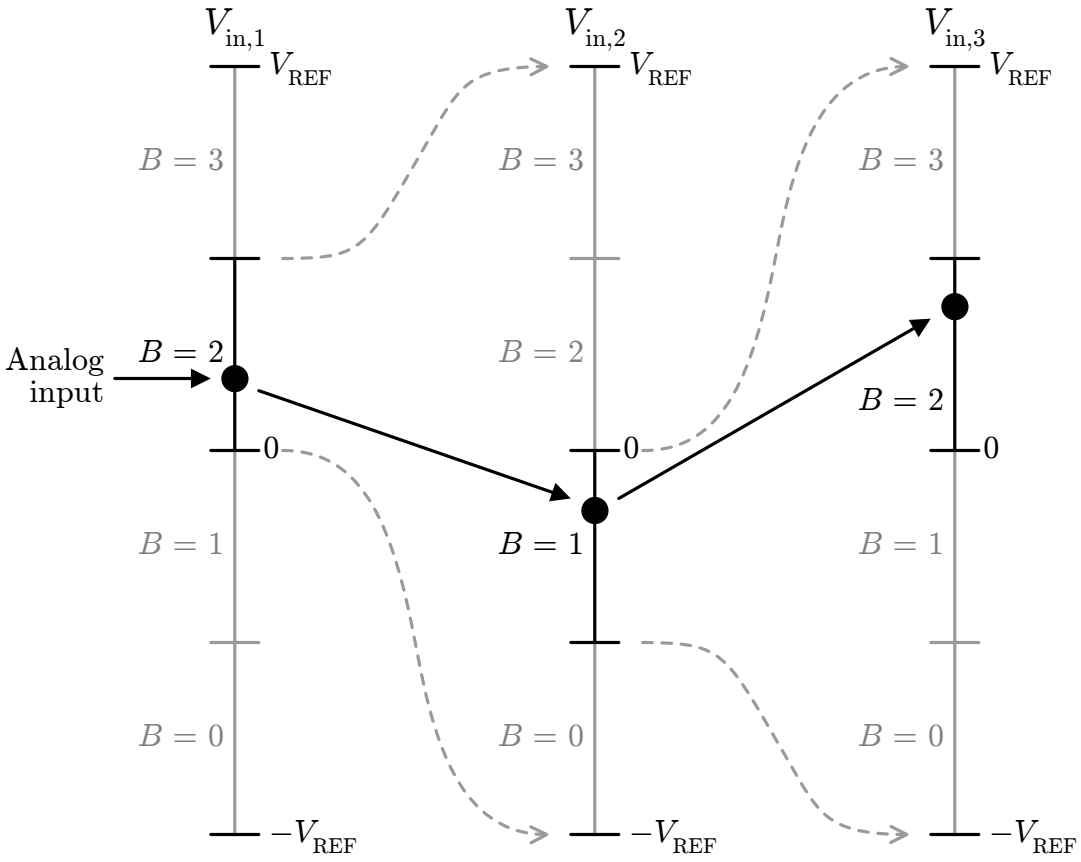

Figure 4.11: Operation of a pipeline in terms of signal ranges for a pipeline composed of three 2-bit stages as per Figure 4.4. Analog input ranges of all three stages shown. Location of analog input indicated by black dot, and partition occupied by input depicted in black (other partitions in gray).

sub-ADC, sub-DAC, and gain blocks.

The following discussions consider the impact of the nonidealities of overrange, offset, analog-digital mismatch, nonlinearity, and noise on converter performance. These nonidealities most commonly limit converter performance in practice.[7] In particular, these discussions emphasize the impact of these nonidealities on the converter SFDR. Note that thus an overall gain or offset in the pipeline input-output characteristic—so long as it is constant and input-independent—is allowed as these quantities constitute but linear deviations. Instead, in high resolution pipeline converters it is typically the uniformity of the quantization—or more precisely the lack thereof—that most limits converter linearity. To summarize these points mathematically, a reconstruction of the form:

$$y' = a(y + V_{\text{err}}) + b \tag{4.17}$$

is permitted so long as a and b are constant and V_{err} is uniform (i.e., akin to Figure 3.7). In this dissertation, such a reconstruction is called a linear reconstruction.

4.2.1 Overrange

To fully accrue the Q_{pipeline} of Equation (4.15) V_{err} must be bounded. Assuming the terminating ADC operates as per Figure 3.7, V_{err} is bounded so long as $V_{\text{in},P}$ lies within the V_{FS} of the terminating stage. As $V_{\text{in},P} = V_{\text{res},P-1}$, this input range requirement translates to an output range requirement on the preceding stage. Propagating this logic upstream, in general the range of the residue output of the p-th stage should lie within the input range of the backend ADC composed of the $(p+1)$-th through P-th stages.

For design purposes, the more stringent requirement that the output range of the p-th stage lie within the input range of the $(p+1)$-th stage is typically adopted.[8]

[7] Naturally, this list is by no means exhaustive: many additional second-order effects can also affect performance. The consideration of many of these secondary concerns can be found in the circuit implementation descriptions of Chapter 5.

[8] Technically, this requirement is sufficient but not necessary for bounding V_{err}. Consider: should the V_{res} output of a stage overrange, so long as the downstream stages eventually bring the sample

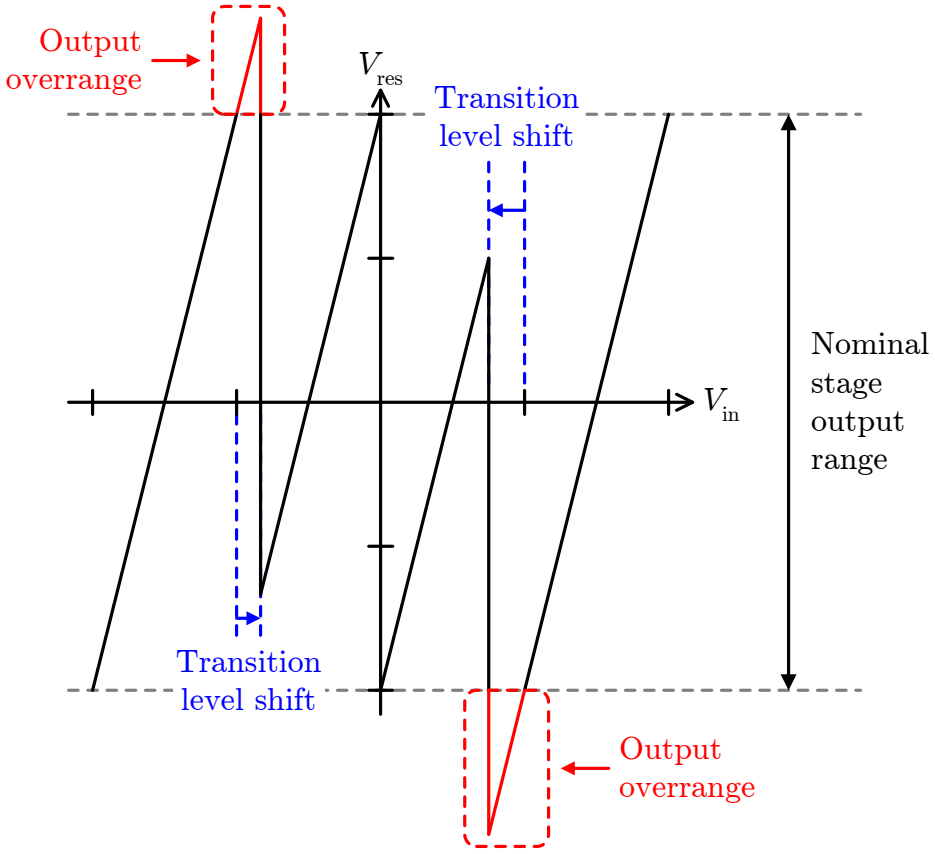

Figure 4.12: Demonstration of overrange due to sub-ADC transition level shifts on a 2-bit stage.

Typically, the two ranges are chosen to coincide (to make full use of each) resulting in a single interstage range. Furthermore, this interstage range is usually kept constant throughout all the stages.[9] If the V_{res} output of a stage surpasses this interstage range, then the backend ADC (and eventually the terminating ADC) likely saturates. In saturation, though the converter input changes, its output does not: in terms of the quantization input-output characteristic, a quantizer bin extends beyond its normal width, compromising the conversion linearity.

back into the input range of the terminating ADC, V_{err} is still properly bounded. Indeed, *Opris et al.* [1998] exploited by interjecting special "overrange" stages throughout a pipeline converter: the transfer functions of these overrange stages returned overranged outputs from upstream stages back to the input range of the terminating ADC.

[9] Of course, counterexamples exist: see footnote 12 of this chapter.

Unfortunately, the stages shown in Figures 4.3 and 4.4 (and indeed all such M-bit stages) are particularly susceptible to overrange from shifts in their sub-ADC transition levels. Such shifts readily occur from offsets in the sub-ADC comparators (see Section 5.5.2). As illustrated in Figure 4.12, such shifts extend and contract the segments, possibly causing a segment to exceed the nominal stage output range. If this elongation surpasses the $\Delta/2$ of the backend ADC, then the terminating ADC saturates. Especially for upstream stages in high resolution pipelines, preventing such overrange amounts to very little margin.

A common solution is to introduce additional segments by adding sub-ADC transition levels. Standard modifications for the stages of Figures 4.3 and 4.4 are shown in Figures 4.13 and 4.14, respectively. In contrast to the 1-bit and 2-bit stages, these 1.5-bit and 2.8-bit stages are robust to transition level shifts of as much as $V_{\text{REF}}/4$ and $V_{\text{REF}}/8$, respectively, before exceeding the output range. The additional transition levels are chosen so that the corresponding sub-DAC voltages are of the form $a/2^b$, for $a, b \in \mathbb{Z}$ to simplify the digital reconstruction implementation.[10] For historical reasons, this technique is often called digital error correction or redundancy [*Lewis et al.*, 1992]: in this dissertation, the (more accurate) latter term is used.[11] Redundancy is used in the pipeline stage design of the SVADC-1 and (unless otherwise stated) is assumed throughout the remainder of these discussions.[12,13]

[10] In particular, sub-DAC voltages of this form, coupled with stage gains of the form 2^M, $M \in \mathbb{N}$, enable the use of the shift-and-add digital reconstruction implementation described in footnote 4 of this chapter.

[11] In the case where the stage input and output range must coincide, the classical M-bit stage, $M \in \mathbb{N}$, contains the minimal number of transition levels to support a stage gain of 2^M without overrange. The 1.5-bit and 2.8-bit stages, then, add extra—that is, "redundant"—transition levels.

[12] Another common solution to overrange is to reduce the stage gain, either slightly (for example from 2 to 1.93 for a 1-bit stage [*Karanicolas et al.*, 1993]) or wholesale (for example from 8 to 4 for a 3-bit stage [*Mehr and Singer*, 2000]). Such reductions guarantee that the stage output range is smaller than the next stage input range. On the other hand, such reductions increase Q_{pipeline}, lowering the overall pipeline resolution.

[13] A subtlety of the 1.5-bit and 2.8-bit stages is that they no longer contain a $V_{\text{in}} = 0$ V transition level. These stages thus defer the determination of the sign of the input to a latter stage. In applications where sign determination is essential, though, there can be concerns of offsets accumulating in the stages prefacing the sign decision.

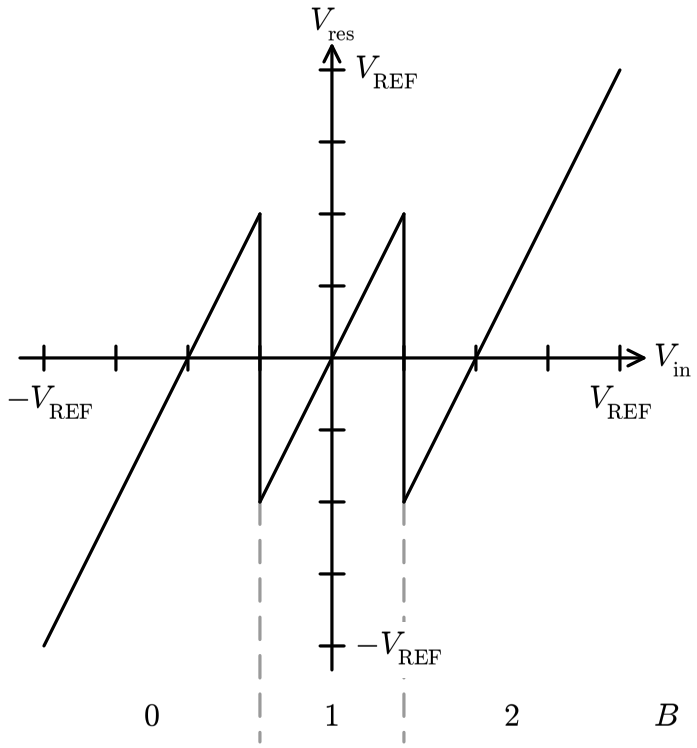

Figure 4.13: Transfer function of a 1.5-bit pipeline stage; compare Figure 4.3. sub-ADC transition levels: $-\frac{1}{4}V_{\text{REF}}$ and $\frac{1}{4}V_{\text{REF}}$. sub-DAC output voltages: $-\frac{1}{2}V_{\text{REF}}$, 0, and $\frac{1}{2}V_{\text{REF}}$. Gain: 2.

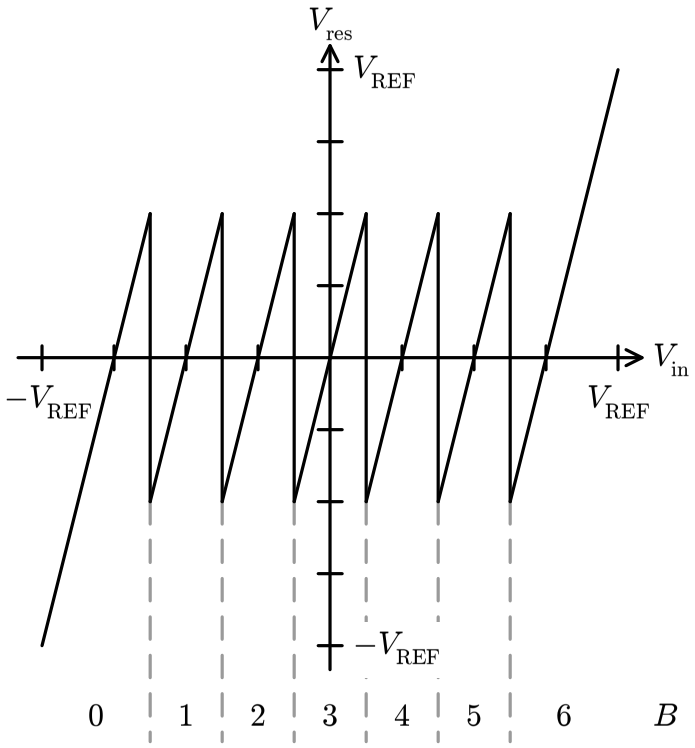

Figure 4.14: Transfer function of a 2.8-bit pipeline stage; compare Figure 4.4. sub-ADC transition levels: $-\frac{5}{8}V_{\text{REF}}, -\frac{3}{8}V_{\text{REF}}, \ldots, \frac{5}{8}V_{\text{REF}}$. sub-DAC output voltages: $-\frac{3}{4}V_{\text{REF}}, -\frac{2}{4}V_{\text{REF}}, \ldots, \frac{3}{4}V_{\text{REF}}$. Gain: 4.

4.2.2 Offset

Offsets in the analog stage transfer function stem from various sources, such as an overall shift in the sub-DAC voltages or an offset incurred during subtraction. As illustrated in Figure 4.15(a), offset can be modeled by an additive term V_{OFF}:

$$V_{\text{res}} = G\left(V_{\text{in}} - V_{\text{DAC}}[B]\right) + V_{\text{OFF}} \tag{4.18}$$

After input-referral, V_{OFF} distributes as in Figure 4.15(b). The resulting $-V_{\text{OFF}}/G$ preceding the sub-ADC is equivalent to a shifting in the ADC transition levels but so long as the redundancy is sufficient to prevent overrange under the sum of $-V_{\text{OFF}}/G$ and any comparator offsets, this offset does not harm converter linearity. The resulting $+V_{\text{OFF}}/G$ added to V_{in} is equivalent to an output-referred offset in the preceding stage, and eventually translates into but a global offset at the pipeline input. Hence offsets can cause shifts in the converter input range, but (assuming no overrange occurs) do not upset linearity.

Offset tolerance also extends to the digital reconstruction as shown in Figure 4.16(a). Here, offsets can occur in either the lookup table ($V_{\text{OFF},a}$) or the backend ADC ($V_{\text{OFF},b}$). In both cases, as shown in Figure 4.16(b) the offsets readily pass to V'_{in} and eventually translate to a global offset that again does not disturb converter linearity. Tolerance to $V_{\text{OFF},a}$ is especially noteworthy, as it implies that only the differences between adjacent $V'_{\text{DAC}}[B]$ are important in digital reconstruction. Indeed, the lookup table can thus be constructed assuming some fortuitous offset—for example, an offset such that all subsequent processing requires only positive numbers—without harming converter linearity.

4.2.3 Analog-Digital Mismatch

The ideal Q_{pipeline} of Equation (4.15) only holds if there is exact matching between the analog stages and digital reconstruction, that is, only if $G_p = G'_p$ and $V_{\text{DAC},p}[B_p] = V'_{\text{DAC},p}[B_p]$. In this case, the pipeline quantization error is solely determined by the last term of Equation (4.12). Unfortunately, deviations in the sub-DAC and gain

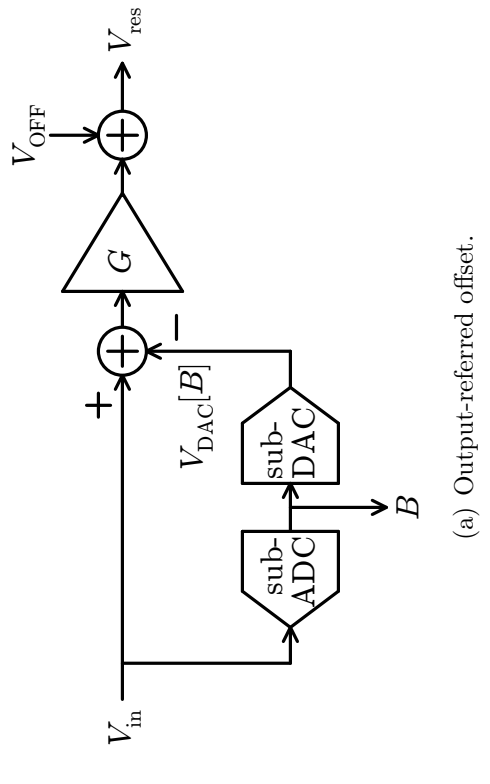

(a) Output-referred offset.

(b) Input-referred offset.

Figure 4.15: Impact of offset on an analog stage. Offset shown (a) output-referred and (b) input-referred.

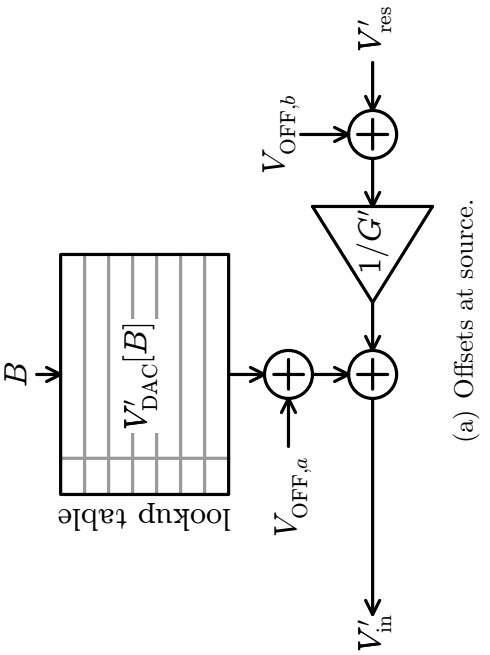

(a) Offsets at source.

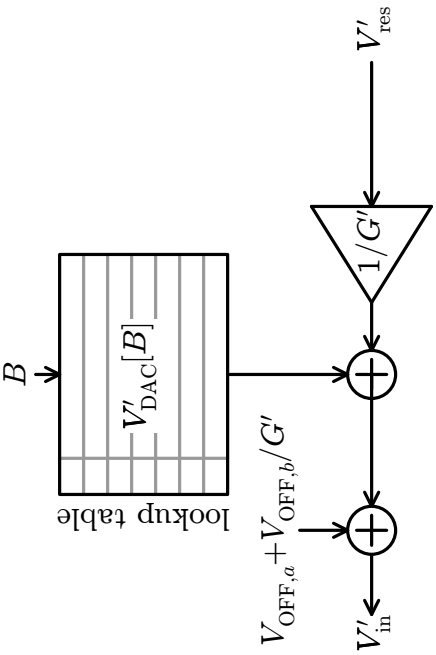

(b) Output-referred offset.

Figure 4.16: Impact of offset on digital reconstruction. Offsets shown (a) at source and (b) output-referred.

blocks of the analog stage often arise in implementation. In this case, analog-digital matching is no longer guaranteed, and all the terms of Equation (4.12) can contribute error. This error causes stage-dependent distortions in the converter staircase that, being systematic, can greatly diminish converter linearity.

An example is shown in Figure 4.17, which simulates a 12-bit pipeline composed of four 2.8-bit stages followed by a 4-bit terminating ADC.[14] Here, the first stage incorporates gain and sub-DAC mismatch but is otherwise ideal.[15] The resulting deviation of the first stage transfer function from ideal is shown in Figure 4.17(a). The offset of the segments is due to sub-DAC deviations, the slope due to gain deviations. The rest of the pipeline stages are ideal. The complete pipeline quantization error is shown in Figure 4.17(b). Note that the quantization error is no longer uniform, and that the nonuniformities can be traced back to the first stage deviations. This systematic distortion greatly limits the SFDR, as demonstrated in the output spectrum of Figure 4.17(c): even with 60-dB $\mathrm{SNR_{FS}}$ and an input sinusoid amplitude of $0.95V_{\mathrm{REF}}$, the SFDR of this 12-bit converter is limited to just 79 dB, a far cry from the over 92 dB expected from Figure 3.20.

In practice, analog-digital mismatch is the predominant cause of linearity loss in pipeline converters. Hence a further consideration of the primary source of this mismatch—the stage residue amplifier implementation—is warranted.

4.2.3.1 Mismatch from Residue Amplifier Implementation

Most high resolution pipeline converters are implemented with switched-capacitor circuitry—often with fully differential signaling—in which the analog stage sub-DAC, subtraction, and gain blocks are achieved in a single residue amplifier circuit. An example residue amplifier circuit is shown in Figure 4.18. It requires two

[14] The ideal quantization error of such a pipeline given by Equation (4.15) is:

$$Q_{\mathrm{pipeline}} = \frac{V_{\mathrm{err}}}{G_1G_2G_3G_4} = \frac{V_{\mathrm{FS}}/2^4}{4^4} = \frac{V_{\mathrm{FS}}}{2^{12}}$$

resulting in a 12-bit conversion.

[15] In particular, the first stage is a 2.8-bit stage implemented via the circuit of Figure 5.21 assuming a capacitor value deviation of $\sigma = 0.1\%$ from ideal.

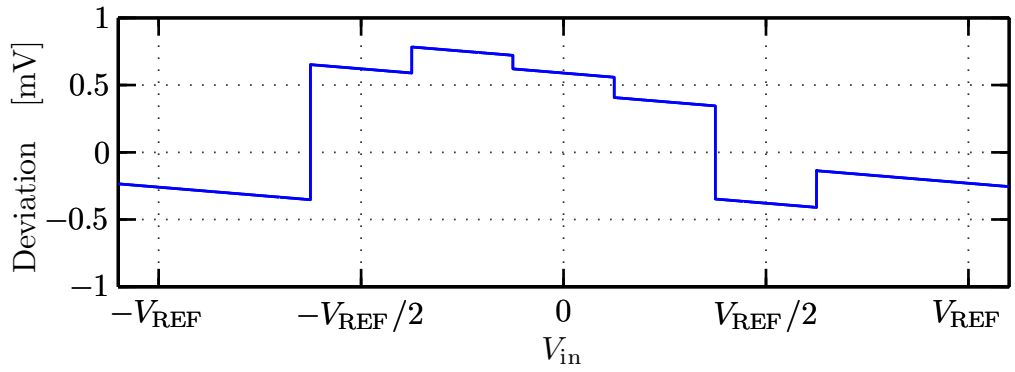

(a) Analog-digital mismatch of first stage.

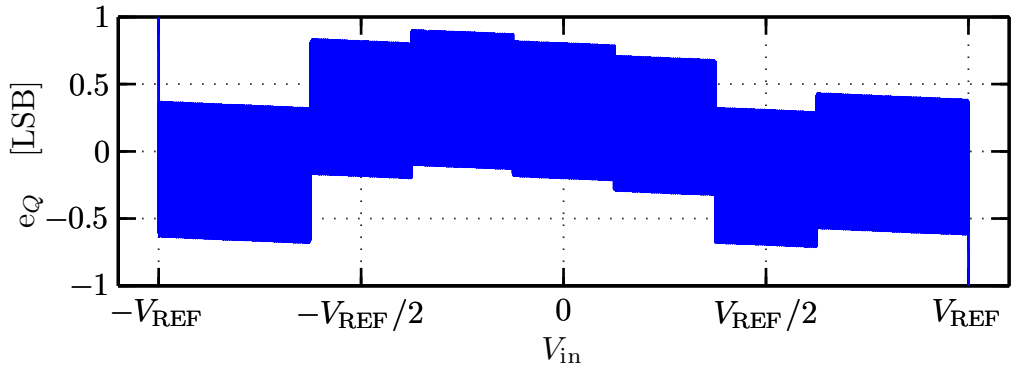

(b) Pipeline quantization error.

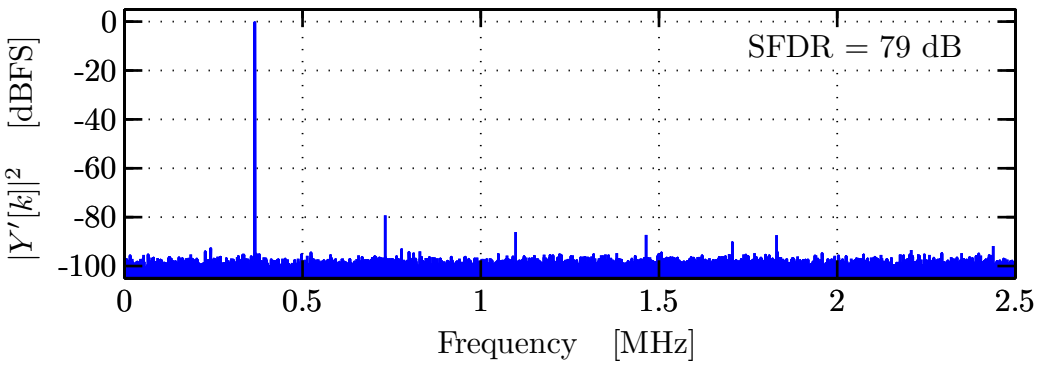

(c) Pipeline output power spectrum.

Figure 4.17: Effect of analog-digital mismatch on a 5-stage pipeline (four 2.8-bit stages, 4-bit terminating ADC). (a) Deviation of first stage's transfer function from ideal (all other stages ideal). Resulting (b) quantization error and (c) output power spectrum. Spectrum computed while processing a $0.95V_{\mathrm{REF}}$ amplitude, 366-kHz input sinusoid of 60-dB $\mathrm{SNR}_{\mathrm{FS}}$. Spectrum in dBFS (a V_{REF} amplitude sinusoid is 0 dBFS).

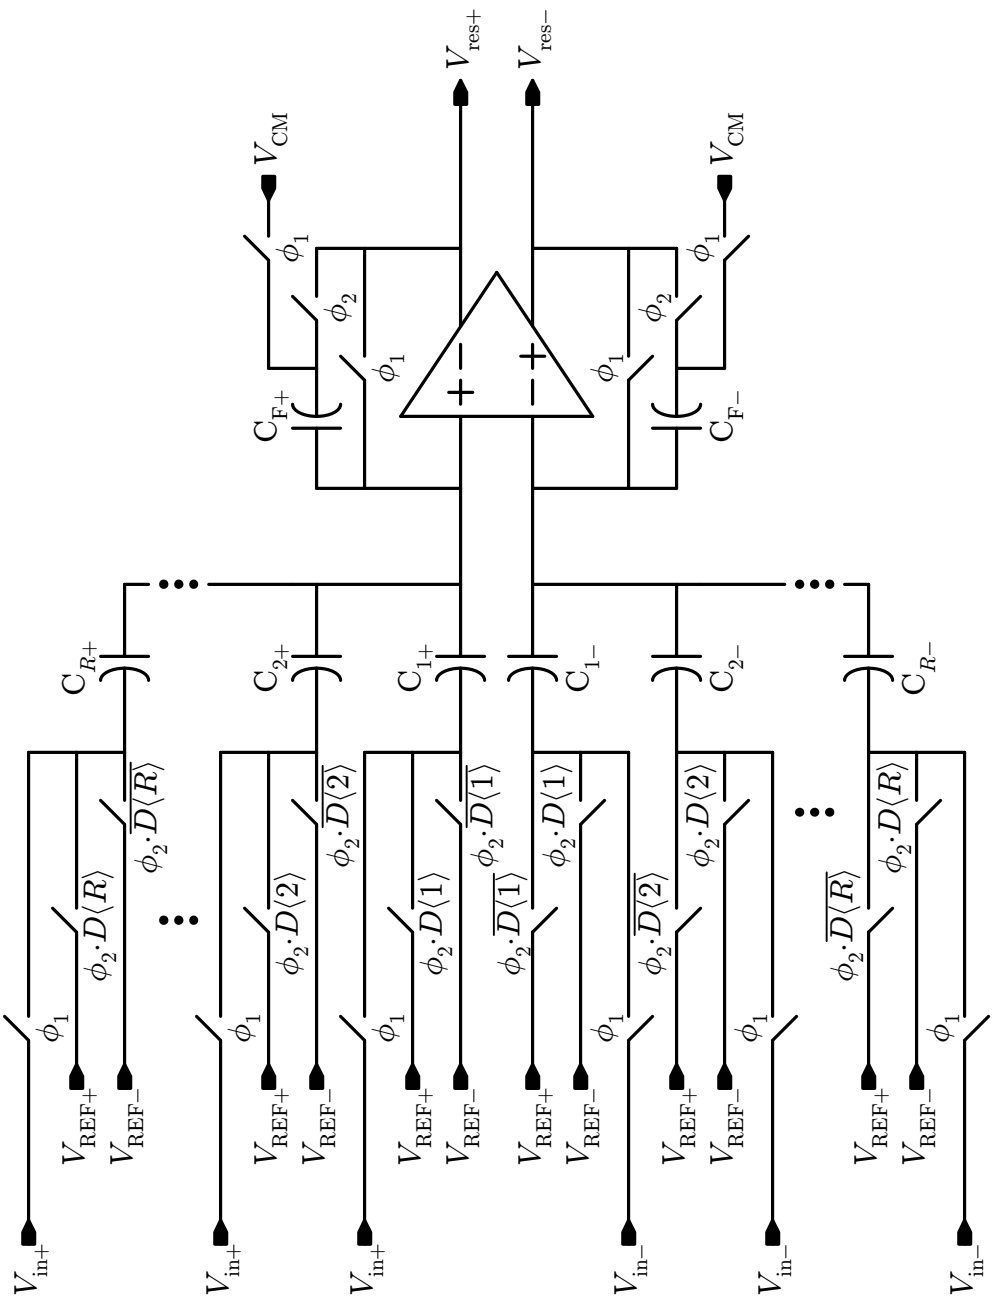

Figure 4.18: An example residue amplifier circuit. See Figure 4.19 for clock signals ϕ_1 and ϕ_2.

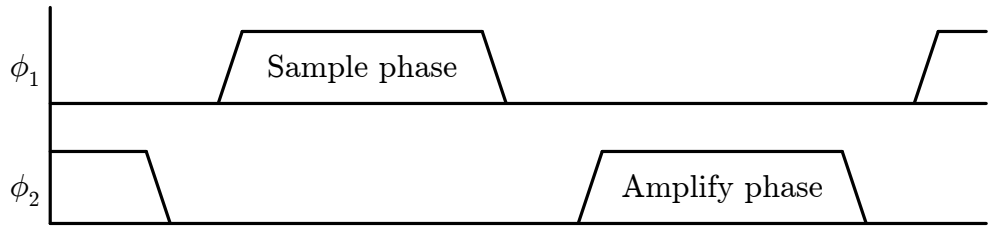

Figure 4.19: Non-overlapping clock signals ϕ_1 and ϕ_2 for the residue amplifier of Figure 4.18. By convention, when the clock signal is high, the driven switch is connected.

non-overlapping clock phases ϕ_1 and ϕ_2—called the sample and amplify phases, respectively—which are depicted in Figure 4.19. By convention, when the clock signal of a switch is high, the switch is connected.

To understand the operation of the residue amplifier, first consider just the positive branch (that is, just the upper circuitry of the differential circuit). In ϕ_1, the inputs and outputs of the operational amplifier are shorted and the input signal $V_{\text{in}+}$ is developed across the sample capacitors C_{1+} through C_{R+}. In ϕ_2, certain sample capacitors are switched to the positive reference voltage $V_{\text{REF}+}$ (the remainder are switched to the negative reference $V_{\text{REF}-}$) and the operational amplifier configured to amplify.[16] The configuration of the sample capacitors in ϕ_2, as dictated by the $D\langle r\rangle$ signals, implements a sub-DAC voltage. Specifically, defining the differential signals:

$$V_{\text{in}} = V_{\text{in}+} - V_{\text{in}-} \tag{4.19}$$

$$V_{\text{res}} = V_{\text{res}+} - V_{\text{res}-} \tag{4.20}$$

$$V_{\text{REF}} = V_{\text{REF}+} - V_{\text{REF}-} \tag{4.21}$$

and assuming matching between the positive and negative branch signal capacitors as per:

$$C_{\text{F}+} = C_{\text{F}-} = C_{\text{F}} \quad \text{and} \quad C_{r+} = C_{r-} = C_r \ , \ r = 1, \ldots, R \tag{4.22}$$

[16]For completeness: the negative branch mimics the positive branch during ϕ_1, developing $V_{\text{in}-}$ across C_{1-} through C_{R-}. The mimicry continues during ϕ_2, but with the change that the sample capacitors connections to $V_{\text{REF}+}$ and $V_{\text{REF}-}$ are inverted: where a positive branch sample capacitor connects to $V_{\text{REF}+}$, the corresponding negative branch sample capacitor connects to $V_{\text{REF}-}$.

it can be shown that V_{res} at the end of ϕ_2 (assuming complete settling) is:

$$V_{\text{res}} = \underbrace{\frac{C_{\text{S}}}{C_{\text{F}}}}_{G} \left\{ V_{\text{in}} - \underbrace{\left[\sum_{r=1}^{R} (-1)^{\overline{D\langle r \rangle}} \frac{C_r}{C_{\text{S}}} \right] V_{\text{REF}}}_{V_{\text{DAC}}[B]} \right\} \tag{4.23}$$

where C_{S} is the total signal capacitance:

$$C_{\text{S}} = \sum_{r=1}^{R} C_r \tag{4.24}$$

As the labeling of Equation (4.23) makes clear, the stage gain G is set by the ratio of the total sample capacitance C_{S} to the feedback capacitance C_{F}, and the stage sub-DAC voltages are determined by the $D\langle r \rangle$ configuration and the ratio of a particular sample capacitance C_r to the total sample capacitance C_{S}. Typically, $D\langle r \rangle$ is related to B through a thermometer coding:[17]

$$\begin{aligned}
B = 0 \quad &\Longleftrightarrow \quad D\langle r \rangle = \{0, 0, 0, \ldots, 0\} \\
B = 1 \quad &\Longleftrightarrow \quad D\langle r \rangle = \{1, 0, 0, \ldots, 0\} \\
B = 2 \quad &\Longleftrightarrow \quad D\langle r \rangle = \{1, 1, 0, \ldots, 0\} \\
&\vdots \\
B = M_{\text{level}} - 1 \quad &\Longleftrightarrow \quad D\langle r \rangle = \{1, 1, 1, \ldots, 1\}
\end{aligned} \tag{4.25}$$

Hence as B increases, more $(C_r/C_{\text{S}})V_{\text{REF}}$ are added to V_{in}, reflecting an increasing sub-DAC voltage.

In other words, the residue amplifier essentially implements the gain and sub-DAC through capacitor ratios. Clearly these implementations are susceptible to capacitor value errors from mismatch. In addition, the gain is also sensitive to operational amplifier nonidealities. In practice, a realizable operational amplifier achieves only a finite gain A, imposes a nonzero parasitic capacitance C_P at each of its input nodes,

[17] As a detail, the context here may suggest that $D\langle r \rangle$ is derived from B. In practice, though, the opposite is true, see Figures 5.19 and 5.23. This subtlety, while pertinent for circuit implementation, does not alter the discussion here.

and incurs an input-referred offset V_{OFF}. Accounting for these nonidealities, Equation (4.23) becomes:

$$V_{\mathrm{res}} = \left[\frac{A\,C_{\mathrm{S}}}{(1+A)C_{\mathrm{F}}+C_{\mathrm{P}}}\right]\left\{V_{\mathrm{in}} - \left[\sum_{r=1}^{R}(-1)^{\overline{D\langle r\rangle}}\frac{C_r}{C_{\mathrm{S}}}V_{\mathrm{REF}}\right]\right\} - \frac{A(C_{\mathrm{F}}+C_{\mathrm{P}})}{(1+A)[(1+A)C_{\mathrm{F}}+C_{\mathrm{P}}]}V_{\mathrm{OFF}} \tag{4.26}$$

Hence operational amplifier nonidealities exacerbate the gain error and add an offset. In fact, in reality the residue amplifier implementation of the gain and sub-DAC deviates even further from ideal than Equation (4.26) suggests, as this Equation still neglects positive and negative branch capacitor mismatch.

While there are other factors that also contribute to analog-digital mismatch, in most cases capacitor mismatch and operational amplifier nonidealities in the residue amplifier implementation dominate.

4.2.4 Operational Amplifier Nonlinearity

Another residue amplifier nonideality contributing to converter performance loss is operational amplifier nonlinearity. As sketched in Figure 4.20(a), typically the operational amplifier gain A is large for small outputs but monotonically diminishes as the output grows.[18] The consequent distortion of the stage transfer function is sketched in Figure 4.20(b): as seen, the distortion is most pronounced for the largest outputs.

Given knowledge of the operational amplifier implementation, it is possible to describe such nonlinear residue amplifier transfer functions analytically (e.g., *Cline* [1995, pp. 69–79]). While such approaches describe the complete shape of the residue amplifier transfer function, in converter design it is often sufficient to just characterize the maximum distortion. This dissertation thus introduces an analysis in the latter

[18]For clarity: the operational amplifier gain A is defined as $V_{\mathrm{y}} = -A(V_{\mathrm{x}})$ where V_{x} and V_{y} are the operational amplifier differential input and output, respectively. Assuming complete settling, A is the DC operational amplifier gain (and varies with the operational amplifier input and output conditions).

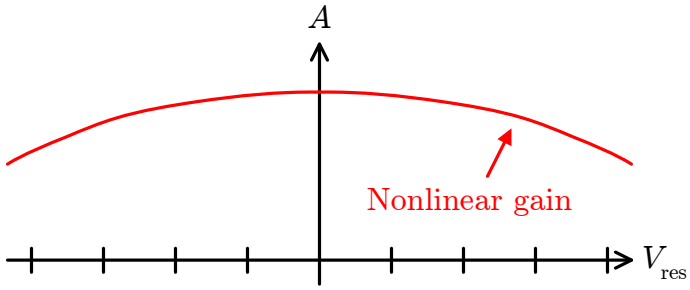

(a) Operational amplifier gain characteristic.

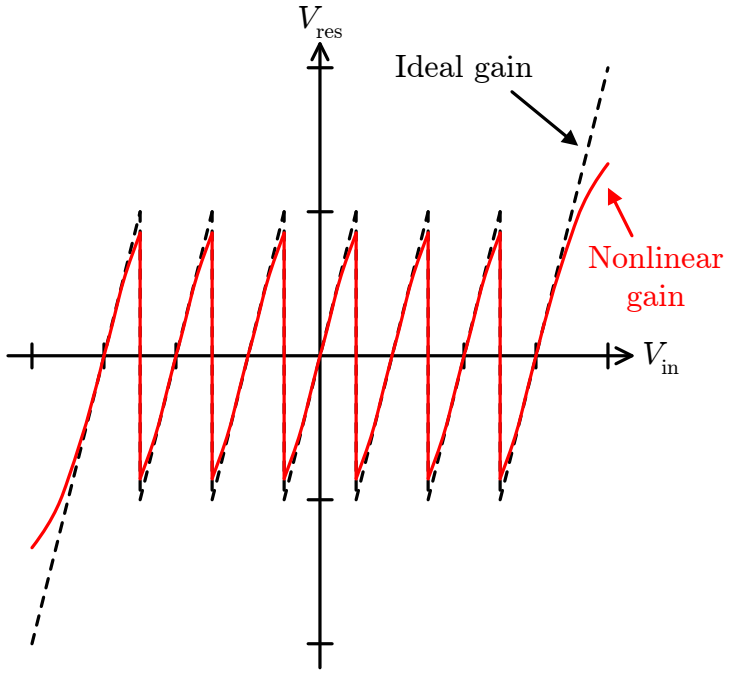

(b) Stage transfer function.

Figure 4.20: Effect of operational amplifier nonlinearity on stage transfer function. Both (a) operational amplifier gain A and (b) resulting transfer function shown. Characteristics assuming ideal operational amplifier shown in (dashed) black, assuming nonlinear operational amplifier shown in (solid) red.

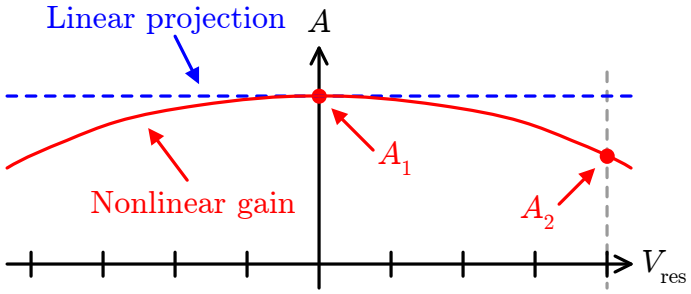

(a) Operational amplifier gain characteristic.

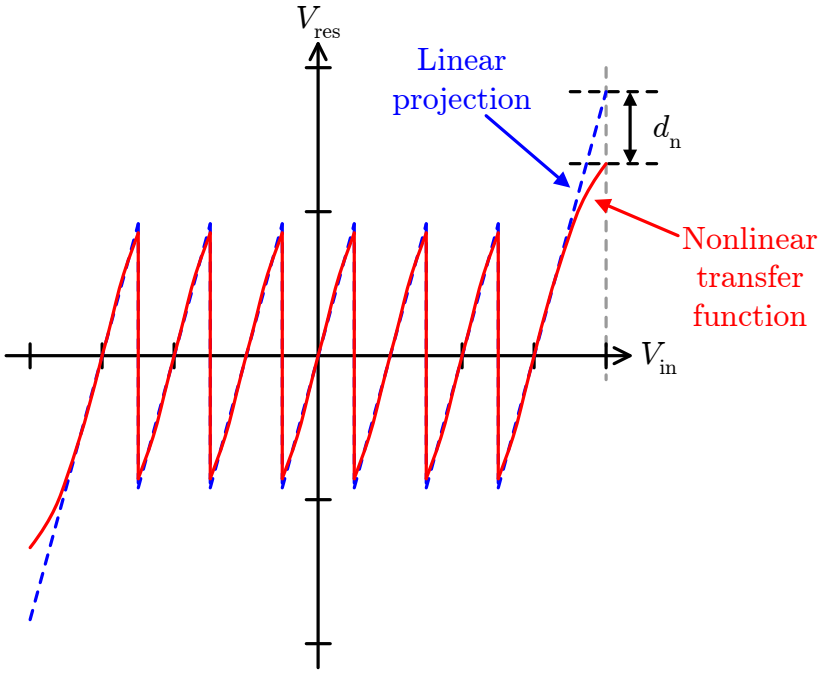

(b) Stage transfer function.

Figure 4.21: Characterization of greatest deviation d_n due to nonlinearity. Actual nonlinear transfer function shown in (solid) red, linear projection of transfer function (from zero crossing) shown in (dashed) blue. Both (a) operational amplifier gain and (b) stage transfer function shown.

vein.

Specifically, the analysis characterizes the maximum deviation of the distorted transfer function from a linear fit since—as is shown in Section 4.3.2.1—under self-calibration an overall stage linear deviation does not limit linearity. The linear fit chosen is the linear projection of the slope of the extremum segment at its zero crossing, as depicted in Figure 4.21. Naturally, other linear fits can be envisioned, but this particular fit makes the analysis more tractable and certainly captures a worst case. As shown in Figure 4.21(a), assume that A is greatest when $V_{\text{res}} = 0$ V, and that it monotonically decreases (in symmetric fashion about $V_{\text{res}} = 0$ V) for nonzero V_{res}. The maximum deviation d_{n} from the linear projection is then at the point of maximum V_{res}, that is, at the far edge of the extremum segments as shown in Figure 4.21(b). If the operational amplifier gain is assumed to be A_1 at $V_{\text{res}} = 0$ V, and A_2 at the point of maximum deviation, then d_{n} can be calculated by solely linear operations. Since this technique uses two linear gains to bound the nonlinearity, it is simply called the two-gain technique in this dissertation.

For example, for a 2.8-bit stage implemented with the residue amplifier of Figure 5.21, and considering the rightmost segment (the leftmost yields a similar result), it can readily be shown that:[19]

$$d_{\text{n}} = 2C_8 \left[\frac{A_1}{(1 + A_1)C_{\text{F}} + C_{\text{P}}} - \frac{A_2}{(1 + A_2)C_{\text{F}} + C_{\text{P}}} \right] V_{\text{REF}} \tag{4.27}$$

Note that if $A_1 = A_2$ then d_{n} goes to 0 as expected. A contour plot of d_{n} (assuming C_8 and C_{P} values as per Figure 5.21) is shown in Figure 4.22. The maximum gain A_1 (in dB) is plotted along the horizontal axis, while the gain loss $A_1 - A_2$ (in dB) is plotted along the vertical axis. Contours of d_{n} are labeled. As an example, for 11-bit maximum distortion an operational amplifier with a maximum gain of 65 dB can experience a gain loss of at most 6 dB over its output range. Naturally, raising or lowering A_1 relaxes or tightens the gain loss requirement, respectively.

This dual-gain method readily characterizes the permissible operational amplifier

[19] For simplicity, this analysis ignores capacitor mismatch and operational amplifier input-referred offset.

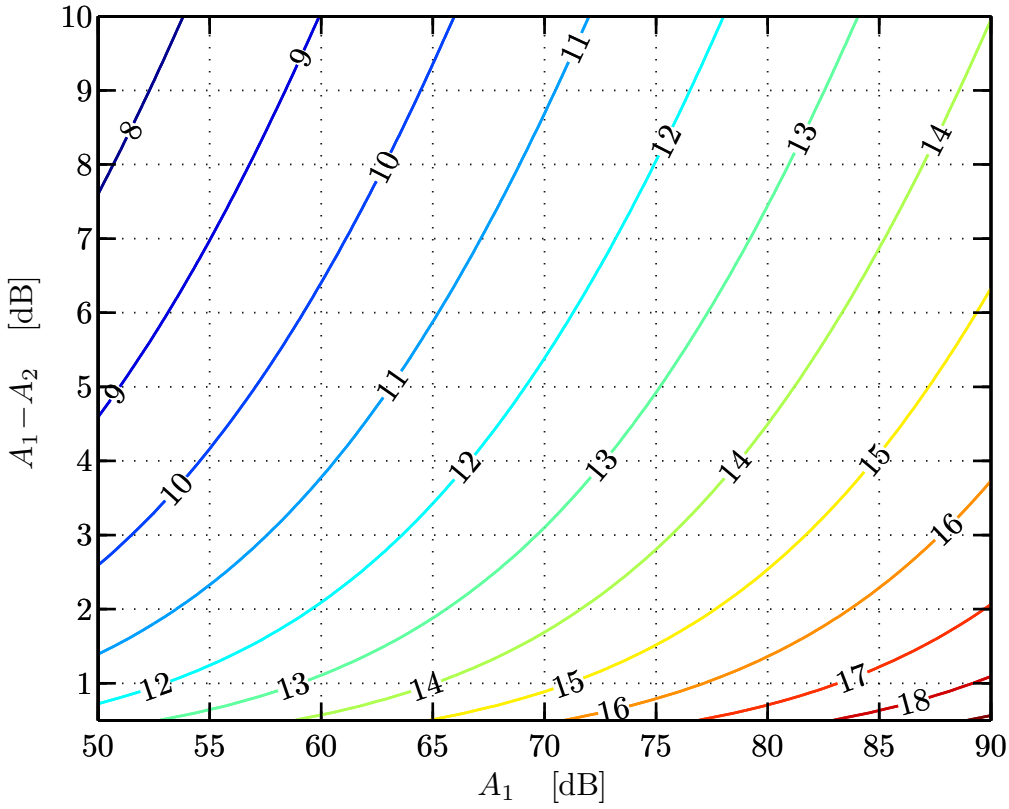

Figure 4.22: d_n computed for a 2.8-bit stage residue amplifier implemented as per Figure 5.21. Contours of constant d_n over maximum gain A_1 and gain loss $A_1 - A_2$ shown. d_n is given in bits: for example, the "11" contour corresponds to $d_n = 2V_{REF}/2^{11}$.

nonlinearity in terms useful for design and can be easily applied to other amplifier circuits.[20] In practice, d_n is typically chosen to be 1–2 bits better than the resolution of the backend ADC to satisfy linearity requirements without incurring egregious

[20]For example, repeating this analysis for the track-and-hold amplifier of Figure 5.16:

$$d_n = C_S \left[\frac{A_1}{(1+A_1)C_F + C_P} - \frac{A_2}{(1+A_2)C_F + C_P} \right] V_{REF}$$

and for the 3.1-bit stage residue amplifier of Figure 5.24:

$$d_n = 2(C_{11} + C_{12}) \left[\frac{A_1}{(1+A_1)C_F + C_P} - \frac{A_2}{(1+A_2)C_F + C_P} \right] V_{REF}$$

In both cases, the rightmost deviation is considered: the leftmost deviation is similar.

operational amplifier overdesign.

4.2.5 Noise

As it processes the signal, the residue amplifier also adds noise.[21] This noise arises from both the on-resistance of connected switches and the operational amplifier. These sources contribute two types of noise: flicker noise (contributed by the operational amplifier) and white noise (contributed by both sources).

Flicker noise is a potential problem in the SVADC-1 since the specifications of Table 3.1 require low noise performance down to 100 Hz. In submicron CMOS technologies, though, the flicker noise corners exhibited by MOSFET devices are usually hundreds of kiloHertz or higher, resulting in elevated noise at such low frequencies. To mitigate this flicker noise, the residue amplifier of Figure 4.18 employs offset cancellation: the switched-capacitor circuit is configured such that the operational amplifier noise is sampled at the end of ϕ_1, and this sampled value is used to cancel the operational amplifier noise during ϕ_2. It can easily be shown that the input-referred operational amplifier noise thus experiences a differentiation of $1 - z^{-1/2}$.[22] This high-pass shaping counteracts the flicker noise. Indeed, in the SVADC-1 offset cancellation sufficiently suppresses flicker noise that white noise sources dominate the residue amplifier noise over the signal bandwidth.

White noise contributions are often characterized by their total integrated (over frequency) noise power. In residue amplifiers such as that of Figure 4.18, after referral

[21] There is also noise in the sub-ADC. Input-referring it to before the sub-ADC, this noise is subsequently coarsely quantized, dulling its impact. However, if the input is near a transition level, this noise may create an "errant" decision. Nonetheless, so long as the resulting V_{res} does not overrange the stage output—unlikely under digital redundancy—the errant decision does not compromise the conversion: rather, the backend ADC provides a new V'_{res} that compensates for the offset introduced by the errant decision.

[22] This differentiation holds exactly if the operational amplifier gain is infinite. If the gain is instead of finite value A, then the input-referred noise is instead shaped by:

$$\frac{A^2(C_{\text{F}} + C_{\text{P}})}{(1+A)^2 C_{\text{F}} + (1+A)C_{\text{P}}} \left(1 - \frac{1+A}{A} z^{-1/2}\right)$$

Given the high gains A used in most designs, approximating with the simpler differentiation often proves more than serviceable in design.

to the proper nodes in the circuit, the complete white noise is shaped by low-pass noise transfer functions. The bandwidths of these transfer functions are typically much greater than the sampling rate,[23] resulting in significant aliasing of these noise contributions. As the total integrated noise is the same regardless of aliasing, it is often used to characterize this noise: the resulting noise power is then assumed evenly distributed across the converter output spectrum.[24]

Finally, recall that as discussed in Section 3.3.3 noise can prove a boon as well as a bane: on the one hand it raises the noise floor and degrades nhSFDR, on the other it linearizes through dither and improves hSFDR. In addition, noise proves useful in signal estimation during self-calibration, as is addressed in Section 4.4.2.2.

4.3 Correction

From an SFDR perspective, analog-digital mismatch in the gain and sub-DAC blocks—that is, mismatch between the analog stage gain and sub-DAC blocks and their corresponding representations in the digital reconstruction—most limits converter performance by disturbing linearity and elevating harmonic distortion. Conceptually, there are two general methods for correcting this mismatch: either the analog stage can be brought in line with the digital reconstruction, or vice versa. In this dissertation, the former are classified as analog correction techniques, the latter as digital correction techniques.[25]

[23] The noise transfer function bandwidths are usually tied to the circuit time constants, which are by design higher than the sampling rate to allow sufficient settling. In operational-amplifier-based amplifiers, settling is essentially exponential. Thus many time constants of settling are needed, especially in high resolution converters. As multiple time constants must nonetheless fit within half a sampling period, the circuit time constants are much greater than the sampling rate.

[24] While the even distribution of noise across the output spectrum is not strictly true, it is often close enough for practical design.

[25] Naturally, this division is somewhat artificial: it is possible to conceive of techniques that are a mix of both. Nonetheless, the division remains useful for organizational purposes here.

4.3.1 Analog Correction Techniques

While this dissertation concentrates on digital correction techniques, analog correction techniques have also been successful in restoring converter linearity. Some examples are presented here to give a flavor for this type of correction.

- *Trim sample capacitors*

  As the analog stage gain and sub-DAC blocks are often implemented by capacitor ratios, sample capacitor values can be adjusted by trim capacitors to perfect the analog stage. Placed in parallel to each sample capacitor, the trim capacitors are connected to either $V_{\mathrm{REF}+}$ or $V_{\mathrm{REF}-}$ to tune the effective sample capacitance [*Lin et al.*, 1991]. While this technique targets capacitor mismatch, it can be extended to compensate for finite operational amplifier gain as well. It requires taking the converter off-line to determine the trim values. Following this calibration, the stored trim values are applied during normal converter operation.

- *Capacitor-invariant gain of 2*

  This technique focuses on gain correction of 1-bit and 1.5-bit stages since, assuming differential circuitry, it is possible to implement a highly linear, 3-level sub-DAC with output voltages $+V_{\mathrm{A}}$, 0, and $-V_{\mathrm{A}}$.[26] Hence the analog-digital mismatch is dominated by gain error. It is possible, then, to design residue amplifiers that produce a gain of 2 independent of capacitor ratios [*Hatanaka and Taniguchi*, 2001; *Lee and Yan*, 2005]. By design invariant to capacitor mismatch, these amplifiers are often still sensitive to finite operational amplifier gain. While they do not require calibration, they do require 4 clock phases instead of 2 to produce V_{res}.[27]

[26] Conceptually, given a ϕ_2 implementation of $+V_{\mathrm{A}}$, $-V_{\mathrm{A}}$ can be had by swapping $V_{\mathrm{REF}+}$ and $V_{\mathrm{REF}-}$, effectively flipping the sign of V_{REF}. In practice, this "swapping" is simply done by inverting the $D\langle r\rangle$. To get the 0 output, $V_{\mathrm{REF}+}$ and $V_{\mathrm{REF}-}$ can be shorted together, sending V_{REF} to 0. This "shorting" is also simply done by adding an extra switch to a common mode voltage V_{CM}. All these new switch connections can be easily handled by digital logic.

[27] Ratio-independent amplifier designs are more popular in cyclic ADCs. Readers interested in this technique are recommended to *Li et al.* [1984] and *Lee* [1983].

- *Capacitor error-averaging*

  This technique also requires no calibration but entails additional clock phases. Applied to a particular 1-bit stage residue amplifier implementation where the sample capacitors reduce to just C_{1+} and C_{1-},[28] this technique uses swapping of the sample and feedback capacitors: in one configuration, $V_{res} = 2V_{in} \pm V_{REF} + \varepsilon$ where ε is a (possibly input-dependent) error, in the other, $V_{res} = 2V_{in} \pm V_{REF} - \varepsilon$. Averaging the two, then, eliminates ε. Averaging can be accomplished either by an active switched-capacitor amplifier (dubbed active capacitor error-averaging [*Song et al.*, 1988; *Chen et al.*, 2001]), or by the sample capacitors of the next stage (dubbed passive capacitor error-averaging [*Chiu et al.*, 2004]). The former requires 3 clock phases, the latter 4, to produce V_{res}. Both active and passive capacitor error-averaging, though, remain sensitive to finite operational amplifier gain.

4.3.2 Digital Correction Techniques

Whereas analog correction techniques modify the analog circuitry to better match the digital reconstruction, digital correction techniques modify the more malleable digital reconstruction to better match the analog circuitry. Conceptually, digital correction techniques measure the analog behavior of the stages and process these measurements to create correction factors. These correction factors are then incorporated into the digital reconstruction.

Generally, digital correction techniques inject a known signal into the analog pipeline and measure the response. For this reason, they are often referred to as digital calibration. Note that an arbitrary converter can be digitally calibrated by taking the converter off-line and inputting, say, a single tone sinusoid. The output can then be processed to produce a calibration table that maps the converter outputs

[28] In particular, capacitor error-averaging builds upon a capacitor-flip-over residue amplifier. In these residue amplifiers, the C_{F+} and C_{F-} of Figure 4.18 are also used as a sample capacitors during ϕ_1. During ϕ_2, C_{F+} and C_{F-} are "flipped-over" to provide the negative feedback capacitance; the amplifier is configured as shown in Figure 4.18 in this phase. For a 1-bit stage, a capacitor-flip-over residue amplifier requires just 2 sample capacitors C_{1+} and C_{1-} and 2 feedback capacitors C_{F+} and C_{F-}, all of which are nominally equal.

to new output levels, improving both SNDR and SFDR [*Händel et al.*, 2000]. This technique, though, requires generating a pure tone sinusoid which can be impractical if the converter is to be calibrated in the field. Two classes, then, of digital calibration techniques that instead employ more easily generated signals are correlation-based techniques and self-calibration techniques.

Correlation-based techniques inject lengthy pseudo-random number sequences (PRNSs) into the pipeline. These signals are digital sequences of -1 and $+1$ with known statistics. The PRNS can be directly injected into the signal path (e.g., *Siragusa and Galton* [2000]), or used to modulate an architectural aspect of a stage (e.g., *Li and Moon* [2003]; *Murmann and Boser* [2003]).[29] In both cases, it is assumed that the PRNS is uncorrelated with the converter input. Thus in the converter output the PRNS can be statistically isolated from the converter input and processed to estimate stage parameters such as gain (e.g., *Siragusa and Galton* [2000]; *Li and Moon* [2003]) or even nonlinearity (e.g., *Murmann and Boser* [2003]). These parameters then inform the digital reconstruction, which is sometimes heavily modified to better accommodate them. Typically, correlation-based techniques operate in the background: the ADC can simultaneously convert input signals during calibration. However, correlation-based techniques can require complex digital logic to process the PRNS and a large number of cycles to converge.

Self-calibration techniques interrupt normal converter operation to inject the calibration signal. Typically, the injected signals are easily generated analog signals (such as 0 V in fully differential circuitry) or predetermined sub-DAC input codes. These signals are routed within the stage so that V_{res} corresponds to some stage parameter. This parameter voltage is digitized by the backend ADC and, following processing, the resulting calibration coefficient incorporated into the digital reconstruction. Hence the parameters of the pipeline are effectively measured by the pipeline itself: the pipeline "self-calibrates".[30] Classically, self-calibration schemes initiate calibration with a downstream stage and proceed stage-by-stage upstream

[29] Specifically, *Galton* [2000] modulate the capacitors corresponding to a DAC code, and *Murmann and Boser* [2003] modulate an overall offset of the sub-DAC voltages.

[30] Less glibly: the parameters of a pipeline stage are measured by its backend ADC.

to ensure the backend ADC is as accurate an estimator as possible.[31] While self-calibration techniques often operate in the foreground,[32] they usually exploit particulars of the residue amplifier and digital reconstruction so that the parameter voltage can be generated, and the calibration coefficients accommodated, with little complexity increase in each.

Correlation-based techniques have many benefits, not the least of which is background operation. Unfortunately, they often assume some knowledge of the converter input signal statistics, for example, that it is uncorrelated with the PRNS. To enable use of the SVADC-1 in a broad array of applications both satellite and terrestrial, it was elected that its performance should not rely on assumptions of the input signal statistics. To some extent, the motivation is tautological: it is the explicit objective of many studies to indeed investigate these statistics. More practically, in many terrestrial applications, the input signal may include strong man-made signals from communication and positioning systems—indeed, it is often desirable to capture these signals alongside the phenomena for timing or analysis purposes—and these signals may use PRNSs in their protocols. Hence, the SVADC-1 instead employs a novel self-calibration technique to correct analog-digital mismatch and improve converter linearity. Self-calibration techniques are thus described in more detail in the following subsections.[33]

[31] Although it is also possible to begin calibration with an upstream stage and proceed downstream [*Singer et al.*, 2000].

[32] It is possible to design self-calibration techniques that instead work in the background. For example, *Moon and Song* [1997] propose a skip-and-fill technique that opts to occasionally ignore an input sample ("skip") to inject a self-calibration signal. The missing input sample is then interpolated ("fill") by a 44th-order polynomial: externally, converter operation appears unbroken (see *Kwak et al.* [1997] for a chip implementation).

[33] Naturally, there are digital calibration techniques beyond correlation-based and self-calibration methods. For example, another technique is to use a separate calibration corral [*Ingino and Wooley*, 1998]. Here, a stage is brought "off-line" and calibrated in a separate "corral" while the remainder of the pipeline continues operation. Being off-line, the stage can be slowly and accurately calibrated. Upon completion of calibration, the stage is restored to the pipeline and another stage rotated into the corral.

4.3.2.1 Discontinuity Height Sufficiency

Self-calibration techniques often focus on measuring the discontinuity heights in the analog stage transfer function. Intuitively, under the segment realignment interpretation of digital reconstruction (Section 4.1.2) and recognizing that overall offsets in the digital reconstruction lookup table do not impugn linearity (Section 4.2.2),[34] it would seem that knowledge of the discontinuity heights is sufficient for linear reconstruction. This intuition is formalized here by an argument similar to that first presented by *Rombouts and Weyten* [1999].

Consider the two-stage pipeline of Figure 4.23. Most self-calibration techniques begin calibration downstream stage and proceed stage-by-stage upstream: the first stage thus represents the stage-to-be-calibrated, the second the already-calibrated backend ADC. Thus the first stage is unknown, in particular, G_1 and $V_{\mathrm{DAC},1}$ are not known. On the other hand, the second stage is assumed well-known to within a linear recombination, that is, $V_{\mathrm{DAC},2}$ is well known (so that $V'_{\mathrm{DAC},2} = V_{\mathrm{DAC},2}$), but F_2 and $V_{\mathrm{OFF},2}$—residual ambiguities from the linear reconstruction—are not. Furthermore, the second stage is assumed linearized so that V_{err} is uniform.

Let $H_{m,1}$ represent the discontinuity heights of the first stage, that is:

$$H_{m,1} = G_1 \left(V_{\mathrm{DAC},1}[m+1] - V_{\mathrm{DAC},1}[m] \right) \quad , \quad m \in [0, \max\{B_1\} - 1] \tag{4.28}$$

Generally, self-calibration techniques use signal injection and stage rerouting to estimate the $H_{m,1}$ with the backend ADC. In the context of Figure 4.23, such strategies can be modeled as causing the $H_{\mathrm{m},1}$ to appear at the output of the first stage. The second stage then estimates the $H_{m,1}$, producing:

$$H'_{m,1} = V'_{\mathrm{res},1}\big|_{V_{\mathrm{res},1}=H_{m,1}} = F_2 H_{m,1} + V_{\mathrm{OFF},2} + e_{\mathrm{est}}[m] \tag{4.29}$$

where $e_{\mathrm{est}}[m]$ is the estimation error stemming from quantization and any other signal processing. Typically, self-calibration techniques use estimation methods that

[34] And also assuming operational amplifier gain nonlinearity is negligible.

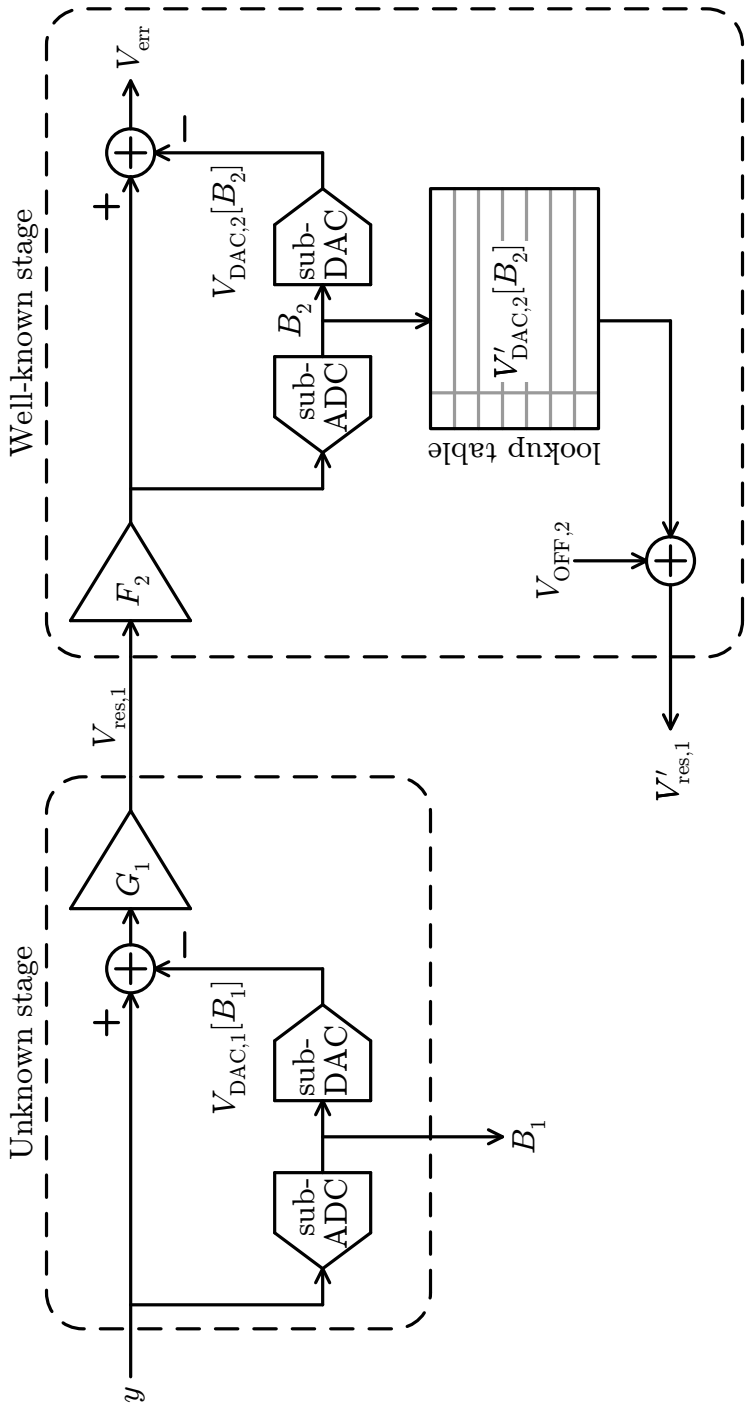

Figure 4.23: Two-stage pipeline model for demonstrating sufficiency of discontinuity height measurement for linear reconstruction.

eliminate residual offsets, so that $V_{\text{OFF},2}$ is removed and Equation (4.29) becomes:

$$H'_{m,1} = F_2 H_{m,1} + e_{\text{est}}[m] \tag{4.30}$$

While the $H'_{m,1}$ estimates retain an F_2 ambiguity, they are nonetheless sufficient for linearly reconstructing the pipeline. To see this fact, assume digital reconstruction uses only known quantities. Then, proceeding from Equation (4.14), digital reconstruction proceeds as:

$$y' = \sum_{m=0}^{B_1-1} H'_{m,1} + V'_{\text{DAC},2}[B_2] \tag{4.31}$$

where the gaining of y' in Equation (4.14) is simply ignored. It can be shown that the pipeline itself operates as:[35]

$$y = \frac{1}{G_1 F_2}\left(F_2 G_1 V_{\text{DAC},1}[0] + F_2 \sum_{m=0}^{B_1-1} H_{m,1} + V_{\text{DAC},2}[B_2] + V_{\text{err}}\right) \tag{4.32}$$

Thus, combining Equations (4.30) through (4.32):

$$y' = G_1 F_2 \left[y - \frac{1}{G_1 F_2}\left(V_{\text{err}} - \sum_{m=0}^{B_1-1} e_{\text{est}}[m]\right)\right] - G_1 F_2 V_{\text{DAC},1}[0] \tag{4.33}$$

Assuming the accumulated estimation error $\sum_{m=0}^{B_1-1} e_{\text{est}}[m]$ is negligible compared to V_{err} (a detail addressed in Section 4.4.2.1):

$$y' \simeq G_1 F_2 \left(y - \frac{V_{\text{err}}}{G_1 F_2}\right) - G_1 F_2 V_{\text{DAC},1}[0] \tag{4.34}$$

As V_{err} is uniform by assumption, Equation (4.34) represents a linear reconstruction (compare Equation (4.17)). Thus, given a linearly reconstructed backend ADC,

[35] Essentially, begin with Equation (4.12), substitute in Equation (4.28), and rearrange the result to factor out the gain $1/(G_1 F_2)$.

by $H_{m,1}$ estimation the stage-and-backend-ADC system can also be linearly reconstructed. Notably, knowledge of the gains and sub-DAC voltages themselves is not necessary: rather, it is a specific combination of the two—in particular, to create the $H_{m,p}$—that needs to be known.

Repeating this argument upstream, a complete pipeline can be linearly reconstructed with accurate knowledge of the $H_{m,p}$ alone. This induction does require a starting backend ADC with uniform V_{err}: this starting ADC can be the terminating ADC or some collection of downstream stages. Notably, V_{err} is attenuated in Equation (4.34): should V_{err} be nonuniform, this nonuniformity is at least diminished by approximately the total upstream gain.

The following sections describe different techniques for measuring $H_{m,p}$. It is worthwhile to keep in mind, then, that Equation (4.34) requires that residual backend ADC offsets (i.e., $V_{\text{OFF},2}$) be eliminated during $H_{m,p}$ measurement, and that the accumulated estimation error be small.

4.3.2.2 Karanicolas Technique

A classic self-calibration technique is that advocated by *Karanicolas et al.* [1993] for a 1-bit stage. As illustrated in Figure 4.24, such a stage has a single discontinuity height H_0. During calibration, instead of processing the input sample during ϕ_1, the residue amplifier is reconfigured to instead process a 0 V signal. During the subsequent ϕ_2, B is forced to a DAC code of either 0 or 1. Let S_0 and S_1 be the resulting V_{res}, respectively. Then:

$$\begin{aligned} S_0 &= V_{\text{res}}|_{B=0} = -G\,V_{\text{DAC}}[0] \\ S_1 &= V_{\text{res}}|_{B=1} = -G\,V_{\text{DAC}}[1] \end{aligned} \tag{4.35}$$

The locations of S_0 and S_1 in the stage transfer function are illustrated in Figure 4.24: they are the top and bottom, respectively, of the discontinuity. Thus their difference:

$$S_0 - S_1 = G\,(V_{\text{DAC}}[0] - V_{\text{DAC}}[1]) \tag{4.36}$$

gives H_0. In this way, the Karanicolas technique requires 2 unique measurements to determine the discontinuity height.

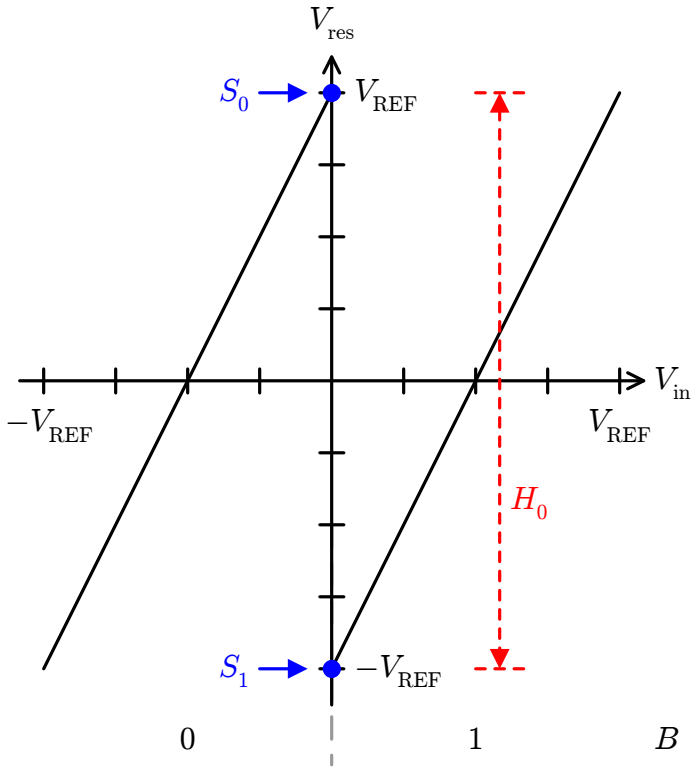

Figure 4.24: Illustration of S_0 and S_1 measurement (as per Equation (4.35)) on a 1-bit stage in the Karanicolas technique.

Practically, to measure S_0, say, the stage is preempted for several cycles and the backend ADC output averaged. While the analysis of this processing is addressed in Section 4.4.2, suffice to say here that the averaging reduces measurement noise and can sometimes permit estimation of S_0 to better than the backend ADC resolution, reducing estimation error. The same is repeated for S_1 and the S_0 and S_1 estimates then subtracted in the digital domain to derive H_0; note that this subtraction has the additional benefit of eliminating any residual backend ADC offset. Overall, the digital domain impact is fairly minor: just an averager and subtractor are needed to generate the updated lookup table entries.[36] The analog domain impact is also fairly minor. Reconfiguration of the residue amplifier of Figure 4.18 to support an S_0 estimation is shown in Figure 4.25; the reconfiguration for S_1 is similar. Additional

[36]Furthermore, if averaging takes place over a power-of-2 number of samples, the averager reduces to just an accumulator and bitshift [*Mayes and Chin*, 1996].

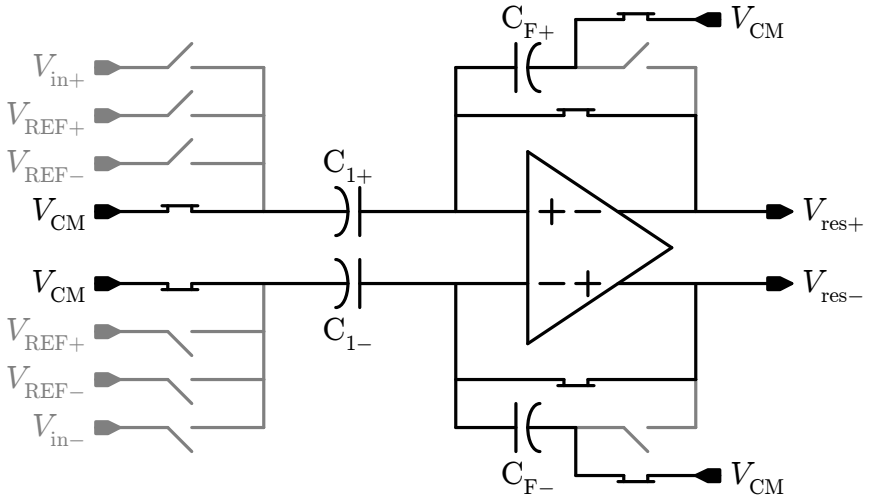

(a) ϕ_1.

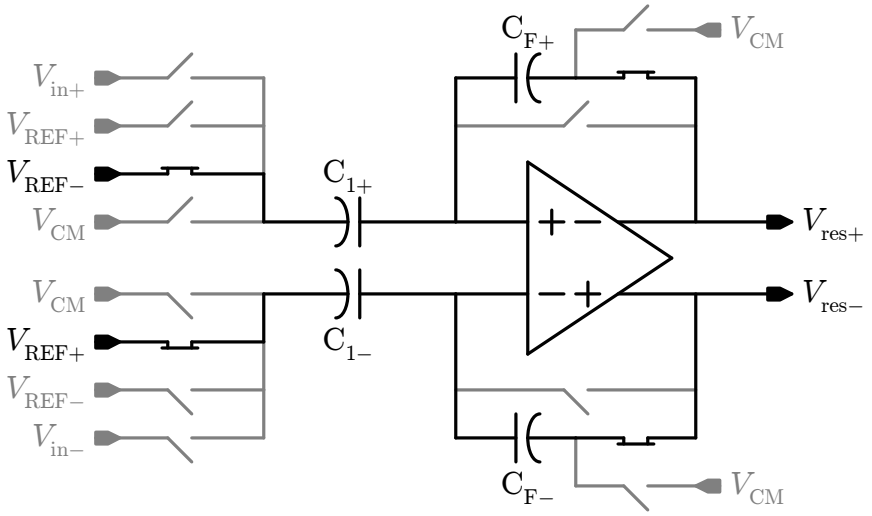

(b) ϕ_2.

Figure 4.25: Karanicolas technique applied to the residue amplifier of Figure 4.18 (which is configured as for a 1-bit stage). Residue amplifier operation during (a) ϕ_1 and (b) ϕ_2 shown for S_0 measurement.

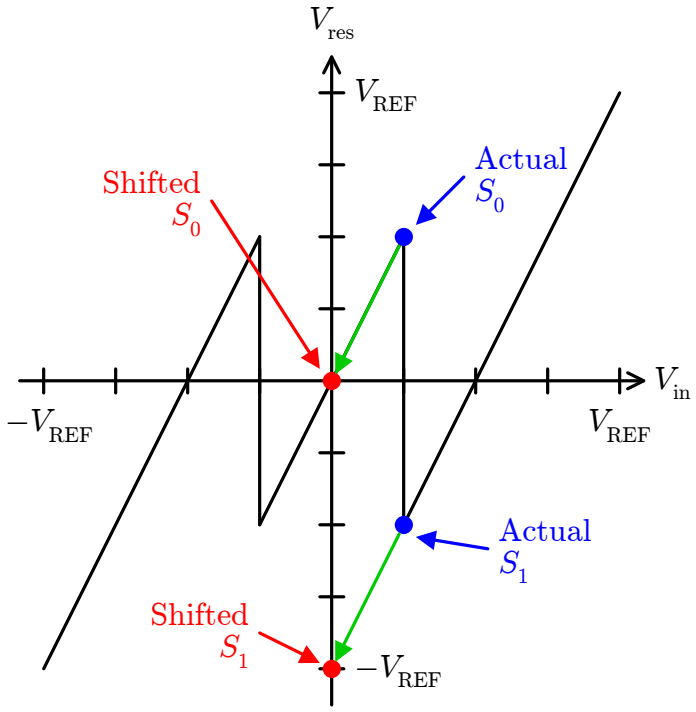

Figure 4.26: Generalization of the Karanicolas technique to a 1.5-bit stage. Actual S_0 and S_1, as would be measured by application of $V_{\text{REF}}/4$ during ϕ_1, shown in blue. Shifted S_0 and S_1, as would be measured by application of 0 V during ϕ_1, shown in red.

switches (one per branch) to a common mode voltage V_{CM} easily accomplish the 0 V needed during ϕ_1. And forcing B to 0 or 1 during ϕ_2 is readily achieved in digital logic. Finally, it is notable that in the ideal 1-bit stage S_0 and S_1 are nominally V_{REF} and $-V_{\text{REF}}$, respectively, and hence highly susceptible to overrange. In overrange, the estimation error can dramatically increase, compromising the technique. To prevent overrange, *Karanicolas et al.* [1993] intentionally reduces the stage gain from 2 to 1.93.

In theory, the Karanicolas technique can be generalized to measure the H_m of arbitrary stage transfer functions. Strictly, though, this extension requires injection of the input voltage at each discontinuity during ϕ_1: for example, a 1.5-bit stage should inject $-V_{\text{REF}}/4$ for one, and $+V_{\text{REF}}/4$ for the other, of its discontinuities [*Chuang and Sculley*, 2002]. If other ϕ_1 input voltages are used, then the locations of

the measured S_0 and S_1 are effectively shifted from their actual locations along the extension of their appropriate transfer function segments, as illustrated in Figure 4.26. However, if the backend ADC is not uniform (not unlikely in actual implementations) the S_0 and S_1 estimates at these shifted locations may differ from the estimates at the discontinuity proper. The increased $e_{\text{est}}[m]$ errors can compromise calibration performance, complicating the extension of the Karanicolas technique to arbitrary transfer functions [*Law et al.*, 2003].[37]

4.3.2.3 DAC-Differencing Technique

An alternative to the Karanicolas technique that is more amenable to higher-bit stages is DAC differencing. This technique takes advantage of the biphasic nature of the residue amplifier: the amplifier is rerouted so that instead of subtracting a sub-DAC voltage from V_{in}, it instead subtracts two sub-DAC voltages. In particular, if the DAC code $m+1$ is applied during ϕ_1, and m applied during ϕ_2, then V_{res} at the end of ϕ_2 is:

$$V_{\text{res}} = G\left(V_{\text{DAC}}[m+1] - V_{\text{DAC}}[m]\right) \tag{4.37}$$

which is directly H_m (compare Equation (4.28)). The stage is again run for several cycles in this preempted state and the backend ADC output averaged. To eliminate any backend ADC offset, an autozero estimation—wherein the same DAC code is applied during both phases—is made and subtracted from the measurement of Equation (4.37).

To concretize this description, consider applying DAC differencing to the residue amplifier of Figure 4.18. Assuming operational amplifier nonidealities of finite gain A, input-referred offset V_{OFF}, and parasitic input capacitance C_{P}, and assuming positive-negative branch capacitor matching as per Equation (4.22), it can be shown that the

[37] As an aside, the increased $e_{\text{est}}[m]$ error arises from nonuniformity during estimation of the stage's *output* voltage during ϕ_2 due to shifting along the transfer function segments. But the shifting is controlled by the choice of the stage's *input* voltage during ϕ_1. In their paper, *Law et al.* [2003] opt to discuss the issue in terms of the input rather than output voltage. In particular, they consider the impact of sub-ADC comparator offsets: even if $-V_{\text{REF}}/4$ and $+V_{\text{REF}}/4$ are exactly generated for a 1.5-bit stage, say, comparator offsets mean that the actual S_0 and S_1 do not corresponding to these input voltages.

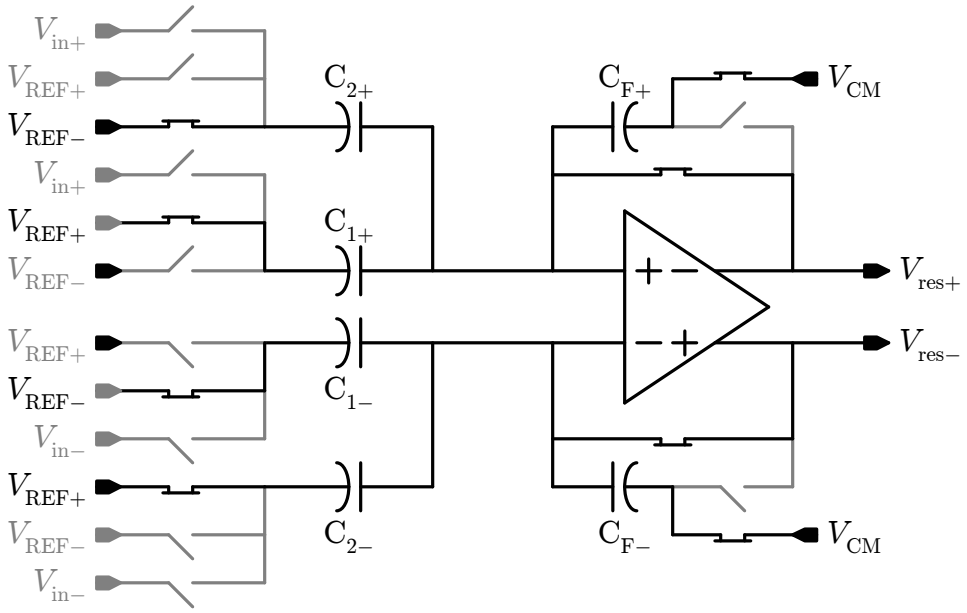

(a) ϕ_1.

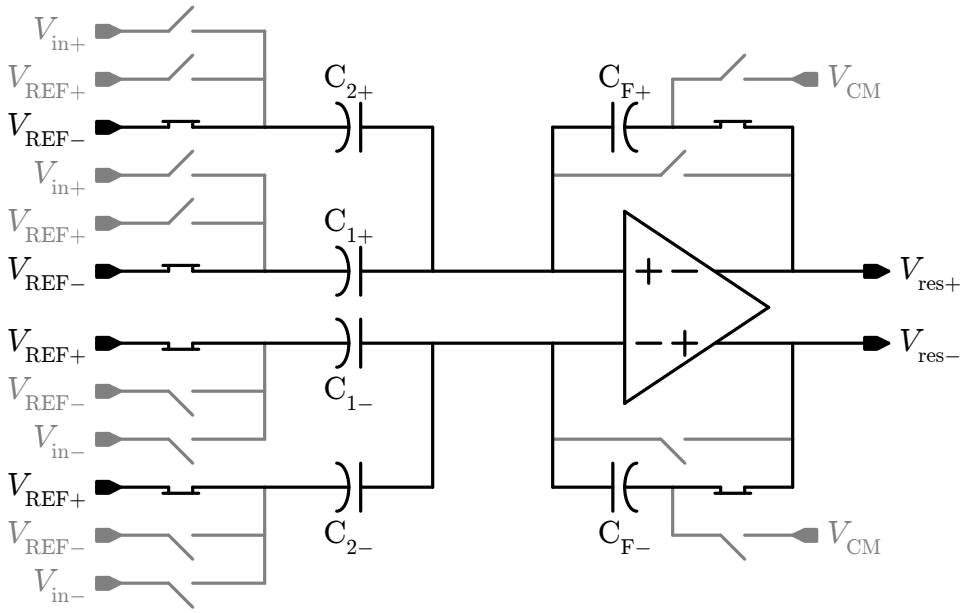

(b) ϕ_2.

Figure 4.27: DAC-differencing technique applied to the residue amplifier of Figure 4.18. Residue amplifier configured for a 3-level sub-DAC.[38] Operation during (a) ϕ_1 and (b) ϕ_2 shown for the $V_{\text{res}}|_{\phi_2:0}^{\phi_1:1}$ DAC-difference output.

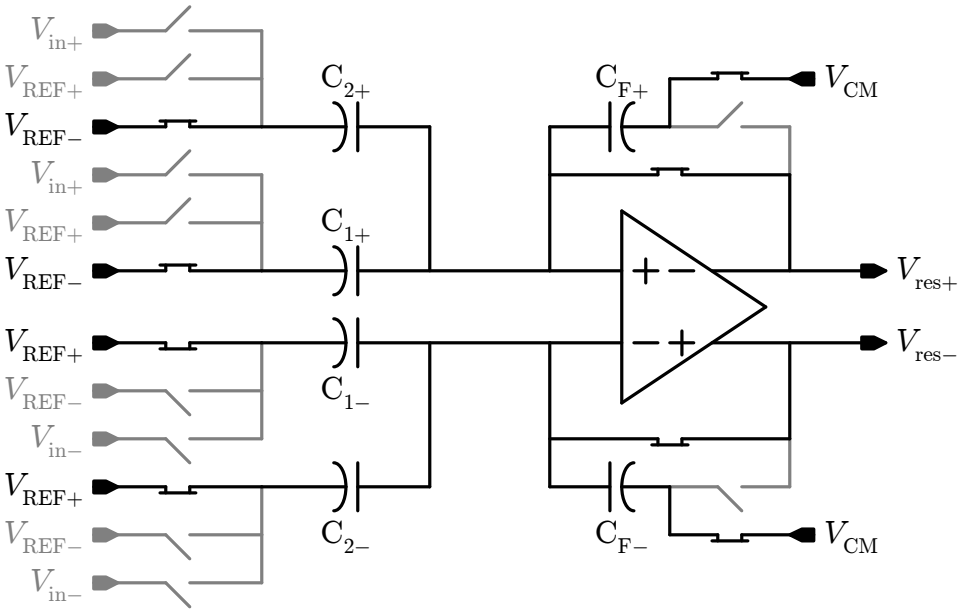

(a) ϕ_1.

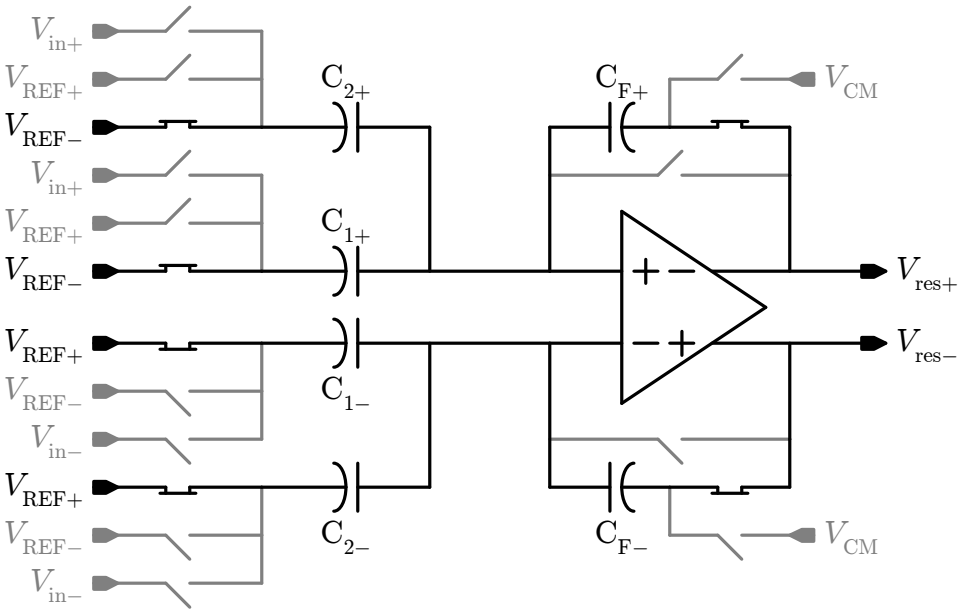

(b) ϕ_2.

Figure 4.28: DAC-differencing technique applied to the residue amplifier of Figure 4.18. Residue amplifier configured for a 3-level sub-DAC.[38] Operation during (a) ϕ_1 and (b) ϕ_2 shown for the $V_{\text{res}}|_{\phi_2:0}^{\phi_1:0}$ autozero output.

H_m of this residue amplifier are:

$$H_m = \frac{2AC_{m+1}}{(1+A)C_\mathrm{F} + C_\mathrm{P}} V_\mathrm{REF} \tag{4.38}$$

To calibrate H_m, the DAC codes $m+1$ and m are applied during ϕ_1 and ϕ_2, respectively. An example of the corresponding reconfiguration is shown in Figure 4.27 for a residue amplifier with a 3-level sub-DAC.[38] The amplifier output V_res at the end of ϕ_2 is:

$$V_\mathrm{res}|_{\phi_2:m}^{\phi_1:m+1} = \underbrace{\frac{2AC_{m+1}}{(1+A)C_\mathrm{F} + C_\mathrm{P}} V_\mathrm{REF}}_{H_m} - \underbrace{\frac{A(C_\mathrm{F}+C_\mathrm{P})}{(1+A)[(1+A)C_\mathrm{F} + C_\mathrm{P}]} V_\mathrm{OFF}}_{\text{offset}} \tag{4.39}$$

This output is called the DAC-difference output and is composed of H_m and an offset. To characterize both this offset and any backend ADC offset, a DAC code m is applied during both phases; Figure 4.28 continues the example of Figure 4.27 for this reconfiguration. The amplifier output is now:

$$V_\mathrm{res}|_{\phi_2:m}^{\phi_1:m} = -\frac{A(C_\mathrm{F}+C_\mathrm{P})}{(1+A)[(1+A)C_\mathrm{F} + C_\mathrm{P}]} V_\mathrm{OFF} \tag{4.40}$$

This output is called the autozero output. Notably, any DAC code can be used for the autozero output: without loss of generality, assume the 0 code here. Together Equations (4.38), (4.39), and (4.40) show that:

$$H_m = V_\mathrm{res}|_{\phi_2:m}^{\phi_1:m+1} - V_\mathrm{res}|_{\phi_2:0}^{\phi_1:0} \tag{4.41}$$

Hence, to derive the H'_m, the appropriate DAC-difference output and the autozero output are estimated by the backend ADC and their digital domain differences computed. Indeed, it can be shown that Equation (4.41) holds even under more general conditions, including full capacitor mismatch where every sample, feedback,

[38] For demonstration purposes, Figures 4.27 and 4.28 implement the 3-level sub-DAC by explicit sample capacitor switching, and not by reference switching as described in footnote 26 of this chapter. While the latter method is typically favored in practice, the former is nonetheless used in Figures 4.27 and 4.28 to illustrate the DAC-differencing principle.

and parasitic capacitor value is unique. The technique is thus robust both to capacitor mismatch and to operational amplifier nonidealities.[39]

Compared to the Karanicolas technique which requires $2N_{\mathrm{H}}$ unique measurements to characterize N_{H} discontinuities, DAC differencing requires only $N_{\mathrm{H}}+1$. Furthermore, DAC differencing does not require generation of particular analog voltages during ϕ_1, and does not require additional analog hardware since it simply repurposes existing switches (which is easily done in digital logic). However, in contrast to the Karanicolas technique, DAC differencing does not measure the actual discontinuity per se. Similar to the generalized Karanicolas technique under S_0 and S_1 shifting, then, DAC differencing is sensitive to backend ADC nonuniformity.

However, often more limiting is the susceptibility of DAC differencing to overrange. For the 1.5-bit and 2.8-bit stages, ideally $V_{\mathrm{res}}|_{\phi_2:m}^{\phi_1:m+1} = H_m = V_{\mathrm{REF}}$: given nonidealities, it is thus highly likely the DAC-difference output surpasses V_{REF} and overranges the backend ADC, increasing the estimation error and compromising the technique. To prevent overrange, previous DAC-differencing implementations forced the calibrated stage output range to be strictly less than the succeeding stage input range. One example is *Lee and Song* [1992], who reduced the calibrated stage gain by a full factor of 2. However, this solution was applied to a two-step converter:[40] such a solution applied to multiple pipeline stages would rapidly swell Q_{pipeline} and likely require extra pipeline stages (and hence additional power consumption) to recoup the resolution loss. Of course, less aggressive gain reductions—such as the 2 to 1.93 trim of *Karanicolas et al.* [1993]—are possible. However, under these reductions the DAC-difference output remains near the calibrated stage output extrema. As practical residue amplifiers often exhibit diminished performance near the output range extrema (due to operational amplifier nonlinearity, for example), calibration outputs near the extrema are best avoided.[41]

[39] Regarding the latter, it is robust to operational amplifier finite gain, parasitic input capacitance, and offset, but not to operational amplifier nonlinearity.

[40] A two-step converter (see footnote 37 of Chapter 3) can be thought of as a two-stage pipeline composed of a pipeline stage followed by a terminating ADC. Of course, the two-step converter is not necessarily built this way: when implemented, various architectural elements of the two stages are time-shared to reduce overall hardware.

[41] For completeness it should be mentioned that, instead of decreasing the output range of the

4.3.2.4 Novel DAC-Differencing Techniques

This dissertation presents three novel solutions for combating overrange in the DAC-differencing technique: increased sub-ADC transition levels, capacitor splitting, and capacitor-based offsetting. All decrease the DAC-difference output without altering the (calibrated or succeeding) stage signal ranges.

The first solution increases the number of sub-ADC transition levels (and correspondingly the number of sub-DAC levels). An example of this solution applied to a 2.8-bit stage is shown in the 3.1-bit stage of Figure 4.29. Graphically, increasing the number of sub-ADC transition levels translates to shorter segments and hence reduced H_m between segments. Any number of sub-ADC levels can be added: in Figure 4.29 enough are added so that the resulting H_m translate to DAC-difference outputs far from the output range extrema (nominally just $\frac{3}{4}V_{\text{REF}}$) while maintaining DAC voltages of the form $a/2^b$ for $a, b \in \mathbb{Z}$ for easier digital reconstruction implementation.[42] There is a hardware cost, though: compared to the 2.8-bit stage, the 3.1-bit stage requires 2 additional comparators, and also requires more measurements—6 DAC-difference and 1 autozero for the 2.8-bit, versus 8 and 1 for the 3.1-bit—to calibrate.

The second and third solutions rely on the superposition-over-sample-capacitor property of the DAC-difference output. To explain this property, assume positive-negative branch capacitor matching as per Equation (4.22) and consider the set of sample capacitors C_m where C_{m+} switches from $V_{\text{REF}+}$ during ϕ_1 to $V_{\text{REF}-}$ during ϕ_2 (and the corresponding C_{m-} switch from $V_{\text{REF}-}$ to $V_{\text{REF}+}$, respectively) and define:

$$\chi_{+\to-} = \sum_m C_m \tag{4.42}$$

Similarly, consider the set of sample capacitors C_l where C_{l+} switches from $V_{\text{REF}-}$

calibrated stage, overrange can also be avoided by increasing the input range of the succeeding stage. For instance, adding transition levels at $V_{\text{in}} = \pm\frac{7}{8}V_{\text{REF}}$ would increase the 2.8-bit stage input range to $\pm\frac{5}{4}V_{\text{REF}}$. However, this solution does not avoid the drawback that the DAC-difference output of the calibrated stage would lie near its output extrema.

[42] See footnote 10 of this chapter.

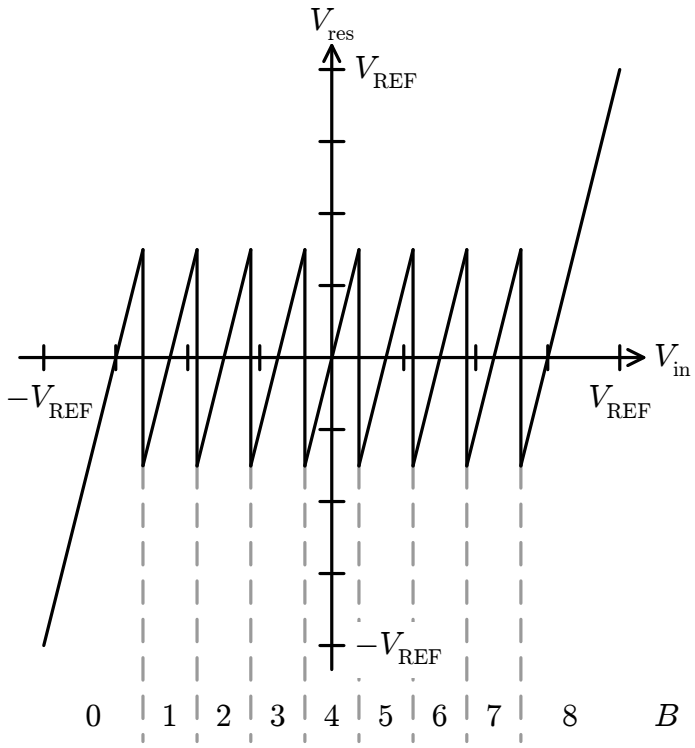

Figure 4.29: Transfer function of a 3.1-bit pipeline stage; compare Figure 4.14. sub-ADC transition levels: $-\frac{21}{32}V_{\text{REF}}, -\frac{15}{32}V_{\text{REF}}, \ldots, \frac{21}{32}V_{\text{REF}}$. sub-DAC output voltages: $-\frac{24}{32}V_{\text{REF}}, -\frac{18}{32}V_{\text{REF}}, \ldots, \frac{24}{32}V_{\text{REF}}$. Gain: 4.

during ϕ_1 to $V_{\text{REF}+}$ during ϕ_2 and define:

$$\chi_{-\to+} = \sum_l C_l \tag{4.43}$$

It can be shown that the DAC-difference output is then:

$$V_{\text{res}} = \frac{2A(\chi_{+\to-} - \chi_{-\to+})}{(1+A)C_{\text{F}} + C_{\text{P}}} V_{\text{REF}} - \frac{A(C_{\text{F}} + C_{\text{P}})}{(1+A)[(1+A)C_{\text{F}} + C_{\text{P}}]} V_{\text{OFF}} \tag{4.44}$$

Hence the DAC-difference output is the (signed) superposition of those sample capacitors whose connections change between ϕ_1 and ϕ_2: it essentially isolates the contribution of these changing sample capacitors. Those sample capacitors whose connections remain unchanged between ϕ_1 and ϕ_2, on the other hand, are effectively

ignored.

In capacitor splitting, a sample capacitor is divided into smaller capacitors such that the DAC-difference output isolating the smaller capacitors individually does not overrange. By Equation (4.44), the complete DAC-difference output of the sample capacitor is then the sum of the DAC-difference outputs of the smaller capacitors. If each sample capacitor is split into a smaller capacitors, then capacitor splitting requires $aN_{\mathrm{H}}+1$ unique measurements to characterize N_{H} discontinuity heights.

In capacitor-based offsetting, the signs of $\chi_{+\to-}$ and $\chi_{-\to+}$ are exploited, so that if $\chi_{+\to-}$ represents H_m, a small $\chi_{-\to+}$ is added to lower the value of the DAC-difference output. This offsetting $\chi_{-\to+}$ contribution can be incorporated into the autozero measurement. Hence capacitor-based offsetting requires as few as $N_{\mathrm{H}}+1$ unique measurements to characterize N_{H} discontinuity heights.

The SVADC-1 adopts the first solution of increasing the number of sub-ADC levels. Note that in capacitor-splitting and capacitor-based offsetting, the sample capacitor configurations during calibration are different from those during ϕ_2 in normal operation. Consequent differences in loading or clocking during calibration versus normal operation can thus result in increased estimation error. However, in the case of increased sub-ADC levels the sample capacitor configurations representing the DAC codes are exactly the same during both calibration and normal operation. Hence even though it incurs a higher analog and digital hardware cost, given the high linearity requirement of the SVADC-1 the more conservative approach of increased sub-ADC levels is adopted.

4.4 Design

With the novel DAC-differencing method of increased sub-ADC transition levels in hand, it is possible to develop a pipeline converter that meets the requirements of Table 3.1, and especially the very demanding SFDR requirement.

4.4.1 Converter Architecture

The converter architecture for the SVADC-1 is shown in Figure 4.30. Conversion is accomplished by five pipeline stages terminated by a 4-bit, midrise ADC. The first three pipeline stages are calibrated, adopting the 3.1-bit transfer function of Figure 4.29. The remaining two uncalibrated stages implement 2.8-bit stages as per Figure 4.14. While the architecture can achieve up to a 14-bit conversion, normally only a 12-bit conversion is utilized: the two LSB of the terminating ADC are ignored. Instead, as described in Section 4.4.2, the "extra" 2 bits of the terminating ADC are only used during calibration. Finally, the total input-referred circuit noise is designed to be 66 dB SNR_{FS} (assuming an input bandwidth of $f \in [0, f_{\text{S}}/2]$), the maximum that maintains the noise-floor dominated regime of Figure 3.20.

To improve power efficiency, the pipeline is scaled. Scaling exploits the fact that stage errors and noise, once input-referred to the beginning of the converter, are attenuated by the upstream gain.[43] Since the input-referred contribution of these inaccuracies decreases with downstream progression, and since a stage of higher accuracy typically consumes more power, there is impetus to relax the accuracy requirements of downstream stages to save on power consumption. The rate of relaxation, or "scaling", is often the subject of extensive optimization study (e.g., *Lewis* [1992]; *Cline and Gray* [1996]; *Goes et al.* [1998]).[44] For the

[43] Hence stages far downstream can incur even large errors without overly diminishing converter performance. Equivalently, it can be said that the backend ADC is of lesser resolution for further downstream stages: stage accuracy requirements hence diminish with downstream progression as fine errors are obscured by the increasing coarseness of the backend ADC.

[44] An oft-cited result is that of *Cline and Gray* [1996], which showed a broad power optimum if the residue amplifier capacitors are scaled by the gain of the stage. Briefly, *Cline and Gray* [1996] consider a pipeline wherein each stage has the same sub-ADC resolution and gain, and wherein the input-referred noise power of the pipeline is dominated by $k_{\text{B}}T/C$ noise from the residue amplifier capacitors. They then derive the optimum scaling of the residue amplifier capacitors to achieve a given input-referred noise budget while consuming minimal power. Two examples help highlight the argument:

No scaling. Under this scaling, the capacitor sizings of each stage are the same. From a power perspective the contribution of each stage is thus equal. From a noise perspective, while each stage produces the same amount of noise, once input-referred the noise contribution of the upstream stages dominates. But since the downstream stage noise contribution is minimal, these stages can be made more noisy with effectively no impact on the noise budget. But such a change is effected by reducing the capacitor sizes of these stages, which in turn can substantially reduce power. Hence this scaling is not optimal.

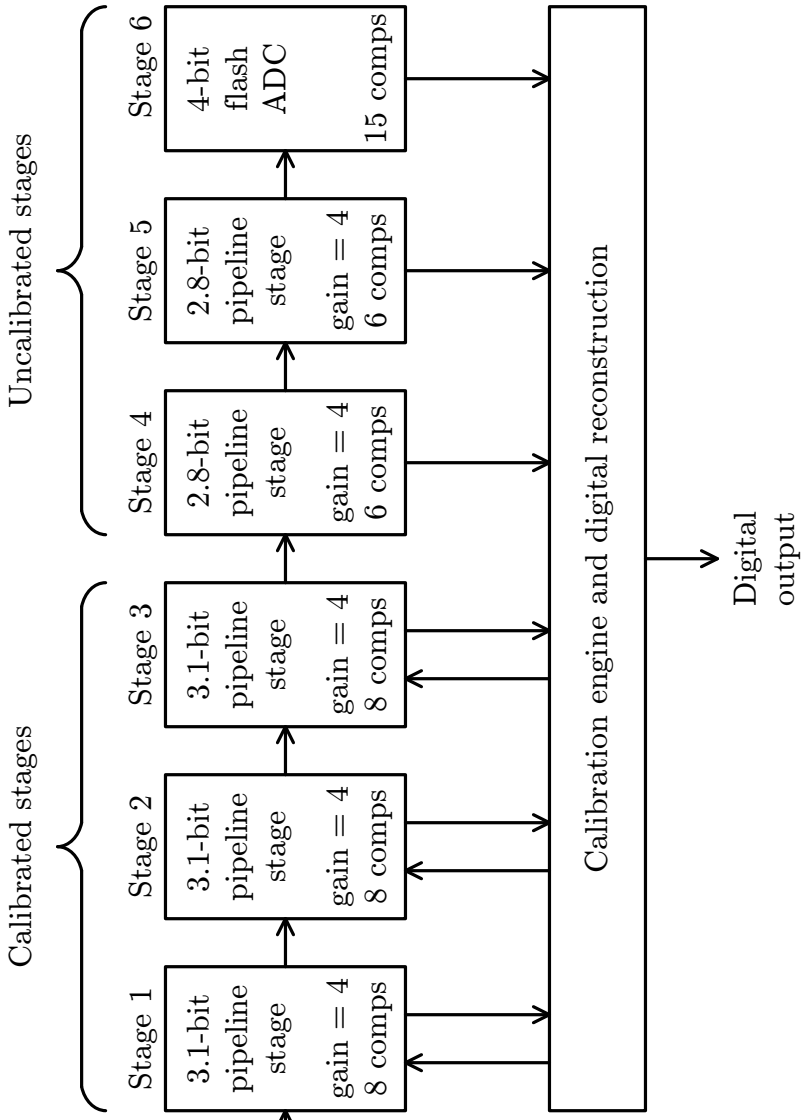

Figure 4.30: Architecture of the SVADC-1 pipeline converter. Gain and number of comparators in sub-ADC (comps) given for each stage (note that a stage with M comps thus produces $M+1$ output levels).

SVADC-1, aggressive scaling was generally eschewed.[45] Nonetheless, the pipeline of Figure 4.30 does employ residue amplifier operational amplifier power scaling between the calibrated versus uncalibrated stages,[46] and residue amplifier capacitor scaling between the first and second pipeline stages.[47] The former saves on power consumption, while the latter reduces the noise power contribution of the first pipeline stage. Otherwise, the calibrated stages are all the same, as are all the uncalibrated stages. The good radiation performance of the SVADC-1 (see Section 6.3), though, suggests that more aggressive scaling may be considered in future revisions.

A dedicated track-and-hold stage prefaces the pipeline proper. This stage is implemented by a biphasic switched-capacitor amplifier that tracks the analog input during ϕ_1—sampling said input at the end of ϕ_1—and holds the sampled value during ϕ_2. The track-and-hold stage thus provides a well-defined input for the first pipeline stage. Typically, track-and-hold stages must be of an accuracy and speed comparable to the entire converter, as such, they are often power hungry. Hence recently there has been much effort to eliminate them entirely (e.g., *Devarajan et al.* [2009]; *Jeon et al.* [2007]; *Mehr and Singer* [2000]). The SVADC-1, though, retains the track-and-hold stage for practical reasons. Note that the input impedance of a switched-capacitor

Scaling by the square of the stage gain. Under this scaling, the capacitor sizings of each stage are divided by the square of the gain of the preceding stage. From a power perspective the contribution of the upstream stages thus dominates (due to the larger capacitances). From a noise perspective, while the downstream stages produce increasing amounts of noise, once input-referred it can be shown that the noise contribution of all the stages is equal. But since the downstream stage power contribution is minimal, these stages can be made more power hungry with effectively no impact on the power budget. But such a change is effected by increasing the capacitors sizes of these stages, which in turn can substantially reduce noise. Hence this scaling is also not optimal.

Rather, the optimum scaling lies between these two extrema. While acknowledging that the actual optimum depends on a number of second-order effects (such as parasitic capacitances), to first order *Cline and Gray* [1996] showed that the optimum is roughly to scale the capacitors by the stage gain, and furthermore that this optimum is fairly broad.

[45] Under the operative philosophy that the more uniform its parts, the simpler it is to understand the radiation response of a complicated circuit.

[46] Specifically, the uncalibrated stage operational amplifiers consume roughly 3/5 the power of the calibrated stage operational amplifiers: compare Figures 5.29 and 5.30. The track-and-hold operational amplifier power consumption is on par with that of the calibrated stages; compare Figures 5.28 and 5.29.

[47] Specifically, the residue amplifier capacitors of the first calibrated stage are twice those of the other calibrated stages: see Figure 5.24.

amplifier changes abruptly when the input is sampled. In a pipeline stage sampling occurs once when the sub-ADC samples the input, and once when the residue amplifier does the same. Assuming the two sampling instants occur at different times—common practice to prevent them from interfering with each other—the circuit driving the pipeline stage must rapidly recover from the first sampling in time for the second. Integrating this driving circuit on chip as the track-and-hold stage relieves the ADC user of such concerns: instead, the user sees but a single sampling instant and ample time (roughly $T_S/2$) to recover from the consequent impedance change. However, future designs may wish to consider removing the track-and-hold stage to reap the consequent power savings.[48]

4.4.2 Calibration Technique

Calibration of the SVADC-1 begins with the furthest downstream stage—stage 3 (this choice is addressed in Section 4.4.2.2)—and proceeds upstream stage by stage.

To calibrate a stage, its residue amplifier is first configured to produce an autozero output. Any DAC code can be used: the $m=4$ code is chosen as thus half the sample capacitors connect to $V_{\mathrm{REF}+}$, half to $V_{\mathrm{REF}-}$. The residue amplifier is kept in this configuration for several cycles while the estimate $V'_{\mathrm{res}}|_{\phi_2:4}^{\phi_1:4}$ is computed by averaging the backend ADC output. To improve estimation accuracy (this choice is also addressed in Section 4.4.2.2), the terminating ADC is interpreted at full 4-bit resolution.

The calibration algorithm then proceeds to fill the digital reconstruction lookup table for the stage. As overall lookup table offsets do not harm linearity (Section 4.2.2), the $B=0$ entry is defaulted to 0. To derive the $B=m$ entry for $m>0$, the residue amplifier is configured to produce the DAC-difference output $V_{\mathrm{res}}|_{\phi_2:m-1}^{\phi_1:m}$. This output is again estimated by averaging of the backend ADC output and the terminating ADC again interpreted at full 4-bit resolution. As dictated by Equation 4.41, the resulting estimate $V'_{\mathrm{res}}|_{\phi_2:m-1}^{\phi_1:m}$ is combined with $V'_{\mathrm{res}}|_{\phi_2:4}^{\phi_1:4}$ to form a

[48] For the SVADC-1 as described in Chapter 5, directly removing the track-and-hold stage would result in a savings of about 5 mW from the total 36 mW consumed by $V_{\mathrm{DD,A}}$.

discontinuity height estimate H'_{m-1}:

$$H'_{m-1} = V'_{\text{res}}|^{\phi_1:m}_{\phi_2:m-1} - V'_{\text{res}}|^{\phi_1:4}_{\phi_2:4} \tag{4.45}$$

The $B = m$ entry is then the accumulation of the autozeroed DAC-difference output estimates for $B = m$ and all lower codes:

$$[B = m \text{ entry of lookup table}] = \sum_{k=0}^{m-1} H'_k \tag{4.46}$$

Given the accumulation, calibration begins with the $B = 1$ entry and proceeds towards greater B.

4.4.2.1 Accuracy Requirements

As noted in Section 4.3.2.1, calibration by H_m measurement works well so long as the measurement error is small. To better quantity "small", reconsider the two-stage pipeline of Figure 4.23. To simplify matters, ignore the second stage ambiguities (i.e., set $F_2 = 1$ and $V_{\text{OFF},2} = 0$). Ignoring overall gain and offset errors (being but linear deviations), the pipeline decomposition and reconstruction operate as (refer to Equations (4.32) and (4.31)):

$$y = \sum_{m=0}^{B_1-1} H_{m,1} + V_{\text{DAC},2}[B_2] + V_{\text{err}} \tag{4.47}$$

and:

$$y' = \sum_{m=0}^{B_1-1} H'_{m,1} + V'_{\text{DAC},2}[B_2] \tag{4.48}$$

Ideally, after calibration the only discrepancy between y and y' should be V_{err}, the quantization error of the second stage. However, assuming each $H'_{m,1}$ is estimated with error $e_{\text{est}}[m]$ as (refer to Equation (4.30)):

$$H'_{m,1} = H_{m,1} + e_{\text{est}}[m] \tag{4.49}$$

the discrepancy between y and y' expands to:

$$y - y' = V_{\text{err}} - \sum_{m=0}^{B_1-1} e_{\text{est}}[m] \tag{4.50}$$

For calibration to work properly the $e_{\text{est}}[m]$ contribution should be much smaller than the V_{err} contribution. Statistically, the greatest $e_{\text{est}}[m]$ contribution occurs when B_1 is maximum. In this case, assuming the $e_{\text{est}}[m]$ are i.i.d. with variance $\sigma^2_{\text{estimate}}$, and assuming V_{err} is distributed with variance $\sigma^2_{\text{backend}}$, for calibration to work properly:

$$\sigma^2_{\text{estimate}} \ll \frac{1}{\max\{B_1\}} \sigma^2_{\text{backend}} \tag{4.51}$$

That is, the H_m estimation resolution must be much better than the backend ADC resolution.[49]

4.4.2.2 Calibration Design

To improve H_m estimation resolution, the SVADC-1 employs both averaging of the backend ADC output and addition of extra bits during calibration.

In a quantizer with circuit noise, averaging of a quantizer output can produce estimates of its average input at accuracies better than its Δ.[50] This phenomenon is illustrated in Figure 4.31, which shows the absolute value of the mean of the error of such an estimate over varying averaging lengths assuming two different DC

[49] For example, if it is assumed that V_{err} is uniformly distributed over $\pm\Delta_{\text{backend}}/2$ (corresponding to the LSB of the backend ADC), and the $e_{\text{est}}[m]$ uniformly distributed over $\pm\Delta_{\text{estimate}}/2$ [*Lee and Song*, 1992], then Equation (4.51) simplifies to:

$$\frac{\Delta^2_{\text{estimate}}}{12} \ll \frac{1}{\max\{B_1\}} \frac{\Delta^2_{\text{backend}}}{12} \quad \Rightarrow \quad \Delta_{\text{estimate}} \ll 0.35\Delta_{\text{backend}}$$

which suggests that Δ_{estimate} (which incorporates both the estimation ADC resolution and any additional signal processing) should be at least 2 bits better than the backend itself.

[50] This phenomenon can be intuitively anticipated by a simple thought experiment. Assume a sign detector that outputs ± 1 (i.e., a 1-bit quantizer) with an input signal that is just slightly positive. If the system is noiseless, the quantizer output is constant, and the estimate of the input is 1. However, if the system is (i.i.d.) noisy, the quantizer output vacillates between 1 and -1 with proportions determined by the noise PDF. It should thus be possible to use these proportions—combined with some assumptions on the noise distribution—to estimate the input to better than ± 1 accuracy.

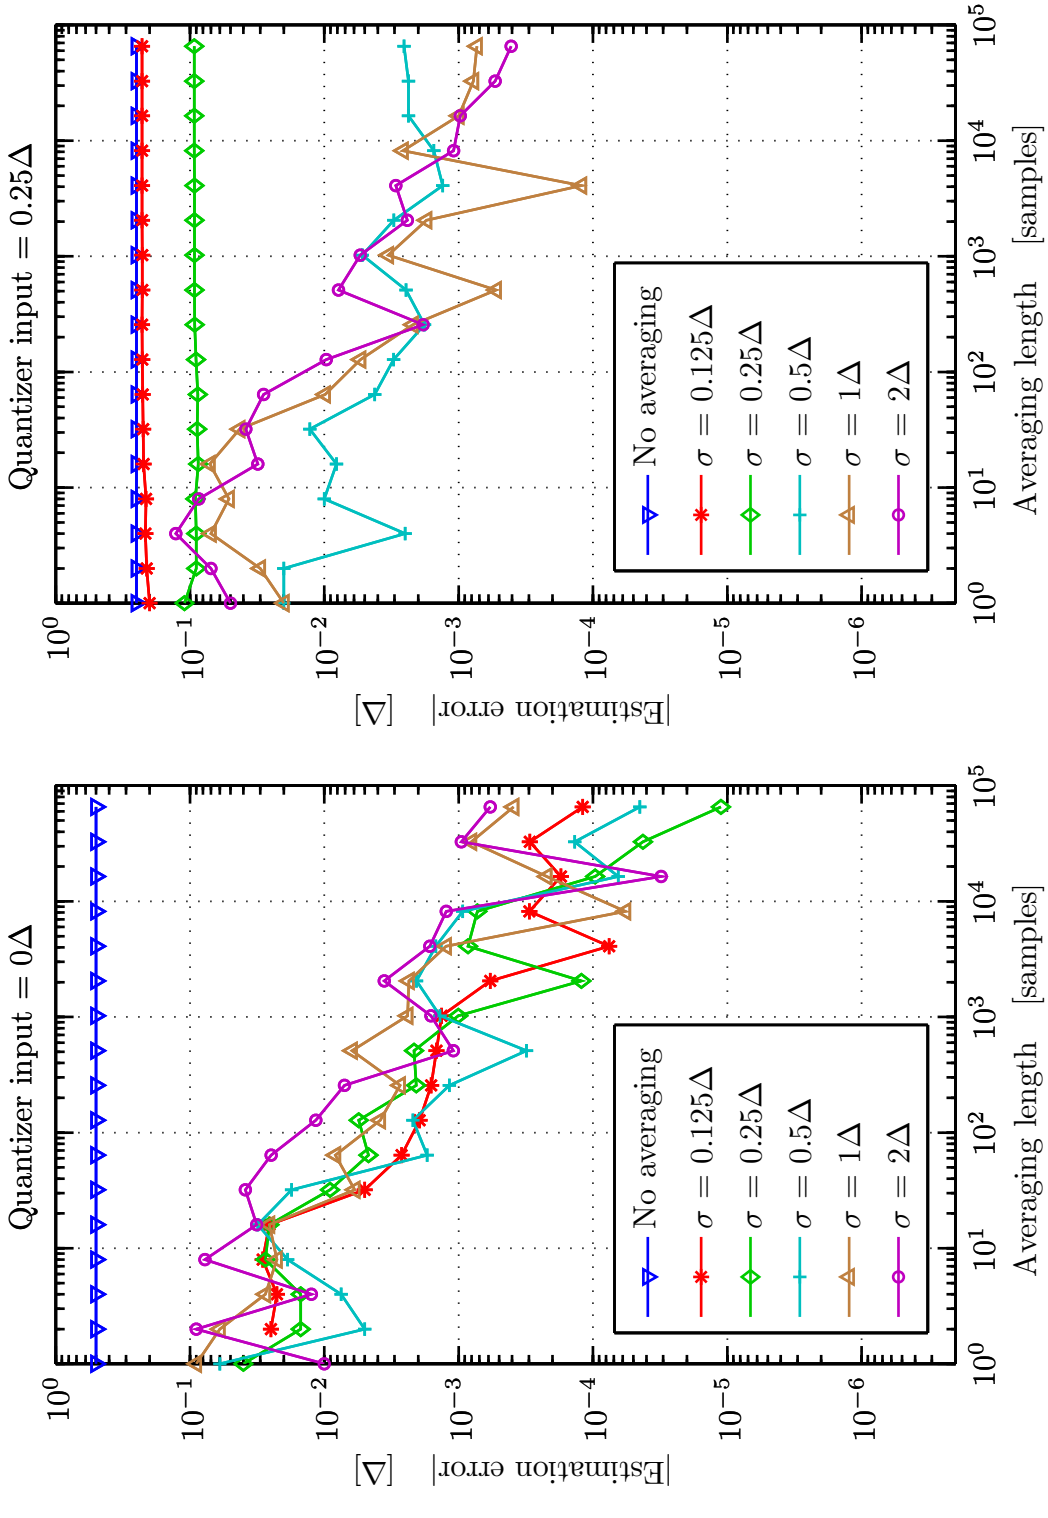

Figure 4.31: Effect of averaging length on estimation error of a uniform, unbounded quantizer of step size Δ for different quantizer inputs. Signal of DC input and noise (i.i.d. Gaussian of indicated standard deviation σ) input to quantizer and input estimated by averaging of quantizer output. Absolute value of average estimation error over 100 Monte Carlo runs plotted. Markers indicate measured data points.

inputs V_{in}. The results of Figure 4.31 are compiled over 100 Monte Carlo runs at each averaging length. As seen in the $V_{\text{in}} = 0\Delta$ (left) case, given sufficient circuit noise, averaging provides significantly more accurate estimates, with accuracy roughly improving with averaging length. However, the effect also depends on the DC input: in the $V_{\text{in}} = 0.25\Delta$ (right) case, the estimation accuracy rapidly levels out for low circuit noise levels (even becoming comparable to the unaveraged case, regardless of average length, at the lowest noise level).[51,52] To ensure estimation accuracy for all possible inputs and noise levels, then, not only is the backend ADC output averaged, but the backend ADC resolution is also improved by extra bits during calibration.

In addition to the averaging length and number of extra bits, the number of calibrated stages also affects SFDR. Notably, the number of extra bits and the number of calibrated stages directly affect the converter architecture. Their interaction is analyzed by simulation: the results are summarized in Table 4.1 which presents the expected average SFDR (and SFDR_5 and SFDR_{95}) of the pipeline under different combinations of the two. These simulations assume the converter is implemented as described in Chapter 5, and incorporate a breadth of nonidealities for all stages, including offsets in the sub-ADC transition levels, and digital-analog mismatch and noise in the residue amplifier.[53] Notably these simulations also account for nonuniformity of the backend ADC quantizing characteristic.[54]

[51] In terms of quantization theory, the argument is akin to that of dither in footnote 29 of Chapter 3. If the system is noiseless, the input signal PDF is $\delta(x - V_{\text{in}})$, whose flat characteristic function violates QT-II. If the system is noisy, the input signal PDF widens, narrowing the characteristic function, and better satisfying QT-II. Should QT-II be satisfied, then Sheppard's correction to the first moment applies, and the means of the backend ADC input and output are equal, regardless of quantizer resolution. But if the circuit noise is too small, QT-II is not well-satisfied, and the correction (as evidenced by Figure 4.31) becomes input-dependent.

[52] For interested readers, while the mean of the estimate saturates in the case of $V_{\text{in}} = 0.25\Delta$, the variance of the estimate decreases exponentially with the averaging length in both cases.

[53] Specifically, in the residue amplifier total integrated noise, capacitor mismatch, and operational amplifier finite gain, input-referred offset, and parasitic input capacitance are modeled. In the sub-ADC, input-referred offsets are modeled for each comparator. In each comparator, this offset is a combination of static offsets in the preamplifier (due to capacitor mismatch and amplifier finite gain, input-referred offset, and parasitic input capacitance) and dynamic offsets (due to transistor mismatch) in the dynamic latch. Notably, with the exception of the dynamic latch, complete settling is assumed in all switched-capacitor circuits. All nonidealities are computed assuming the circuit implementations described in Sections 5.4 and 5.5.

[54] In addition, since Table 4.1 uses nonidealities derived from circuit implementations, it should

Number of extra bits	Number of calibrated stages: 1	2	3	4
1	$78.9\ ^{87.4}_{69.5}$	$86.1\ ^{91.5}_{80.0}$	$84.9\ ^{90.1}_{80.3}$	$84.0\ ^{88.9}_{78.5}$
2	$79.2\ ^{85.4}_{72.6}$	$93.7\ ^{97.5}_{88.9}$	$93.7\ ^{96.8}_{88.2}$	$92.5\ ^{96.0}_{86.7}$
3	$78.5\ ^{85.3}_{70.2}$	$95.5\ ^{98.1}_{90.3}$	$97.8\ ^{98.3}_{97.0}$	$97.2\ ^{98.2}_{94.8}$
4	$78.4\ ^{86.5}_{70.3}$	$95.6\ ^{98.0}_{90.5}$	$98.1\ ^{98.4}_{97.7}$	$98.0\ ^{98.4}_{97.5}$

Table 4.1: Expected peak SFDR, in dB, of a calibrated converter given number of extra bits during calibration and number of calibrated stages. Each entry gives mean peak SFDR (and SFDR_{95} and SFDR_{5} in superscript and subscript, respectively) over 100 Monte Carlo test converters. For each test converter, peak SFDR is assessed by measuring the SFDR over a range of input signal amplitudes; for each amplitude, SFDR is defined as average SFDR assessed over 10 separate spectra. Statistics shown are then computed over these peak SFDRs. Input signal frequency placed on-bin in first nonzero-frequency FFT bin. SFDR computed assuming 50,000-point FFT, and spectrum assessed over complete $f \in [0, f_S/2]$ bandwidth. Long averaging (16,384 samples) used during calibration to eliminate averaging length effects.

To achieve 90-dB SFDR, Table 4.1 reveals that 2 extra bits and 2 calibrated stages are sufficient. The SVADC-1 adopts the former, but extends the latter to 3 calibrated stages. This margin is to maintain performance should the analog-digital mismatch prove worse than anticipated.[55] The result is the architecture of Figure 4.30. In this configuration, the pipeline SFDR is ultimately limited by the estimation accuracy: whereas increasing the number of bits during calibration increases the SFDR, increasing the number of calibrated stages does not.

be kept in mind that it is the end result of an iterative design process, wherein stage circuitry is designed, incorporated into architecture-level testing, and then adjusted as necessary.

[55] In the SVADC-1, analog-digital mismatch is mostly limited by interdigital capacitor matching. Unfortunately, the manufacture of these structures is not necessarily well-controlled during fabrication (see Section 5.3.3), motivating the heightened caution.

4.5 Conclusion

This chapter introduced the pipeline architecture for analog-to-digital conversion and focused on the quantization aspects thereof. As the pipeline adopts a staged approach to quantization, the chapter began by considering the operation of each stage algebraically and geometrically, and then cascaded several stages to describe the pipeline itself. This approach provided insights into the operation of both the analog stage and digital reconstruction portions of the pipeline, including the segment realignment interpretation of digital reconstruction that identified the transfer function discontinuity heights as integral to maintaining converter linearity.

This chapter then considered the effect of multiple analog nonidealities on converter performance. Many of these (such as overrange) can be compensated through well-understood means, while others (such as offset) do not damage conversion linearity. However, operational amplifier nonlinearity and analog-digital mismatch—and especially the latter—often limit converter linearity. To compensate for the former, this chapter introduced a two-gain technique for assessing the effect of the nonlinearity. To correct the latter, this chapter considered digital self-calibration techniques, in particular, DAC-differencing techniques which reroute the residue amplifier to measure the discontinuity heights of the analog stage transfer function via the backend ADC. DAC-differencing techniques, though, often suffer from overrange problems. Hence three novel solutions (increased number of sub-ADC transition levels, capacitor splitting, and capacitor-based offsetting) were proposed, and the first—DAC differencing applied to a stage with an increased number of sub-ADC transition levels—chosen for use in the SVADC-1. The resulting SVADC-1 pipeline converter architecture is shown in Figure 4.30. Designing this architecture required a consideration of the necessity of extra bits during calibration and the number of calibrated stages: simulations show that adopting 2 extra bits and 3 calibrated stages is sufficient for achieving the 90-dB SFDR required in Table 3.1.

The circuit-level implementation of the SVADC-1 is the topic of the next chapter. In addition to incorporating the self-calibration technique of this chapter, the SVADC-1 implementation also compensates for the radiation effects described in

Chapter 2. The result is a high-fidelity, radiation-hard ADC that satisfies the specifications of Table 3.1.

Chapter 5

Circuit Implementation

This chapter describes the implementation of the SVADC-1 pipeline converter. In particular, it describes the SVADC-1 chip that implements the analog portions of the converter as indicated in Figure 5.1.

The chapter begins by discussing the choice of manufacturing process. As the SVADC-1 is fabricated in a commercial manufacturing process, it adopts a radiation-hardness-by-design approach. Particular radiation-hardness-by-design techniques to prevent latchup and switch leakage are then explained. It should be stressed, though, that radiation-hardness by design is not confined to these techniques: it is instead a broader philosophy encompassing many design decisions—from circuit architectures to specification overdesign—all of which are made with radiation effects in mind. The applications of this broad philosophy should be clear from the next portions of the chapter, which establish some general conventions, and then delve stage by stage, and block by block, through the circuit implementation of the SVADC-1.

It should be noted that, although many design decisions are motivated by radiation concerns, the radiation performance of the manufacturing process was not known at design time.[1] Hence many decisions involving radiation effects were based upon worst-case expectations gleaned from the literature, for example, assuming significant radiation-induced g_m loss, in contrast to the measured results of Section 2.3.3.3.

[1] Indeed, for the vast majority of design time, the process itself was in beta release, stabilizing only a few months before tapeout.

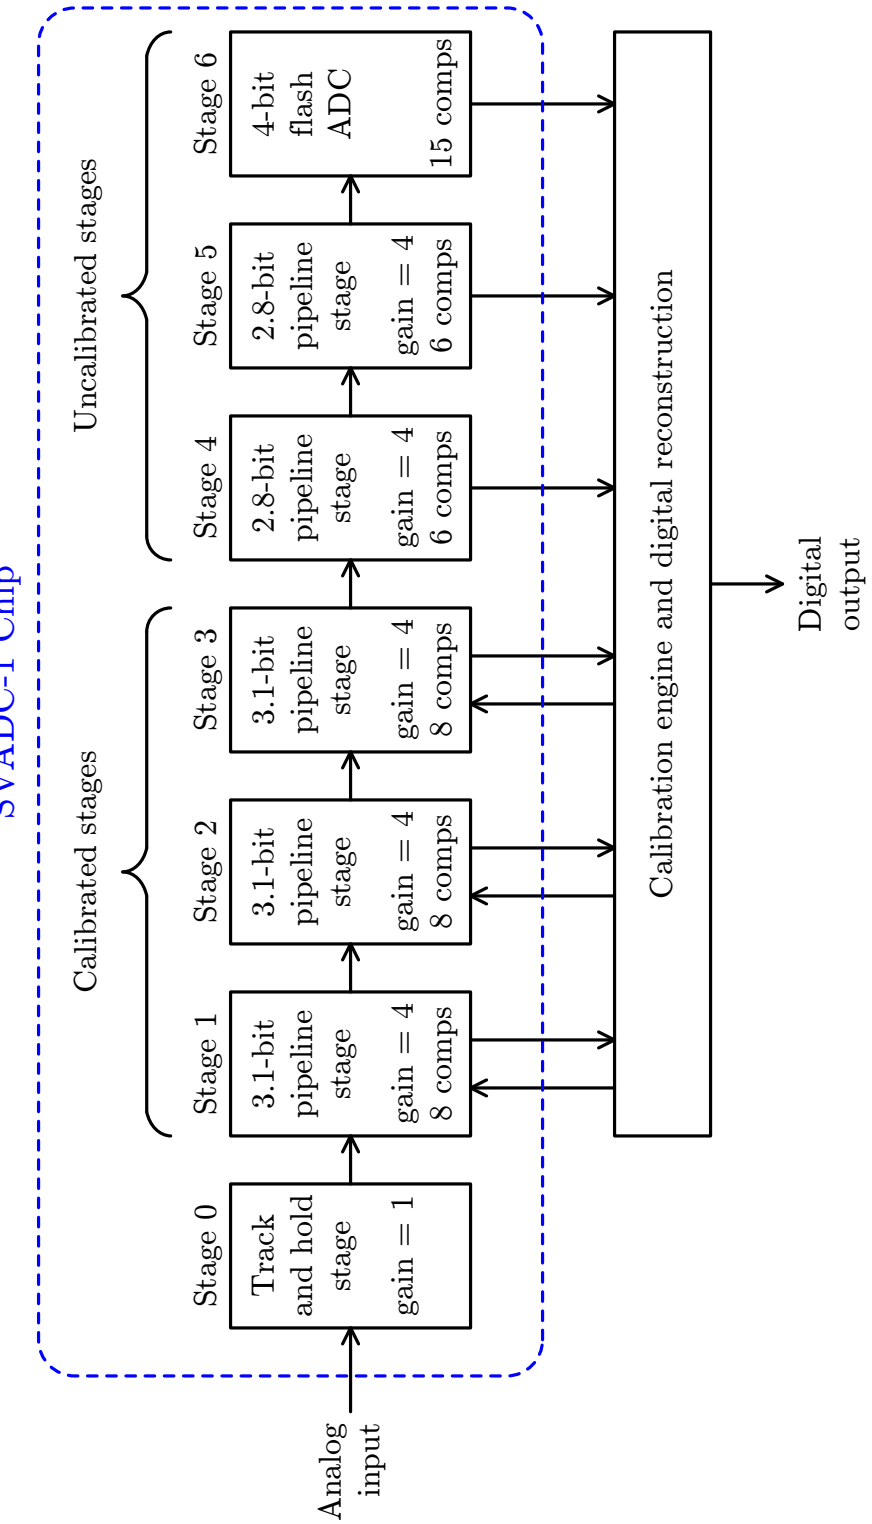

Figure 5.1: Architecture of the SVADC-1 pipeline converter: portions implemented in SVADC-1 chip enclosed by blue dashed box. Gain and number of comparators (comps) given for each stage.

Indeed, given the radiation uncertainty of the process, the circuit design adopts an overall more conservative approach.[2]

5.1 Process Choice

There are two broad strategies for developing radiation tolerant integrated circuits: radiation-hardness by process and radiation-hardness by design.

In radiation-hardness by process, integrated circuits are fabricated in a manufacturing process custom designed to mitigate radiation effects. Traditionally, such processes focus on total-dose hardening of gate and field oxides, paying especial attention to not only the manufacturing of these oxides, but also the impact of stresses from prior and subsequent process steps.[3] The resulting constraints often trade off between radiation tolerance, device performance, chip reliability, and manufacturability (including both ease of production and device yield). Consequently, radiation-hard processes often lag a generation or two behind commercial technologies. Furthermore, owing to their specialized use, they are typically more expensive and more difficult to obtain.

The alternative is to use a commercial manufacturing process and apply design practices to ensure radiation tolerance. Such radiation-hardness by design leverages the (much) reduced cost and (again, much) increased availability of commercial

[2] In practice, lack of radiation characterization—if not of process stability—is a common enough scenario in integrated circuit design that the approach described here, though perhaps not optimal, may nonetheless prove valuable.

[3] Indeed, the entire flow—from the chemicals used to clean the wafer prior to manufacturing, to the growth of passivation over the completed integrated circuit—can affect gate and field oxide radiation tolerance. For example, annealing (a common practice deployed extensively throughout manufacturing to deposit layers or reflow damage) when performed for long durations, at high temperatures, or in the presence of certain ambients, can easily soften already-manufactured gate oxides [*Dressendorfer*, 1989b, pp. 349–354]. Thus, when considering radiation performance the entire flow functions as an integrated whole: many disparate steps can affect the radiation performance of a single feature. Readers further interested are recommended to Chapter 6 of *Dressendorfer* [1989b]—which discusses the hardening of a standard CMOS flow—as an introduction. While addressing older, micron feature size processes, *Dressendorfer* [1989b] nonetheless well illustrates the matrix of interrelations that must be considered. Another, more modern example of the challenges of hardening commercial processes can be found in *Shaneyfelt et al.* [1998], which discusses the hardening of STI oxides.

versus radiation-hard processes. Historically, though, commercial processes are fairly radiation-soft, especially in their gate and field oxides. However, the unrelenting pursuit of scaling has led to ever thinner oxides, such that modern CMOS processes often employ gate oxides less than 10 nm thick. Such oxides are naturally less susceptible to radiation damage (as described in Section 2.3.2), prompting much recent interest in the viability of commercial submicron CMOS processes for radiation-hard applications [*Anelli et al.*, 1999; *Lacoe*, 2003].[4]

The SVADC-1 is fabricated in a commercial SiGe BiCMOS process—dubbed BiCMOS8iED—generously provided by National Semiconductor Corporation, and adopts radiation-hardness by design to ensure radiation tolerance. Although the process includes bipolar transistors, the SVADC-1 uses only the CMOS layers.[5] These layers correspond to a single-well, 1P5M, 0.25-μm CMOS process fabricated in bulk Si[6] on a non-epitaxial substrate. Interdevice separation is implemented by STI oxides.

To establish conventions for the following discussions, Figure 5.2 shows the

[4]A subclass of commercial CMOS processes of much interest for radiation applications are silicon-on-insulator (SOI) technologies (including silicon-on-sapphire). *Schwank et al.* [2003] provide a good introduction to the use of SOI processes for radiation applications. Granted, from a total-dose perspective, SOI devices incur all the radiation degradations of their bulk counterparts and then some: for example, the buried oxide layer can accumulate trapped charge and act as a second radiation-degraded gate oxide, modulating the channel. Indeed, this second gate can sufficiently alter the dominant physics that the worst case irradiation bias for SOI MOSFETs can be substantially different than that standard CMOS devices. However, from a single-event perspective, SOI devices can be more robust than their bulk counterparts, owing to the reduced charge collection of their shallower bulks. They are also immune to the classical PNPN-based latchup of Section 2.2.2.2 (although, without proper body ties, they are susceptible to single-event snapback, a type of single-device latchup). As single-event effects—especially latchup—are often an overriding concern in radiation-hard design, the inherently better single-event performance of SOI makes it a popular commercial choice for many radiation applications (e.g., *Redman-White et al.* [1990]; *Edwards et al.* [1999]; *Irom et al.* [2006]). However, for various reasons—including process availability and access to bipolar transistors (see footnote 5 of this chapter)—an SOI process was not used for the SVADC-1.

[5]The bipolar device layers are instead used in the LNA/AAF ASIC developed in the same process by fellow Ph.D. student Benjamin J. Mossawir [*Mossawir et al.*, 2006; *Mossawir*, in preparation]. Initially, the entire LNA-AAF-ADC signal path of Figure 1.11 was to be implemented by a single ASIC: a process was thus chosen that would best accommodate all three designs. In particular, the LNA—owing to the 100 Hz low frequency cutoff specification—heavily leverages the low flicker noise of the bipolar transistors [*Mossawir et al.*, 2006]. However, as design progressed the receiver was divided into two separate ASICs—the LNA/AAF and the ADC—to decouple their development schedules. The upshot is that, though it does not use the bipolar layers, the SVADC-1 is nonetheless fabricated in a BiCMOS process.

[6]The Ge portion of the process is used only in the bipolar devices.

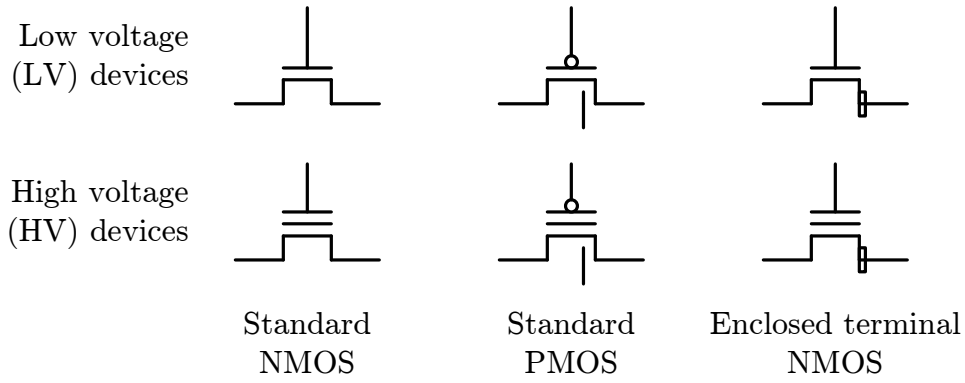

Figure 5.2: MOS transistor symbols. Both LV (top) and HV (bottom) device symbols shown. Symbols for standard NMOS and PMOS devices shown. Also shown are symbols for enclosed terminal terminal NMOS devices: for these symbols, a box is affixed to the enclosed terminal (see Section 5.2.2.2).

symbols used to denote both low-voltage (LV) and high-voltage (HV) NMOS and PMOS devices from the process. HV devices can reliably sustain higher gate-to-channel and junction voltages than their LV counterparts (up to 3.3 V versus 2.5 V), but their thicker gate oxides are more radiation susceptible (see Section 2.3.3). Also shown are symbols for enclosed terminal devices; these devices are addressed in Section 5.2.2.2. Since the process is single well, NMOS bulks are always electrically tied to the substrate (which is in turn tied to ground) and hence NMOS bulk connections are omitted for clarity. PMOS bulk connections, being flexible, are shown explicitly.

5.2 Radiation-Hardness by Design

Being manufactured in a commercial process, the SVADC-1 employs radiation-hardness by design. The following subsections single out some specific radiation-hardness-by-design techniques that warrant especial consideration. It bears repeating, though, that radiation-hardness by design is not limited to these techniques: for the complete view, the reader is referred to Sections 5.3, 5.4, and 5.5 as well.

5.2.1 Preventing Latchup

Since the SVADC-1 is manufactured on a non-epitaxial substrate, latchup is an overriding concern. The reader is referred to Section 2.2.2.2 for a more detailed description of the latchup process. Typically, the parasitic circuits formed from PNPN structures remain in the blocking state, described by (see Equation 2.1):

$$\frac{\alpha_{\mathrm{fs,Q_1}}}{1+\dfrac{R_{\mathrm{e,Q_1}}}{R_{\mathrm{WELL}}}}+\frac{\alpha_{\mathrm{fs,Q_2}}}{1+\dfrac{R_{\mathrm{e,Q_2}}}{R_{\mathrm{SUB}}}}<1 \tag{5.1}$$

where R_{WELL} and R_{SUB} are the well and substrate resistances, respectively. However, ionizing radiation can trigger the circuit, activating the positive feedback and driving the circuit to latchup. This switching point occurs when:

$$\frac{\alpha_{\mathrm{fs,Q_1}}}{1+\dfrac{R_{\mathrm{e,Q_1}}}{R_{\mathrm{WELL}}}}+\frac{\alpha_{\mathrm{fs,Q_2}}}{1+\dfrac{R_{\mathrm{e,Q_2}}}{R_{\mathrm{SUB}}}}=1 \tag{5.2}$$

Unfortunately, non-epitaxial substrates incur relatively high R_{SUB} (especially when compared to epitaxial substrates) and hence are more prone to loop activation and subsequent latchup.[7]

Many systems combat latchup by current limiting the supply, or detecting when latchup has occurred and cycling the supply. In addition to practical concerns such as reaction speed and the system impact of cycling typically shared supplies [*Johnston*, 1996], such strategies are risky as they often compensate an already-activated loop. A more conservative approach is to belay even loop activation: in latchup terms, instead of managing the hold point, the safer course is to prevent the switching point [*Troutman*, 1986, p. 205]. The SVADC-1 adopts this latter strategy. Specifically,

[7] Those desirous of a more detailed investigation of the effect of high substrate resistivity on latchup are referred to *Voldman* [2007, pp. 302–312]. Notably, *Voldman* [2007] provides an explanation that assumes a dual-well process wherein the substrate resistance lies in parallel with a p-well resistance, which is not the case in BiCMOS8iED. Nonetheless, it addresses many of the physical mechanisms and concerns associated with high R_{SUB}.

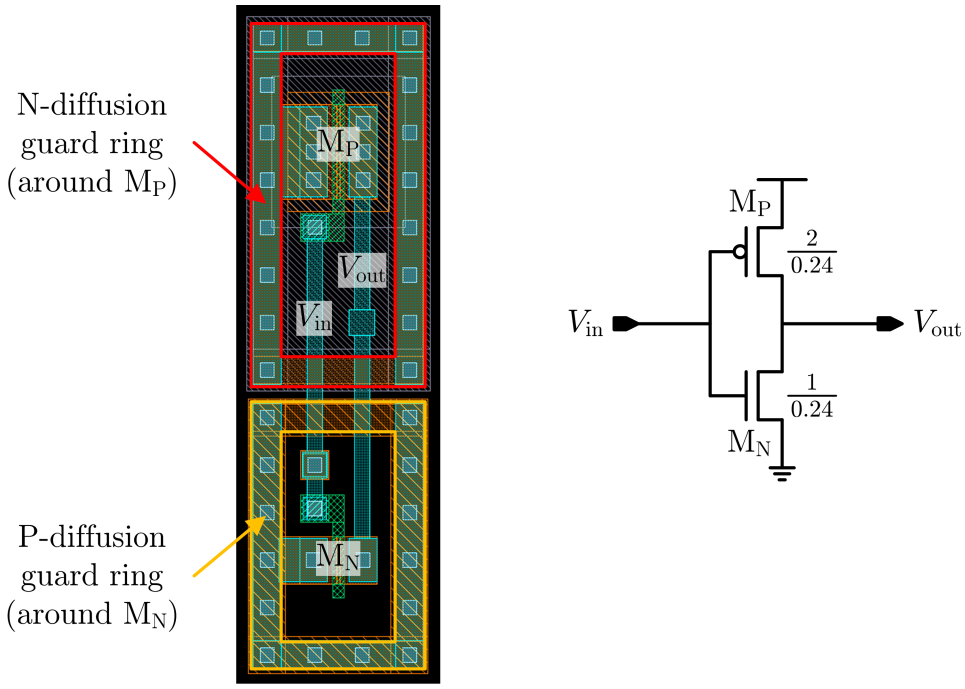

Figure 5.3: Demonstration of guard ring deployment in the layout of an inverter: layout to left, corresponding circuit schematic to right.

it uses extensive deployment of diffusion guard rings close to active devices. These diffusions work to reduce R_{WELL} and R_{SUB} in Equations 5.1 and 5.2, increasing the blocking state space. An example is shown in Figure 5.3 for an inverter. As can be seen, the strategy does incur an area penalty.[8]

Notably, the guard rings are tied to local supplies, including local grounds for the p-diffusions about the NMOS devices. Coupling the substrate with the local return can cause increased substrate noise injection.[9] However, using local supplies helps ensure small R_{SUB} relative to $R_{\text{e,Q}_2}$ (a condition more difficult to guarantee if the guard

[8] An even more robust latchup prevention strategy is to use double guard diffusions of alternating types, for example, adding an n-diffusion ring about the p-diffusion ring surrounding an NMOS device. *Osborn et al.* [1998] compared the latchup robustness afforded by unringed, single-ringed, and double-ringed strategies in a commercial CMOS process. Their findings suggest that double ringing does indeed provide a significant boost to latchup robustness. Practically, though, it also incurs a much more significant area penalty. Hence the SVADC-1 elects the single-ring strategy.

[9] Indeed, since a noisy substrate can derail converters of otherwise high IDR, in many ADC designs the substrate is connected to its own, unique metallization network separate from all other supply returns to reduce local substrate noise injection from those returns.

rings connect to ground through different, long metallization paths), expanding the blocking state space. As latchup is a paramount concern in radiation applications,[10] and as the SVADC-1 is fabricated on a non-epitaxial substrate prone to latchup, in this case stronger latchup prevention outweighs performance loss concerns.[11]

5.2.2 Preventing Leakage

The radiation-induced increase of the drain-to-source leakage current I_{LEAK} of NMOS devices is described in Section 2.3.3.2. This increase can compromise switched-capacitor circuits the leakage current of as nominally disconnected switches drains significant charge from nominally isolated signal nodes. As an example, consider the simple single-ended switched-capacitor amplifier of Figure 5.4. Assuming the switches are implemented by LV 6/0.25 NMOS devices and modeled as variable resistors with off-resistance $R_{\mathrm{off}} = 2.5/I_{\mathrm{LEAK}}$, where I_{LEAK} evolves with dose as in Figure 2.15, Figure 5.5 shows simulation results of the amplifier output at different doses.[12] The droop of V_{out} during $\phi_{2\mathrm{D}}$ at higher doses is clear.

The droop stems primarily from leakage in the feedback switch $\mathrm{S_F}$ draining the charge stored at V_{X} during $\phi_{2\mathrm{D}}$. Analytically, V_{out} during $\phi_{2\mathrm{D}}$ can be well-approximated by a growth exponential that describes the ideal (i.e., sans leakage) amplifier output and a decay exponential that describes the leakage-induced droop. Assuming the operational amplifier model of Figure 5.4, the time constant of the former is approximately [*Razavi*, 2001, pp. 436–437]:

$$\tau_{\mathrm{growth}} = \frac{(1 + g_{\mathrm{m}} R_{\mathrm{out}}) C_{\mathrm{F}} + C_{\mathrm{S}} + C_{\mathrm{P}}}{R_{\mathrm{out}} \left[(C_{\mathrm{S}} + C_{\mathrm{P}}) C_{\mathrm{out}} + C_{\mathrm{F}} (C_{\mathrm{S}} + C_{\mathrm{P}} + C_{\mathrm{out}}) \right]} \tag{5.3}$$

[10] Indeed, in most part selection, latchup is the first gate: if a part is susceptible to latchup, then it does not fly, irrespective of other compelling performance properties.

[11] BiCMOS processes often include additional structures—such as the subcollector and deep trench structures—that can be used to further enhance latchup robustness. A review of these practices can be found in *Voldman* [2007, Chap. 6]. In the SVADC-1, though, as only the CMOS layers are used, these strategies are not adopted.

[12] This R_{off} model is admittedly very rough: for example, since I_{LEAK} is measured with fixed bias voltages, this model ignores R_{off} variations with channel bias. Nonetheless, it is sufficiently illustrative for this discussion.

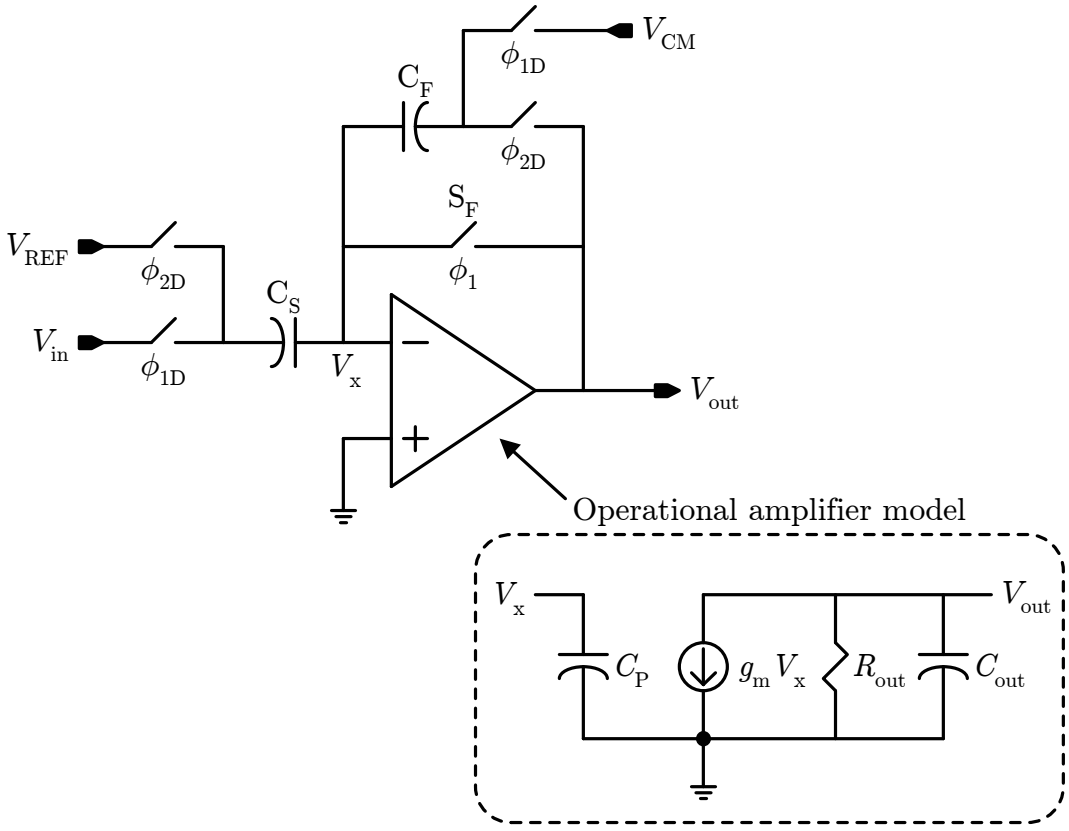

Figure 5.4: Switched-capacitor amplifier for demonstrating I_{LEAK} impact. Operational amplifier model shown in insert. See Figure 5.15 for clock timing diagram.

whereas the time constant for the latter can be approximated as:

$$\tau_{decay} = \frac{1}{R_{off,S_F} C_F} \tag{5.4}$$

Importantly, the decay exponential is always present, even without radiation. Nominally, though, R_{off,S_F} is large and τ_{decay} sufficiently slow compared to the amplifier clock rate that the droop is negligible. However, radiation-induced I_{LEAK} increase—which translates to R_{off,S_F} decrease—hastens the droop. Perhaps notably, the droop can also be exacerbated by decreasing the amplifier clock rate which would grant the decay exponential more time to develop. Consequently, augmented I_{LEAK} may cause down-clocked switched-capacitor amplifiers to perform worse under radiation.

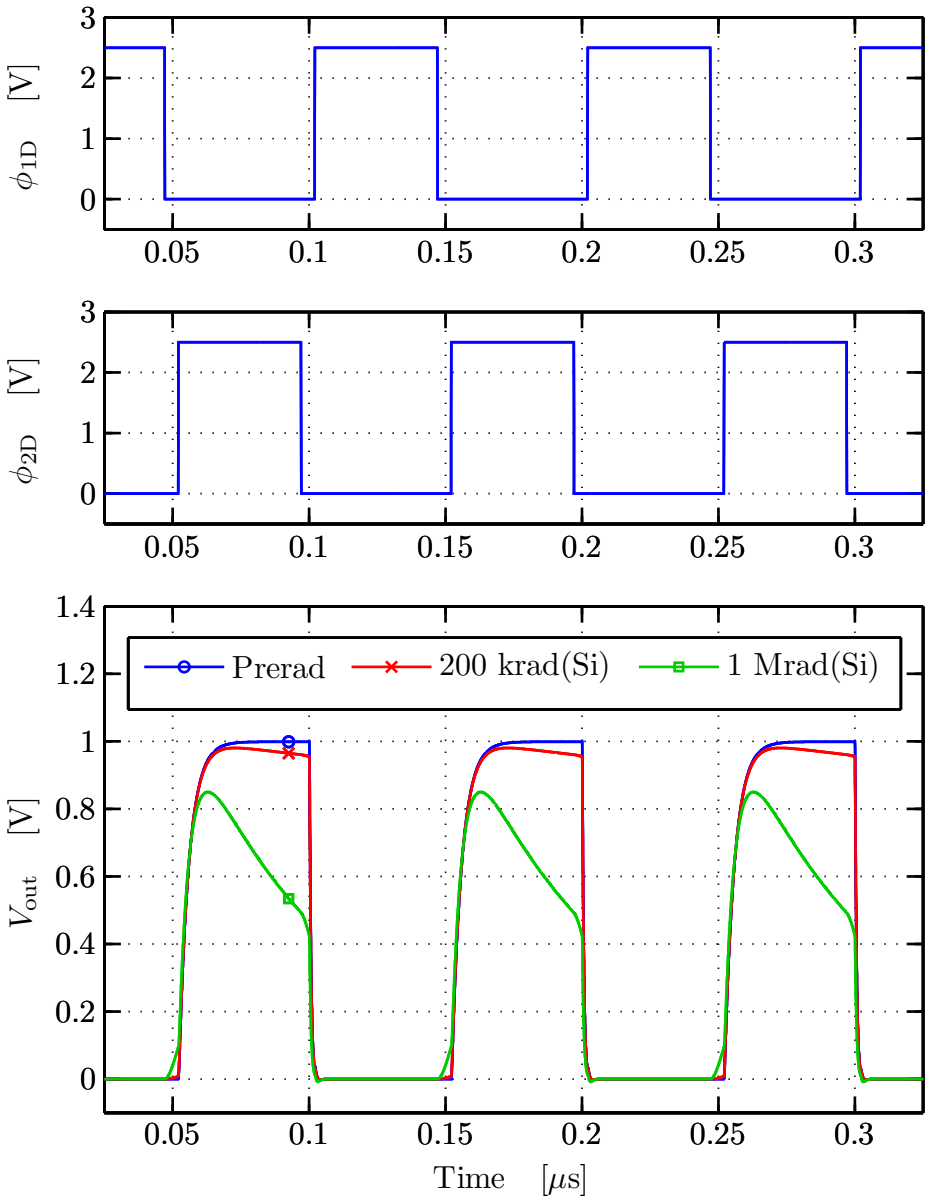

Figure 5.5: Impact of I_{LEAK} on the V_{out} of the switched-capacitor amplifier of Figure 5.4 under various radiation doses. ϕ_{1D} and ϕ_{2D} waveforms included for reference: amplifier clocked at 10 MHz. Amplifier conditions: $C_S = 1$ pF and $C_F = 0.25$ pF. Input signal conditions: $V_{in} = 0.25$ V and $V_{REF} = V_{CM} = 0$ V. Operational amplifier conditions: $g_m = 4.6$ mS, $R_{out} = 1.08$ MΩ, $C_{out} = 2.61$ pF, and $C_P = 170$ fF. Switch conditions: LV 6/0.25 switch assumed, $R_{on} = 333$ Ω (average of switch R_{on} over 0.75–1.75 V signal range) and $R_{off} = 2.5/I_{LEAK}$, for I_{LEAK} as per Figure 2.15. Other switch sizes yield similar results.

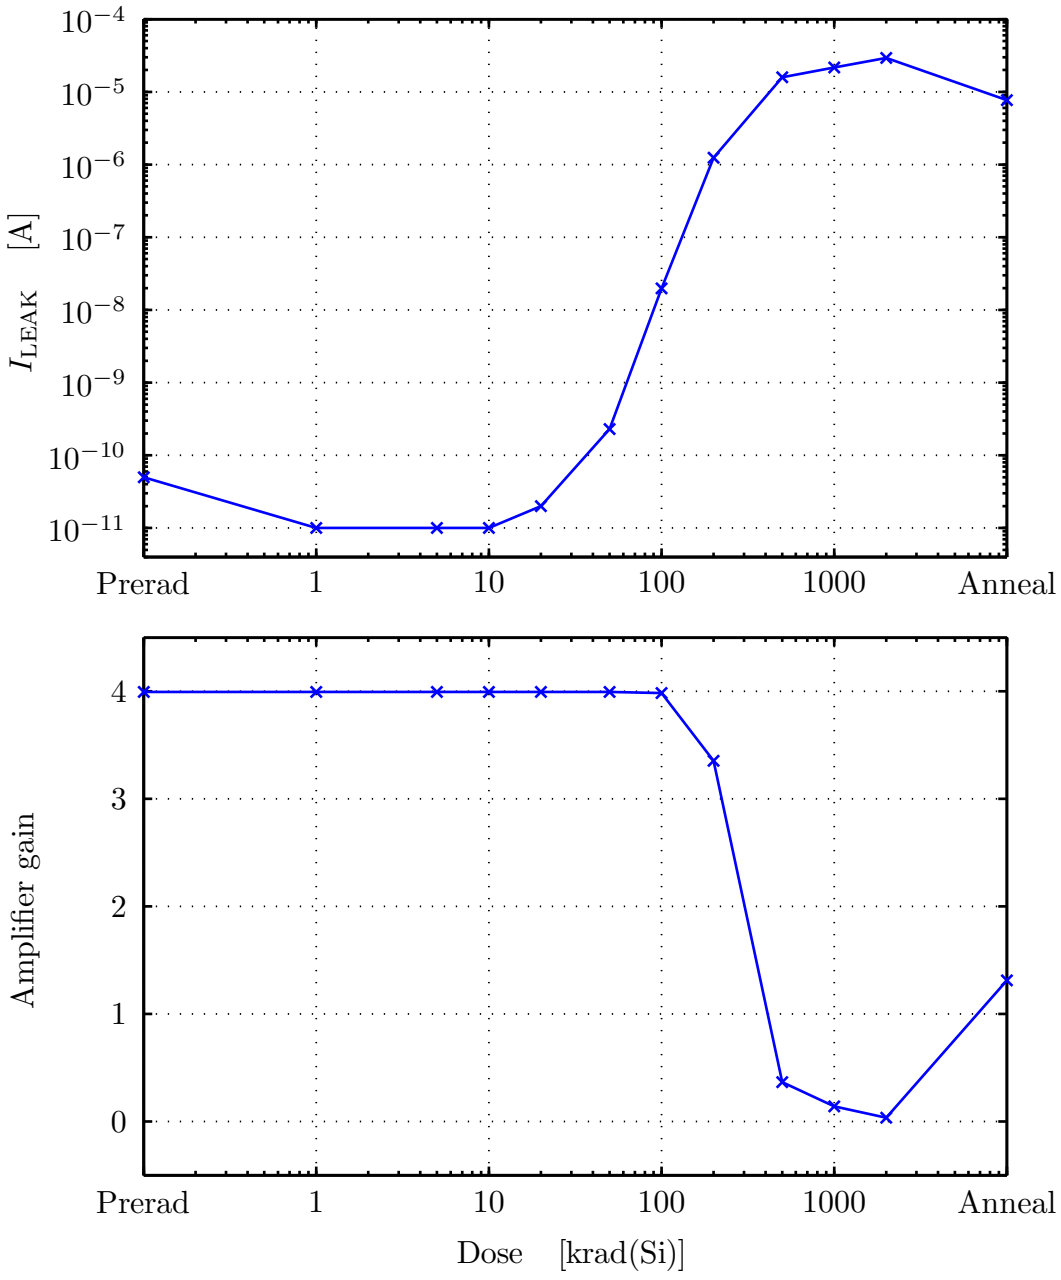

Figure 5.6: Impact of I_{LEAK} on the gain of the switched-capacitor amplifier of Figure 5.4 over dose. Radiation evolution of gain shown at bottom, of I_{LEAK} included at top for reference. Circuit conditions are the same as in Figure 5.5. Again, other switch sizes yield similar results. Markers indicate measured data points.

The first order effect of droop is loss of amplifier gain, as shown in Figure 5.6: given the measured I_{LEAK} dose profile, the nominal amplifier gain of 4 drops to subunity by 500 krad(Si). In an uncorrected pipeline ADC, this gain loss alters the stage transfer function discontinuity heights H_m (see Equation 4.28), increasing analog-digital mismatch and compromising linearity. In the SVADC-1, though, calibration corrects such H_m changes. Rather, such gain loss augments Q_{pipeline} (see Equation 4.15): instead of quantization error nonuniformity, the increase in quantizer Δ itself limits the SFDR.

5.2.2.1 Nonstandard Transistor Layout

To reduce radiation-induced I_{LEAK} increase, the SVADC-1 employs a nonstandard transistor layout for select NMOS devices. Two such layouts—namely the annular layout and enclosed terminal layout—are shown in Figure 5.7. The standard layout is also included for reference. As discussed in Section 2.3.3.2, in the standard layout the vast majority of I_{LEAK} flows through a leakage path underneath the radiation-soft thick field oxide adjacent to the nominal transistor. In contrast, owing to its smaller thickness, very little I_{LEAK} flows underneath the more radiation-hard thin gate oxide. The nonstandard layouts of Figure 5.7, then, mitigate I_{LEAK} by fully enclosing a source or drain terminal in thin gate oxide. Insofar as any current in or out of the enclosed terminal must thus pass underneath thin gate oxide, and insofar as the thin gate oxide is radiation-hard compared to the thick field oxide, I_{LEAK} is reduced compared to the standard layout.

Many recent designs have favored the annular layout (e.g., *Snoeys et al.* [2000]; *Rivetti et al.* [2001]; *Anelli* [2000]). However, the SVADC-1 instead adopts the enclosed terminal layout. Two concerns motivate this choice: asymmetry and modeling. Regarding the first, in general the enclosed terminal layout is less asymmetric than the annular, especially for the wider devices expected in switched-capacitor circuit switches. For example, with the channel connected, annular device output conductances are known to vary by 20–75% depending on whether the drain (i.e., the higher potential terminal) corresponds to the inside or outside diffusion [*Anelli*, 2000, p. 109]. While proper design can accommodate such signal dependent

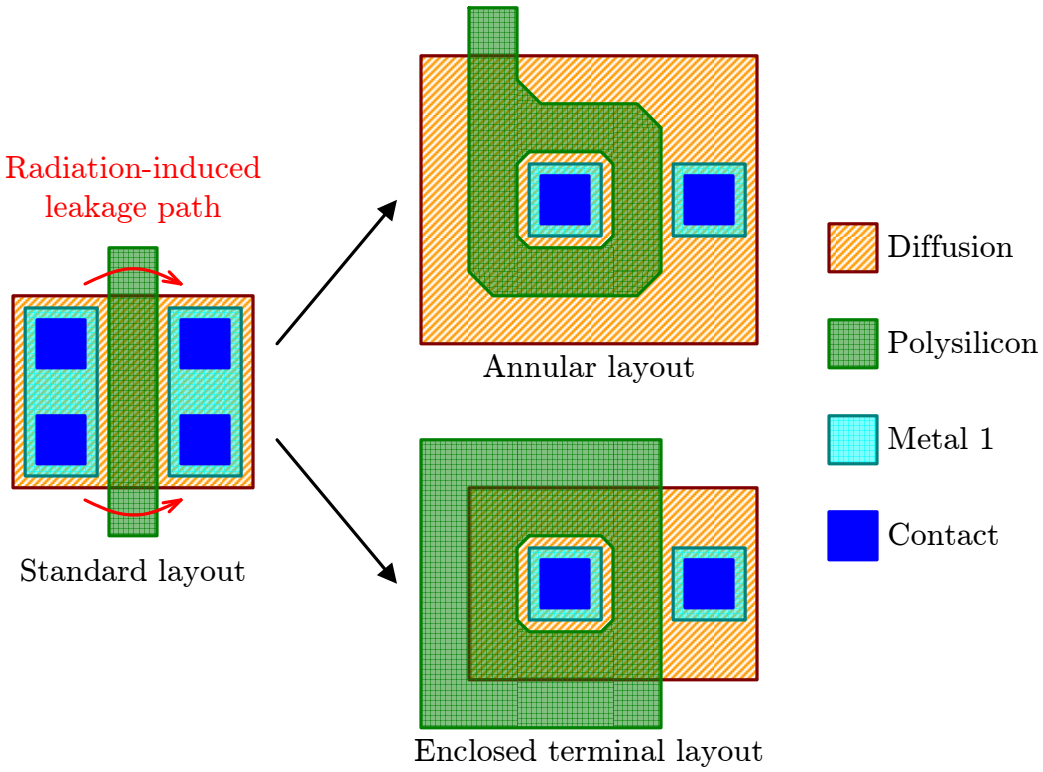

Figure 5.7: Nonstandard NMOS transistor layouts for mitigating I_{LEAK}. Standard device—with parasitic leakage paths highlighted in red—shown at left for reference.

variation, it is generally undesirable.[13] Regarding the second, as the enclosed terminal layout cleaves closer to the standard, its model ought also cleave closer to the standard device model. As perhaps to be expected, manufacturers do not generally provide models for nonstandard devices and designers must recourse to custom models built upon manufacturer-provided models. It is thus generally advisable to select nonstandard devices that best match the standard. In particular, note that the channel current of the enclosed terminal and standard devices should be fairly similar, especially for wider devices: over the center of the device width—which handles the majority of the channel current—the current flow should be the largely the same for

[13]For the enclosed terminal layout, such resistance variation, while existent, should be much less, especially for wider devices.

both devices. This similarity again recommends the enclosed terminal layout.[14,15]

5.2.2.2 Enclosed Terminal Transistor Layout

The enclosed terminal layout used in the SVADC-1 is detailed in Figure 5.8. The thin gate oxide of the enclosure is made at least the minimum channel length—0.25 μm for LV devices, and 0.4 μm for HV devices—to ensure its manufacturability. By the same token, the polysilicon overhang associated with the enclosure is at least the minimum design rule dimension. Inside the enclosure itself, the polysilicon of the inner corners is chamfered to 45°: this angle not only enhances manufacturability, but also improves reliability by guarding against the high gate oxide electric fields associated with a sharp 90° corner. Finally, this layout technique is only applied on NMOS devices since (as explained in Section 2.3.3.2) for PMOS devices I_{LEAK} only decreases with dose.

The efficacy of the enclosed terminal layout for preventing I_{LEAK} is shown in Figure 5.9, which shows measured radiation testing data of test NMOS devices from the BiCMOS8iED process (as described in Appendix F). In particular, Figure 5.9 shows I_{LEAK} over dose for the same device drawn in either standard (left) or enclosed terminal (right) fashion for a variety of device sizings. While the standard device shows a clear increase in I_{LEAK} by 6 orders of magnitude, the enclosed terminal device shows no manifest change in I_{LEAK} up to the full 2 Mrad(Si) tested. In fact, as its channel current is now predominantly controlled by radiation-hard thin gate

[14]To be sure, it is possible to derive equivalent W/L ratios [*Giraldo et al.*, 2000; *Grignoux and Geiger*, 1982] and matching statistics [*Anelli*, 2000, Chap. 5] for the annular layout. (Indeed, interested readers are referred to Chapters 4 and 5 of *Anelli* [2000] for a good summary of annular device modeling.) These derived models, though, often include fitting parameters that cannot be known *a priori* and require experimental measurement.

[15]It is also possible to derive simple models for the enclosed terminal layout device. For example, *Nowlin et al.* [2005] advocated modeling the device as a channel MOSFET (representing the nominal transistor and drain-to-source current flow under the corners of the enclosure) in parallel with a parasitic MOSFET. The parasitic MOSFET is composed of two FETs in series that follow the parasitic current flow under the gate ring and then along the edge of the field oxide (the latter segment being depicted by the channels of the orange FOXFET devices in Figure 2.13(b)): the parasitic device is thus modeled as a ring-MOSFET followed by an edge-FOXFET. *Nowlin et al.* [2005] used this model to explain increases in maximum transconductance with dose for the enclosed terminal devices (see footnote 44 of Chapter 2). However, they note that more precise calibration is needed for accurate prediction of device behavior under radiation.

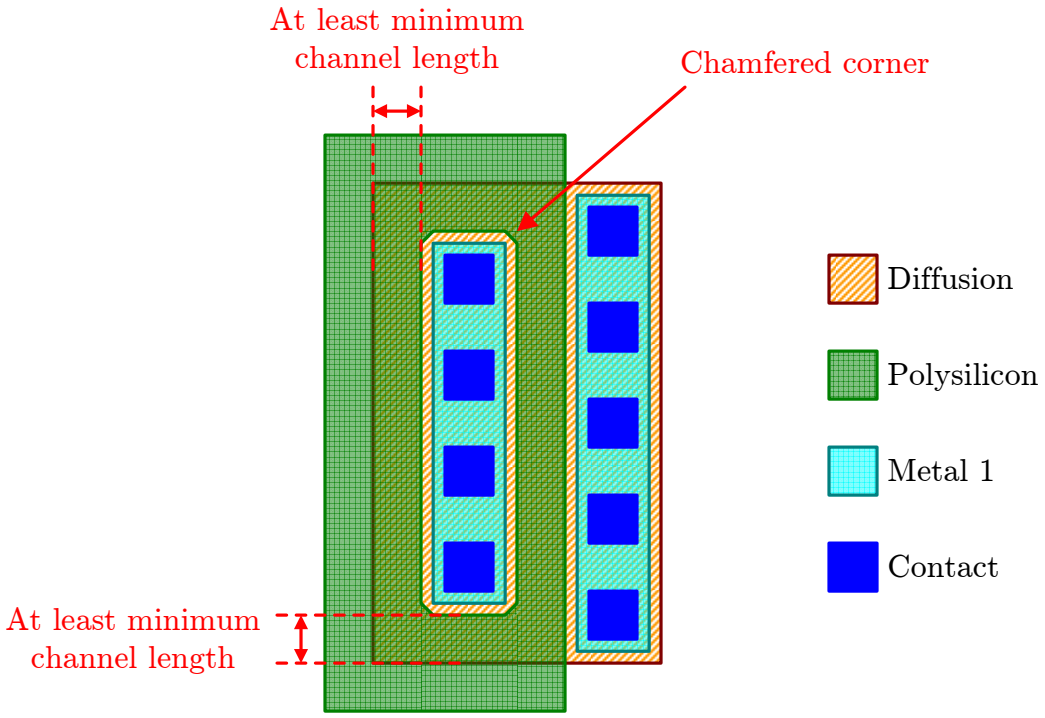

Figure 5.8: Detail of enclosed terminal layout for an NMOS device.

oxide, it is expected that the enclosed terminal device should generally show little change in channel current, I_{LEAK} or otherwise, with dose. This effect is confirmed in Figure 5.10, which compares the radiation evolution of the I_{D}-V_{GS} curves for the LV 6/0.25 device of Figure 5.9; the I_{D}-V_{GS} curves of the other devices are similar. The alignment of the standard device characteristics to the left, and enclosed terminal device characteristics to the right, is preserved. On the whole, the enclosed terminal device shows little variation in its I_{D}-V_{GS} curve with dose. In contrast, the standard device shows dramatic subthreshold I_{D} increase with dose. These results not only demonstrate the efficacy of enclosed terminal layout for mitigating radiation-induced I_{LEAK}, but also confirm that the majority of I_{LEAK} indeed flows under the thick field oxide and that the thin gate oxide of this process is indeed radiation-hard.

Figure 5.10 also shows that the standard and enclosed terminal devices display similar post-threshold behavior. However, while it is similar to the standard device in this regime, the enclosed terminal device is nonetheless still different and, given

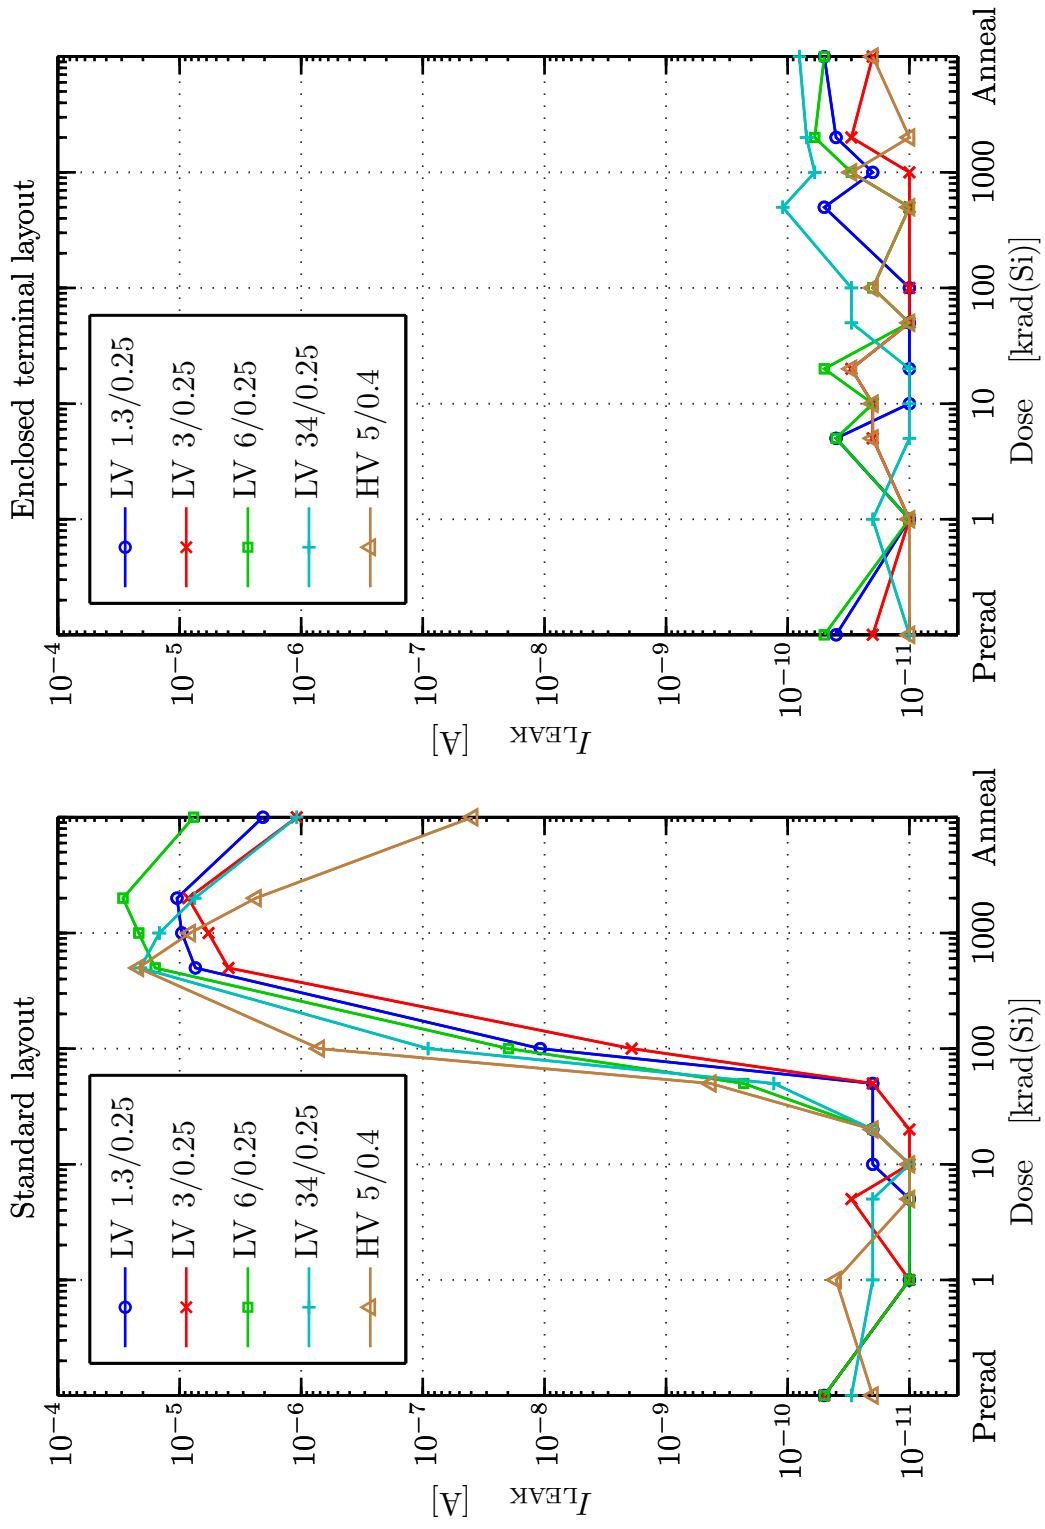

Figure 5.9: Measured I_{LEAK} versus total dose for devices drawn in either standard (left) or enclosed terminal (right) fashion for multiple devices. Device type and W/L given in legend. I_{LEAK} measured as I_{D} when $V_{\mathrm{D}} = 2.5$ V and $V_{\mathrm{G}} = V_{\mathrm{S}} = V_{\mathrm{B}} = 0$ V. All devices irradiated under worst-case bias. Markers indicate measured data points.

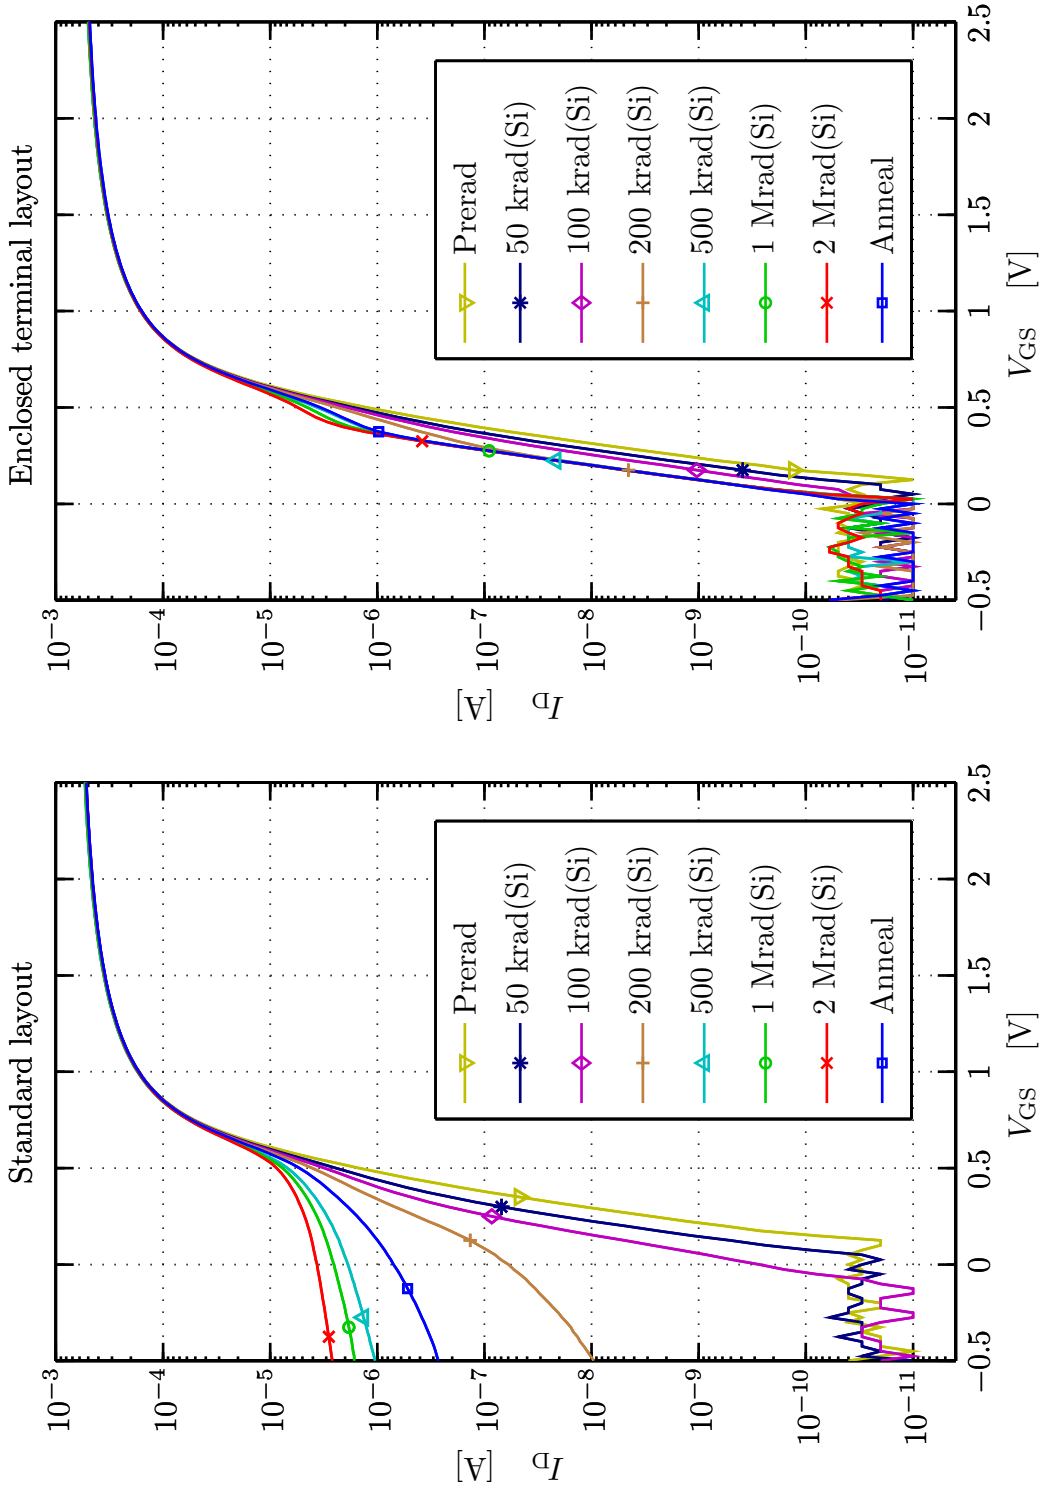

Figure 5.10: Measured radiation evolution of I_D-V_{GS} curves for an LV 6/0.25 device drawn in either standard (left) or enclosed terminal (right) fashion. Select total doses shown. Other devices yield similar results. Measured with $V_S = V_B = 0$ V and $V_D = 0.1$ V; V_G swept in 25 mV increments. Devices irradiated under worst-case bias.

a lack of rigorous modeling, equivalences between the two can only be drawn so far. Hence the enclosed terminal layout is only deployed for select NMOS devices (such as the switches of switched-capacitor circuits) whereas the standard layout is maintained where sensitive analog models are necessary for design (such as in operational amplifiers and dynamic latches). To differentiate between the two layouts in schematics, an enclosed terminal device is indicated by the standard NMOS symbol with a box affixed to the enclosed terminal as depicted in Figure 5.2.

5.3 Common Practices

The next few sections describe the implementation of the SVADC-1 on a circuit and transistor level. While a strictly top-down approach would be conceptually neat here, in practice the circuit descriptions work best when prefaced with some introductory comments that underpin the subsequent discussions. The following subsections lay this foundation.

5.3.1 Supplies

The SVADC-1 uses five separate supplies: an analog supply $V_{\mathrm{DD,A}}$ (for analog circuitry), a digital supply $V_{\mathrm{DD,D}}$ (for digital logic), a clock supply $V_{\mathrm{DD,CLK}}$ (for circuits—buffers, switches, and so on—handling clock signals), a digital I/O supply $V_{\mathrm{DD,IO}}$ (see Section 5.5.5), and a reference supply $V_{\mathrm{DD,REF}}$ (see Section 5.4.3). Each supply is coupled with its own return. Ostensibly the returns should be separate on chip and only tied together off chip. However, as per the local supply philosophy for latchup robustness of Section 5.2.1, all returns do indeed connect on chip at the substrate. Nonetheless, separate metallization trees are maintained for each return.

5.3.2 Switches

Switched-capacitor circuits are used extensively throughout the SVADC-1. To accommodate the full signal range, the switches of these circuits are implemented by CMOS transmission gates as shown in Figure 5.11. This switch design prevents I_{LEAK}

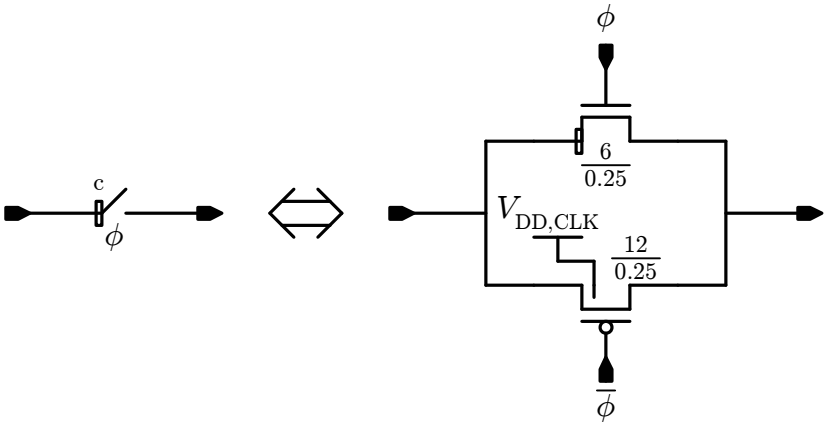

Figure 5.11: CMOS switch symbol (left) and corresponding circuit (right). In the switch symbol, the enclosed terminal of the NMOS device is indicated by a box affixed to the appropriate switch terminal. The switch sizing is indicated by the letter above the switch (see Table 5.1 for corresponding sizing).

by drawing the NMOS device in enclosed terminal fashion. The enclosed terminal is indicated in the switch symbol by a box affixed to the appropriate switch terminal. Different switch sizings are used throughout the SVADC-1 as summarized in Table 5.1. The decidedly discrete set of switch sizings is chosen for practical reasons: since the enclosed terminal devices are not properly extracted by the layout-versus-schematic (LVS) tool of the process, the enclosed terminal devices must be substituted by hand and checked by eye and design reuse helps reduce errors during this process.

Switch type	NMOS size	PMOS size
a	1.3/0.25	2.6/0.25
b	3/0.25	6/0.25
c	6/0.25	12/0.25
d	28/0.25	56/0.25
e	34/0.25	68/0.25

Table 5.1: CMOS switch sizings. All devices LV, all device sizings in μm.

All switches are composed of LV devices. For simplicity only the NMOS gate signal ϕ is given when the switch symbol is used (this convention is consistent with the convention of Section 4.2.3.1).[16]

5.3.3 Capacitors

The capacitors of the switched-capacitor circuitry are implemented as interdigital metal capacitors, shown in Figure 5.12 in both overhead and cross-section views. The circuit symbol is also shown. Metals 2 through 5 implement the capacitor proper and metal 1 functions as a bottom plate shield.

Note that the interdigital capacitor uses both lateral and vertical capacitance. While this structure boosts density, reliance on vertical capacitance is typically not advised since this dimension is not well-controlled during manufacturing (in contrast, the lateral dimension—being determined by photolithography—is) [*Aparicio and Hajimiri*, 2002]. However, the manufacturing process only provides matching data for interdigital capacitors employing both fields, consequently, both are used here. In addition, the process provides measurement data for only a single interdigital capacitor size. Hence a custom, scalable capacitance model was developed for this structure that computes capacitance as a function of length and number of fingers.[17,18]

[16] As a detail, note that the PMOS bulk is always tied to $V_{\mathrm{DD,CLK}}$: the PMOS R_{on} reduction afforded by bulk connection switching [*Limotyrakis et al.*, 2005] is not necessary in the SVADC-1.

[17] The model multiplies the overhead area of the capacitor by the capacitance density per unit area, the area being a simple function of the number of fingers and finger length. The capacitance density is estimated by two-dimensional Poisson field solver results of the cross-section view of Figure 5.12: the capacitance is divided by the width of the structure to give the density. Linear corrections, based on two-dimensional Poisson field solver results of the overhead view of Figure 5.12, are then added to compensate for the capacitance at the ends of the fingers and for the edge capacitance of the outermost fingers. To complete the model, a parasitic capacitance to substrate is included, based on the parallel-plate capacitance between the metal 1 ground shield and the substrate. The model has proven fairly accurate—albeit only indirectly confirmed—in that it has well predicted not only the circuit noise of the SVADC-1, but also the cutoff frequency of a 6th-order, Chebyshev low-pass filter also designed in this process [*Mossawir*, in preparation].

[18] In layout, interdigital capacitor arrays are always surrounded by dummy metal on metals 1 through 5 to reduce edge effects on the outer capacitors during fabrication, enhancing both capacitor accuracy and matching.

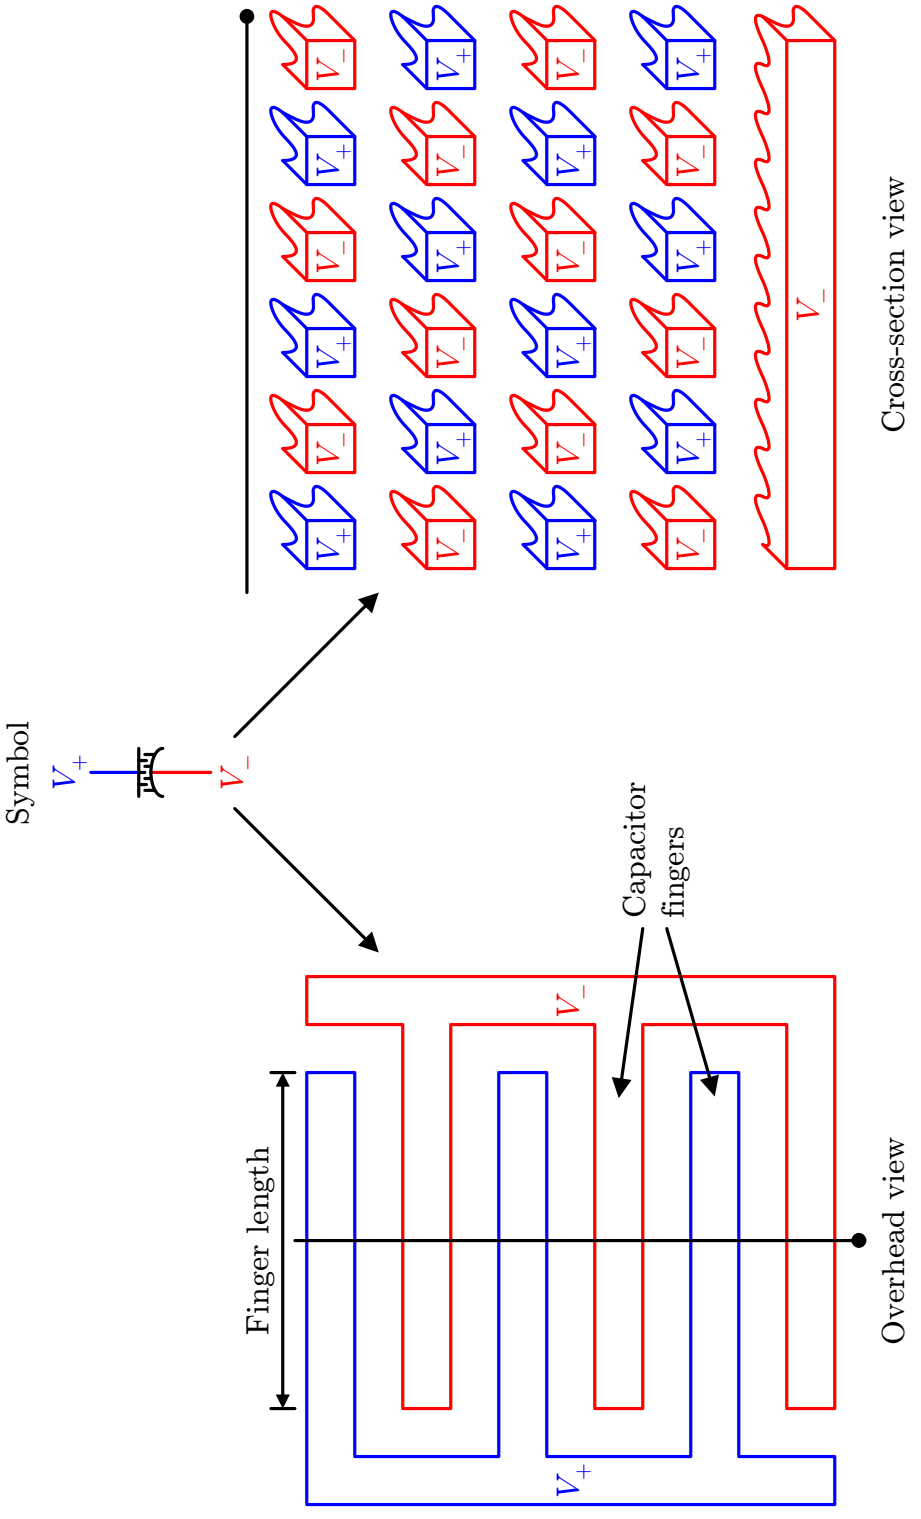

Figure 5.12: Depiction of interdigital capacitor used in the SVADC-1. Circuit symbol shown at center: overhead and cross-section views shown at left and right, respectively. Bottom terminal highlighted in red, top in blue. Drawings not necessarily to scale.

5.3.4 Clocking

All on-chip clocking is ultimately sourced by an off-chip clock signal ϕ_{CLK}. The system-level distribution of ϕ_{CLK} is shown in Figure 5.13. Through a series of buffers, a single in-phase clock ϕ_{STG} is distributed to each stage. In addition, a similar in-phase clock ϕ_{IO} is distributed to the I/O circuitry (see Section 5.5.5).[19]

Each stage locally derives all necessary phases from ϕ_{STG}. In particular, the non-overlapping clock phases are derived from ϕ_{STG} by the circuitry of Figure 5.14. As shown in Figure 5.14(a), the single-ended ϕ_{STG} is converted to differential by parallel paths of 2 and 3 inverters sized for equal delay. As shown in Figure 5.14(b), the resulting differential ϕ_{LOC} and $\overline{\phi_{\text{LOC}}}$ drive a cross-coupled NAND circuit to generate the non-overlapping phases [*Gregorian and Temes*, 1986, pp. 516–517]. The resulting clocks are depicted in Figure 5.15; the delay around the loop of Figure 5.14(b), fostered by the in-loop inverters, sets the non-overlap duration. Generating the non-overlapping clocks locally obviates many dephasing issues associated with distributing such phases over long distances. Still, the final ϕ_1, $\phi_{1\text{D}}$, ϕ_2, and $\phi_{2\text{D}}$ are locally distributed as a ground-shielded bus within the stage and care taken to equally load, tap, and buffer each line. Any additional needed clock phases (such as pulsed clocks for the comparator) are derived from the tapped-and-buffered $\phi_1/\phi_{1\text{D}}/\phi_2/\phi_{2\text{D}}$ bus.[20]

As depicted in Figure 5.14, ϕ_1 and $\phi_{1\text{D}}$ are in-phase with ϕ_{STG}, and ϕ_2 and $\phi_{2\text{D}}$ anti-phase. This configuration is used in stages 0, 2, 4, and 6. Since the ϕ_2 of stage p should correspond to the ϕ_1 of stage $p+1$, the ϕ_1 ($\phi_{1\text{D}}$) and ϕ_2 ($\phi_{2\text{D}}$) clocks of adjacent stages should be delayed by half a ϕ_{STG} cycle, for stages 1, 3, and 5, in Figure 5.14(b), $\phi_1/\phi_{1\text{D}}$ are instead taken to be $\phi_2/\phi_{2\text{D}}$, and vice versa. However, to simplify the discussions of this chapter, by convention ϕ_1 and $\phi_{1\text{D}}$ always refer to

[19] As a detail on the layout of the clock tree: in general, clock lines are distributed by ground-shielded lines on upper metals and care taken to minimize lines crossing above or below. In addition to reducing clock signal coupling, this practice establishes small but well-defined wire capacitance. The clock tree is then sized to support a fanout between 3 and 4 while incorporating these often dominating wire parasitics. Finally, given the asymmetric organization of the stages (see Figure 5.40), the clock lines of Figure 5.13 are sometimes "snaked" to maintain equal wire capacitance on branching lines.

[20] The one exception to this practice is the clock driving the sampling switch of the track-and-hold stage which uses the early clock phase ϕ_{MAS} as described in Section 5.4.1.1.

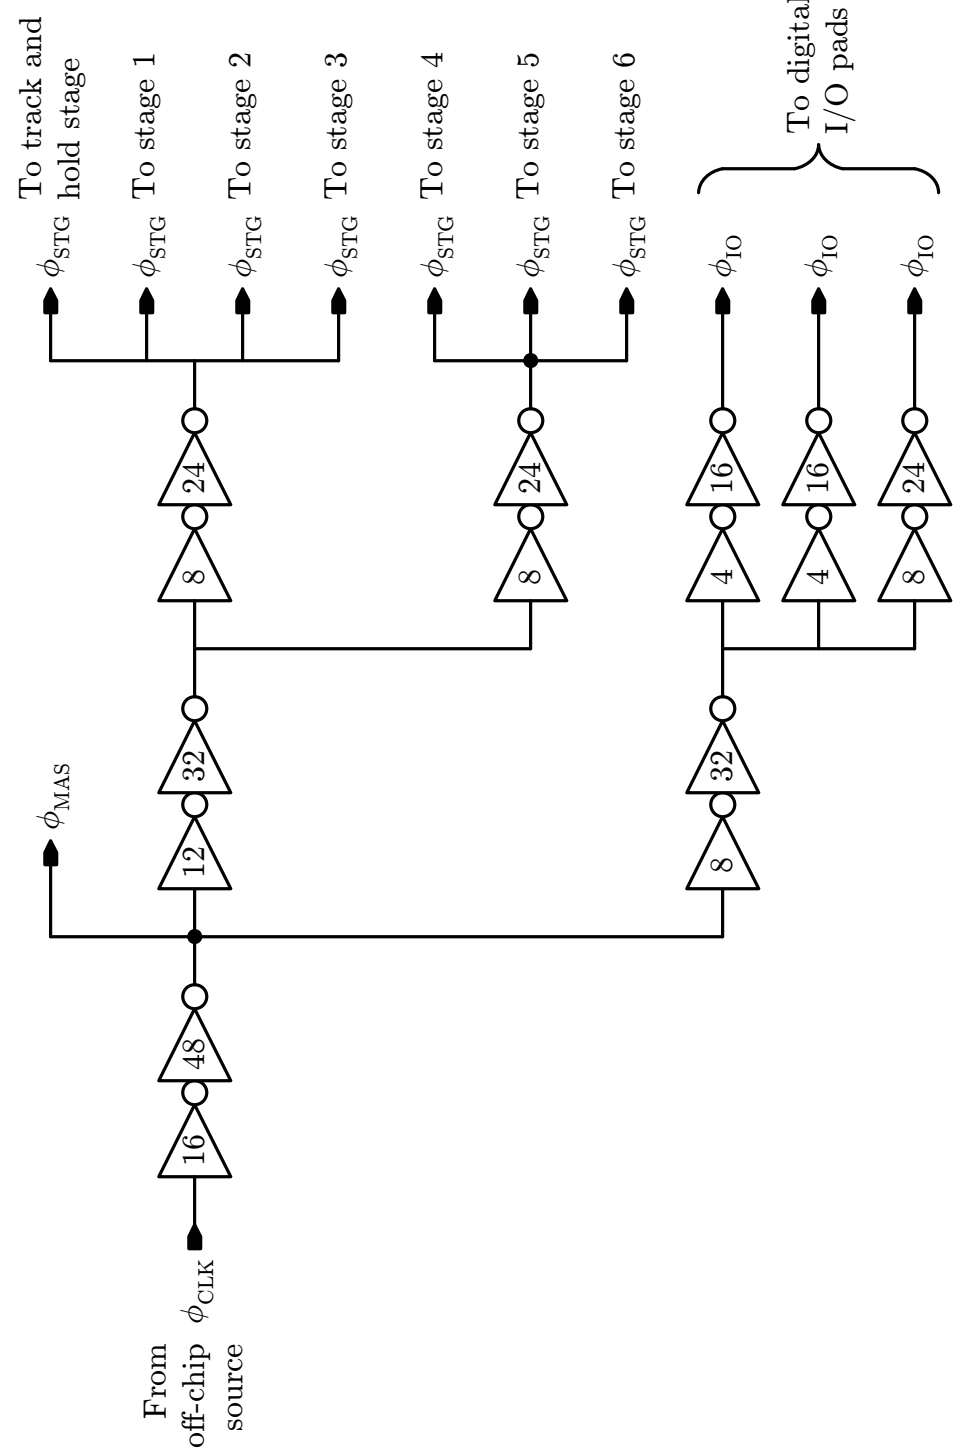

Figure 5.13: Clock tree for distributing off-chip clock ϕ_{CLK} to stages (ϕ_{STG} and ϕ_{MAS}) and I/O circuitry (ϕ_{IO}).

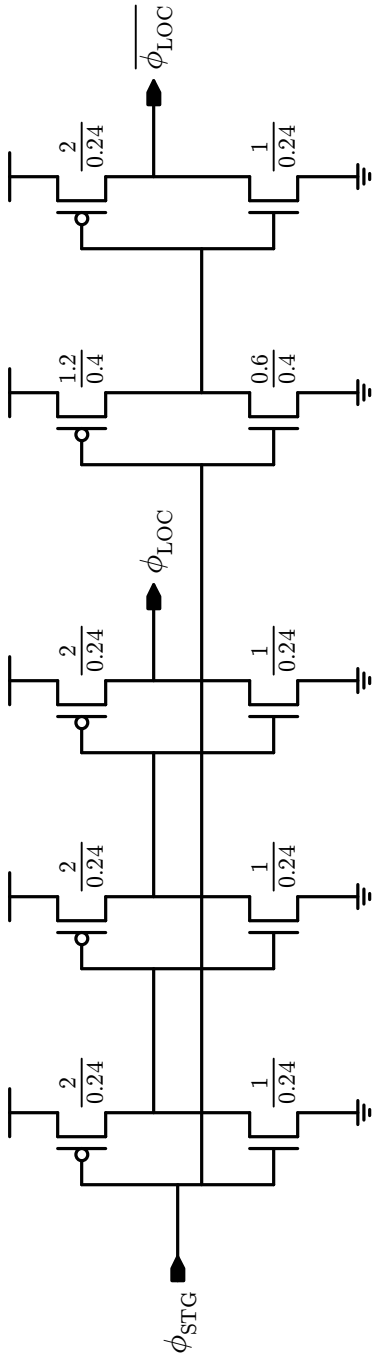

(a) Single-ended to differential converter.

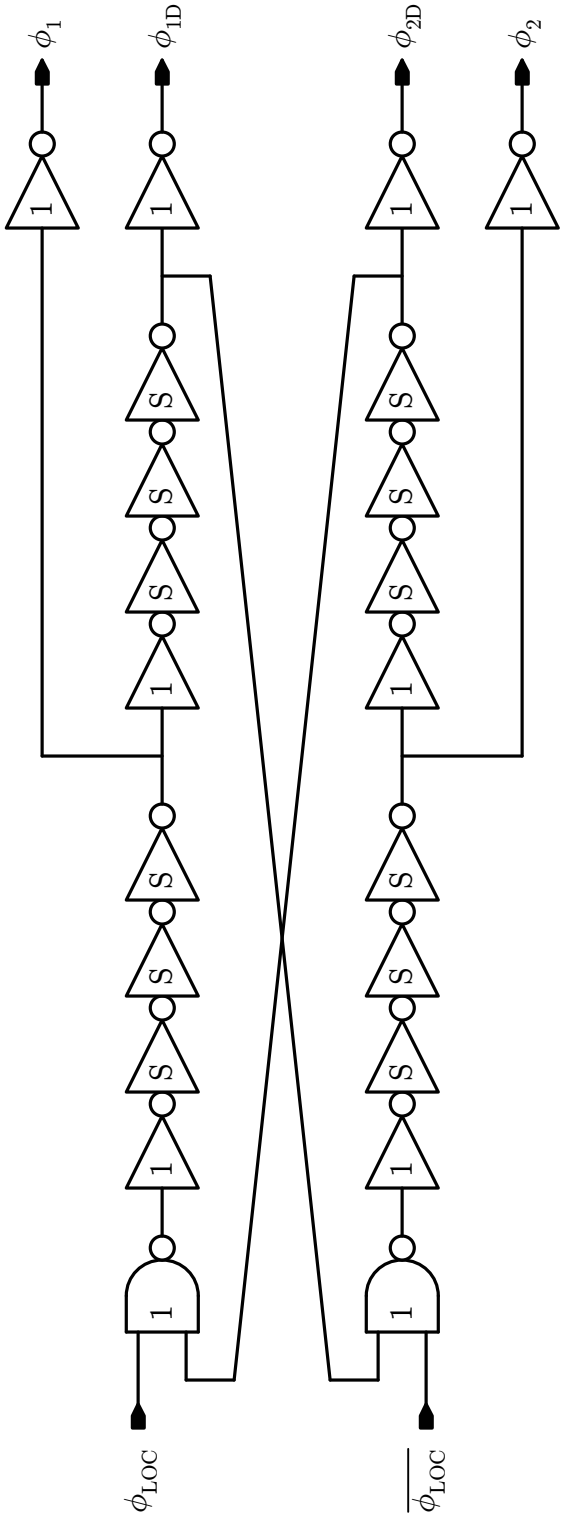

(b) Non-overlapping clock generator.

Figure 5.14: Non-overlapping clock generation circuit used in each stage. "S" inverter refers to slowed inverter (i.e., a 1X inverter with output loaded by a large gate capacitance). All device sizings in μm. Resulting waveforms shown in Figure 5.15.

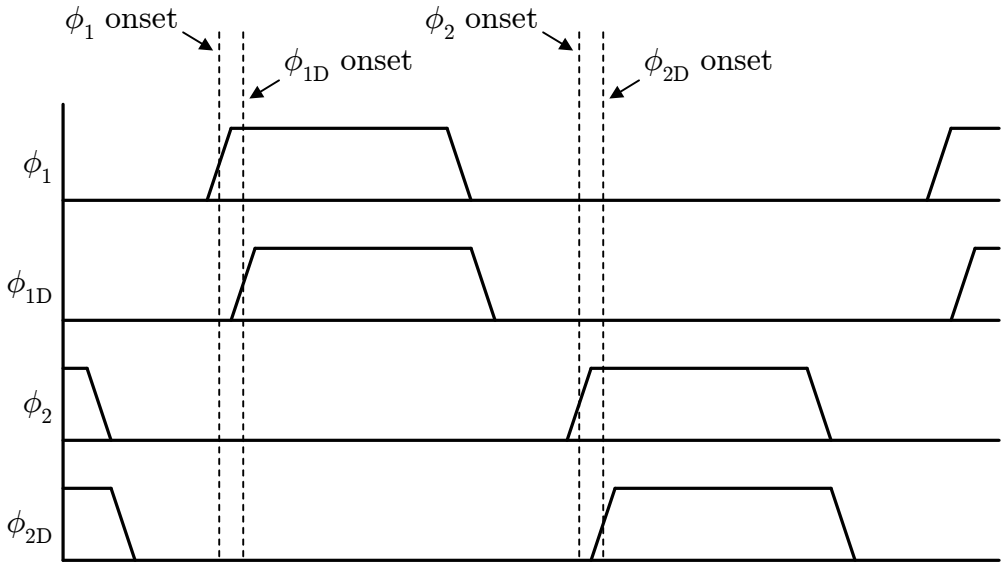

Figure 5.15: Clock phases produced by circuit of Figure 5.14. Onsets of phases labeled for reference.

the sample phase of the stage residue amplifier, and ϕ_2 and ϕ_{2D} always refer to the amplify phase of the stage residue amplifier, regardless of stage number.

Generally, all clock generating and buffering circuitry, as well as all switches driving by clocks, are supplied by $V_{DD,CLK}$.

5.3.5 Signal Path

The analog signal path of the SVADC-1 is fully differential. By convention, in the following discussions fully differential signals composed of the difference of the positive voltage V_{x+} and negative voltage V_{x-} are denoted as V_x (that is, $V_x = V_{x+} - V_{x-}$). All stages are designed to accommodate the same input and output analog signal range of ± 1 V differential around a 1.25 V common mode. This range is established by the reference voltages V_{REF+} and V_{REF-} (nominally set to 1.75 V and 0.75 V, respectively, for a V_{REF} of 1 V) and the common mode voltage V_{CM} (nominally set to 1.25 V). The reference voltages are global to the entire converter (see Section 5.4.3), while V_{CM} is locally generated in each stage (see Section 5.5.3).

Finally, the entire SVADC-1 is designed to operate at rates of 10 MS/s, twice the

specified 5 MS/s. This overdesign compensates for possible radiation-induced losses in circuit bandwidth and current drive (for example, due to radiation degradation of operational amplifier g_{m}).

5.4 Stage Design

With the groundwork laid in the previous section, this section now turns to the converter proper. In particular, this section describes the converter on a more architectural scale, proceeding stage by stage. It concludes with a description of reference (both voltage and current) generation.

5.4.1 Track-and-Hold Stage

The track-and-hold stage consists of the unity-gain, switched-capacitor amplifier of Figure 5.16. Broadly, this track-and-hold amplifier tracks the input V_{in} during the sample phase, sampling V_{in} onto $\mathrm{C_{S+}}$ and $\mathrm{C_{S-}}$ at the end of the phase. During the succeeding amplify phase, the amplifier develops the sampled signal at V_{out} and holds it for the remainder of the phase. As the amplifier illustrates many design principles deployed throughout the SVADC-1, it is worth investigating in some detail.

Sampling is triggered when the feedback switches $\mathrm{S_{F+}}$ and $\mathrm{S_{F-}}$ disconnect. Note that the bias across these switches is strongly set by the operational amplifier, guarding against input-signal dependency of the sampling instant (and the consequent nonlinearity). As shown in Figure 5.17, the sampling instant itself is defined by the falling edge of the early clock ϕ_{MAS} (as defined in Figure 5.13): this clock choice ensures a quiescent environment at the sensitive sampling instant. The choice also minimizes the active circuitry processing the sampling edge, reducing jitter. To further promote clean, linear sampling, bootstrap circuits are applied to the input switches $\mathrm{M_{S+}}$ and $\mathrm{M_{S-}}$: these circuits are detailed in Section 5.4.1.1.

Shorting the operational amplifier during the sampling phase confers benefits. Generally, amplifier resetting reduces unwanted input-signal memory effects between samples. However, in radiation environments, resetting is even more important as

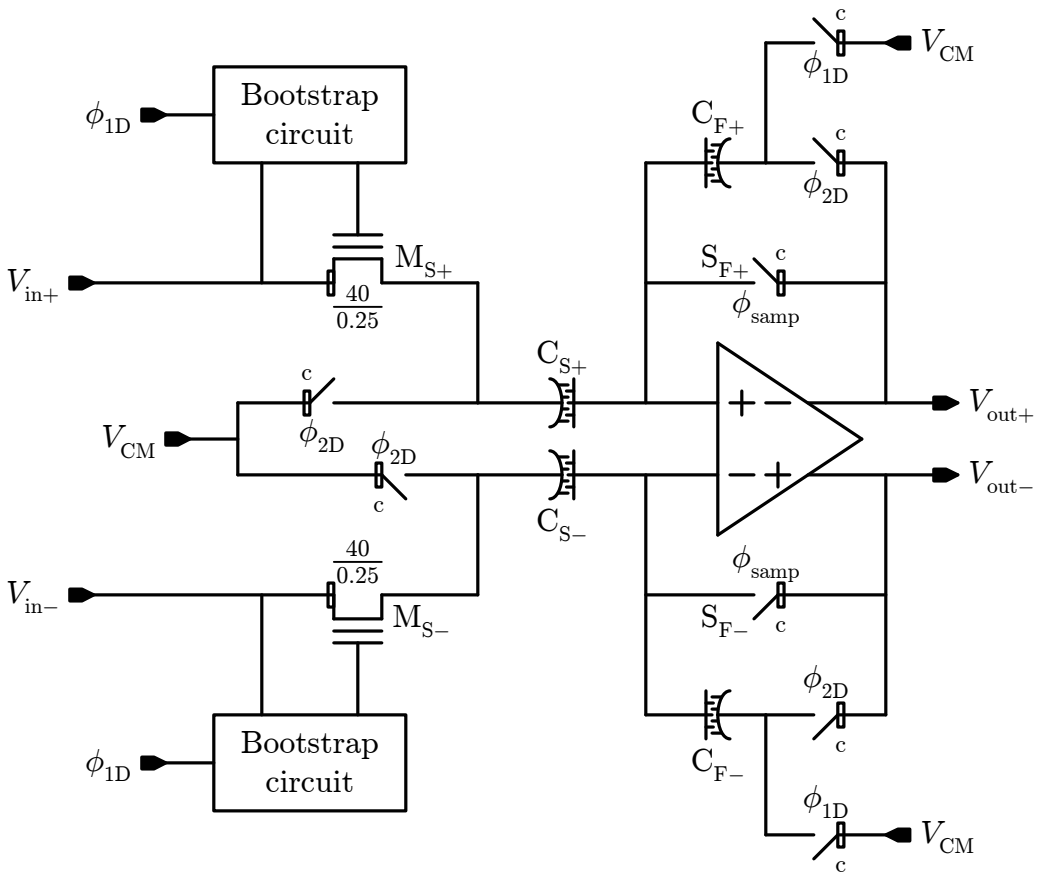

Figure 5.16: Track-and-hold amplifier circuit. $C_{S+} = C_{S-} = 2$ pF and $C_{F+} = C_{F-} = 2$ pF. For bootstrap circuit see Figure 5.18, for operational amplifier circuit see Figure 5.28, for ϕ_{samp} derivation see Figure 5.17.

it prevents charge accumulation at any node from ionizing strikes and curtails the duration of any single-event upsets. Self-resetting provides this resetting in an efficient manner, simultaneously reusing the operational amplifier current to reset the amplifier and charge the sample and feedback capacitors. Note that each node of the amplifier is thus driven by a low-impedance source (which quickly removes any excess radiation-induced charge) at least once a clock cycle. In addition, the charging of C_{F+} and C_{F-} is used for offset cancellation, which reduces the flicker noise of the operational amplifier and enables performance down to the 100 Hz required of the SVADC-1.

Finally, all switches are implemented with enclosed terminal layout for the NMOS

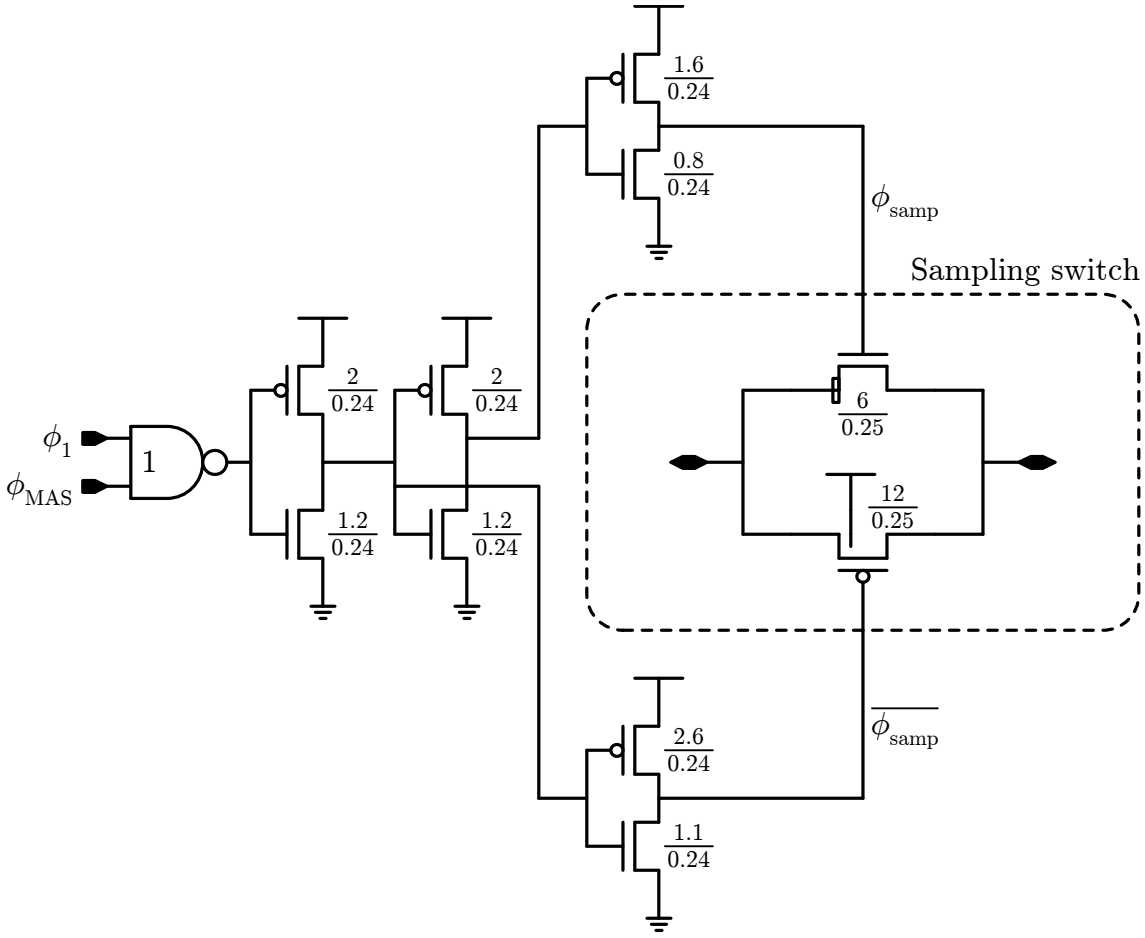

Figure 5.17: Derivation of the ϕ_{samp} clock driving S_{F+} and S_{F-} (drawn explicitly here and labeled "sampling switch") of track-and-hold amplifier circuit of Figure 5.16. ϕ_{samp} adopts the rising edge of ϕ_1 and the falling edge of ϕ_{MAS}. All supplies tied to $V_{\text{DD,CLK}}$, all device sizings in μm.

devices. The asymmetric switch is always oriented so that the enclosed terminal is driven by a low-impedance source, reducing the effect of its additional parasitic capacitance. In particular, the enclosed terminal is never placed at the operational amplifier inputs, as such placement increases charge sharing at these sensitive nodes during the amplify phase.

5.4.1.1 Bootstrap Circuit

The amplifier of Figure 5.16 divorces the sampling operation, accomplished by S_{F+} and S_{F-}, from the tracking operation, accomplished by M_{S+} and M_{S-}. A bootstrap circuit, shown in Figure 5.18, ensures the linearity of the latter: the circuit of Figure 5.18(a) bootstraps the sampling switch M_S via the capacitor C_B, while the multiplier circuit of Figure 5.18(b) provides the higher-than-V_{DD} voltage needed to charge C_B.

The circuit of Figure 5.18 is taken from *Abo* [1999, pp. 54–59], but modified for radiation. First, while Abo used LV devices, the SVADC-1 implementation instead opts for more robust, HV devices for all but the sampling switch.[21] Second, the bootstrap circuit is sensitive to discharge of C_B during ϕ_{1D}: to prevent radiation-induced I_{LEAK} from causing such discharge, enclosed terminal devices are used throughout. The enclosed terminals are judiciously oriented to reduce the impact of their added parasitic capacitance: they either avoid nodes sensitive to charge sharing (such as in the case of M_1 and M_2) or are driven by low-impedance sources when such avoidance is impossible (such as in the case of M_3). Finally, as the multiplier circuit is sensitive to charge sharing at the cross-coupling nodes, intentional parasitic capacitances C_{E+} and C_{E-} are added to these nodes to promote simulation-to-silicon accord.

In all, when ϕ_{1D} is high, simulations show that the bootstrap circuit establishes a V_{GS} of about 2.23 V, with a variation of about ±15 mV over the full 0.75–1.75 V input analog signal range. Subsequent simulations of the entire track-and-hold amplifier—including the bootstrap circuit, the operational amplifier (see Figure 5.28), and the

[21] Properly designed, no junction voltages in the bootstrap circuit should rise above V_{DD} (indeed, Abo amended the standard bootstrap circuit to help guarantee this property and consequently enhance reliability), although individual nodes (such as the gate of M_S) may. Hence, LV devices could be used: the use of HV devices simply further enhances reliability.

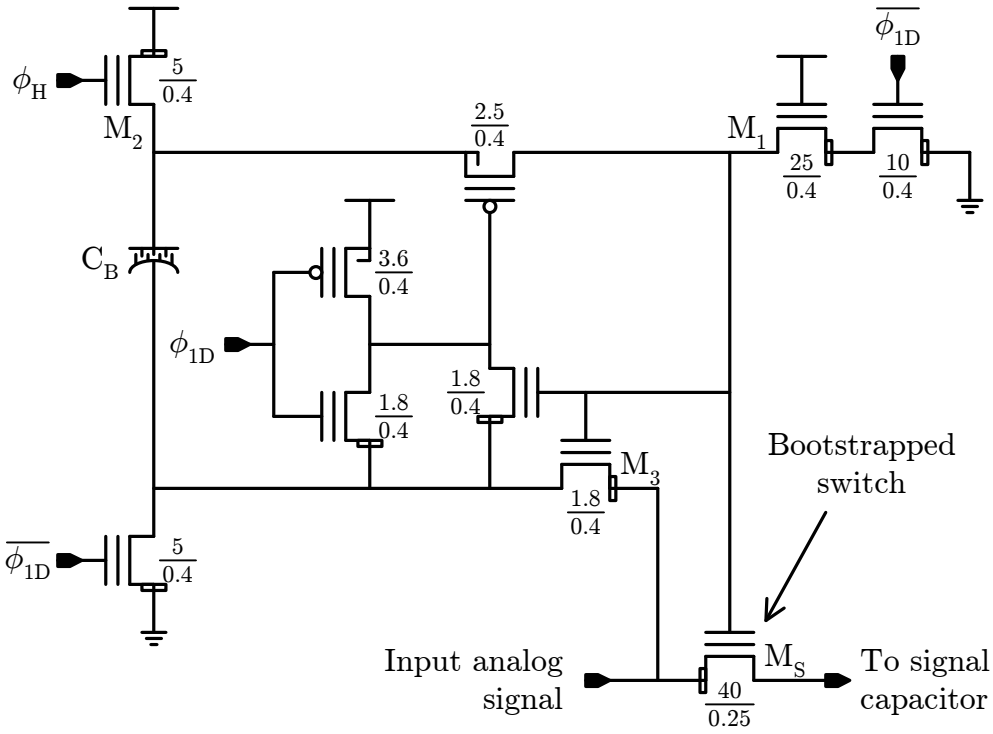

(a) Bootstrap circuit.

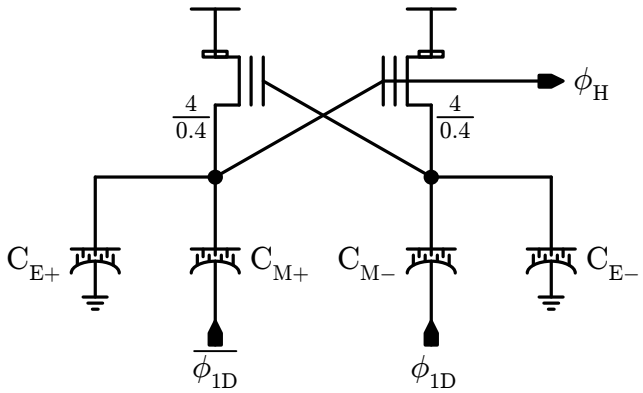

(b) Multiplier circuit.

Figure 5.18: Bootstrap circuit, including (a) the bootstrap circuit proper (one instance of this circuit is used for M_{S+}, and another for M_{S-}, of Figure 5.16) and (b) multiplier circuit (just one instance of this circuit is used to drive ϕ_H for both bootstrap circuits). $C_B = 1.5$ pF and $C_{M+} = C_{M-} = C_{E+} = C_{E-} = 200$ fF. All supplies tied to $V_{DD,CLK}$, all device sizings in μm.

nonlinear input capacitance of the ESD cells attached to the input pads (see Section 5.5.5)[22]—suggest that the amplifier achieves an SFDR on the order of 100 dB.[23]

5.4.2 Pipeline Stages

To describe the pipeline stage architecture, it is easiest to begin with the uncalibrated stages and then modify them to produce the calibrated stages.

5.4.2.1 Uncalibrated Stage

The uncalibrated stages are 2.8-bit stages with transfer functions as per Figure 4.14. The block diagram of the uncalibrated stage is shown in Figure 5.19 and the corresponding timing diagram is shown in Figure 5.20. During the sample phase, V_{in} is sampled by the residue amplifier and converted by the sub-ADC. The sub-ADC is fired early to ensure its decision is ready by the beginning of the amplify phase. Since V_{in} is still being developed by the upstream stage amplifier, the sub-ADC effectively operates on an offsetted version of V_{in}, however, the stage redundancy is sufficient that this offset does not harm converter linearity (see Section 4.2.2). During the amplify phase, the thermometer-encoded sub-ADC decisions $D\langle 1\rangle$ through $D\langle 6\rangle$ determine the sub-DAC voltage implemented by the residue amplifier. In addition, a binary-encoded version of the decisions $B\langle 0\rangle$ through $B\langle 2\rangle$ is flopped and transferred to the I/O pads for output.

The uncalibrated stage residue amplifier is shown in Figure 5.21. The general operation of this residue amplifier was described in Section 4.2.3.1. Here, $D\langle 1\rangle$ through $D\langle 6\rangle$ feed only the middle sample capacitors, highlighted in blue. The extremum capacitors, on the other hand, are defaulted to set decisions to create the wider extremum segments of the 2.8-bit transfer function. Notably, while C_{1+} (and C_{8-}) and C_{8+} (and C_{1-}) do not need switches to $V_{\text{REF}-}$ and $V_{\text{REF}+}$, respectively, these switches are nonetheless included so that the switch networks attached to each

[22] Although it is unlikely the nonlinear capacitance of the reverse-biased ESD diodes is limiting at the input frequencies of the SVADC-1 [*Chun and Murmann*, 2006], it is nonetheless included in these simulations.

[23] Specifically, the amplifier is clocked at 10 MHz while processing a full-scale, 1-MHz input sinusoid, and the resulting harmonic distortion used to estimate the SFDR.

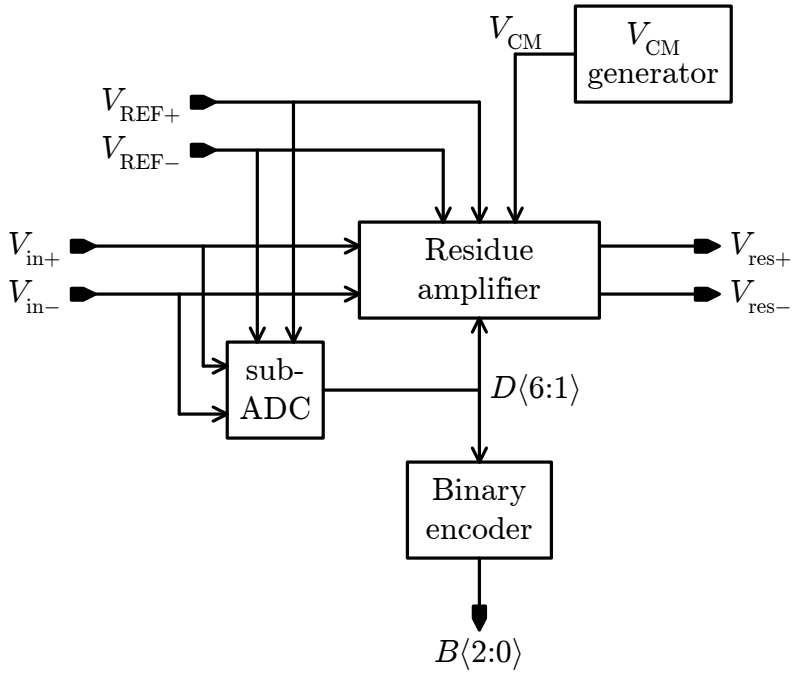

Figure 5.19: Uncalibrated stage block diagram. Transfer function (2.8-bit) as per Figure 4.14. For stage timing see Figure 5.20, for residue amplifier see Figure 5.21.

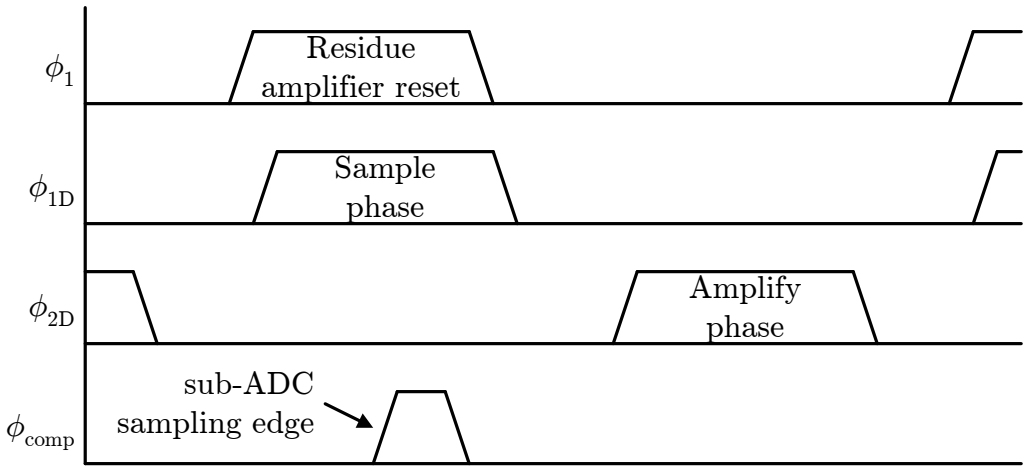

Figure 5.20: Timing diagram for pipeline stages. ϕ_{1D} and ϕ_{2D} labeled as per their use in the residue amplifier. sub-ADC fires on rising edge of ϕ_{comp}.

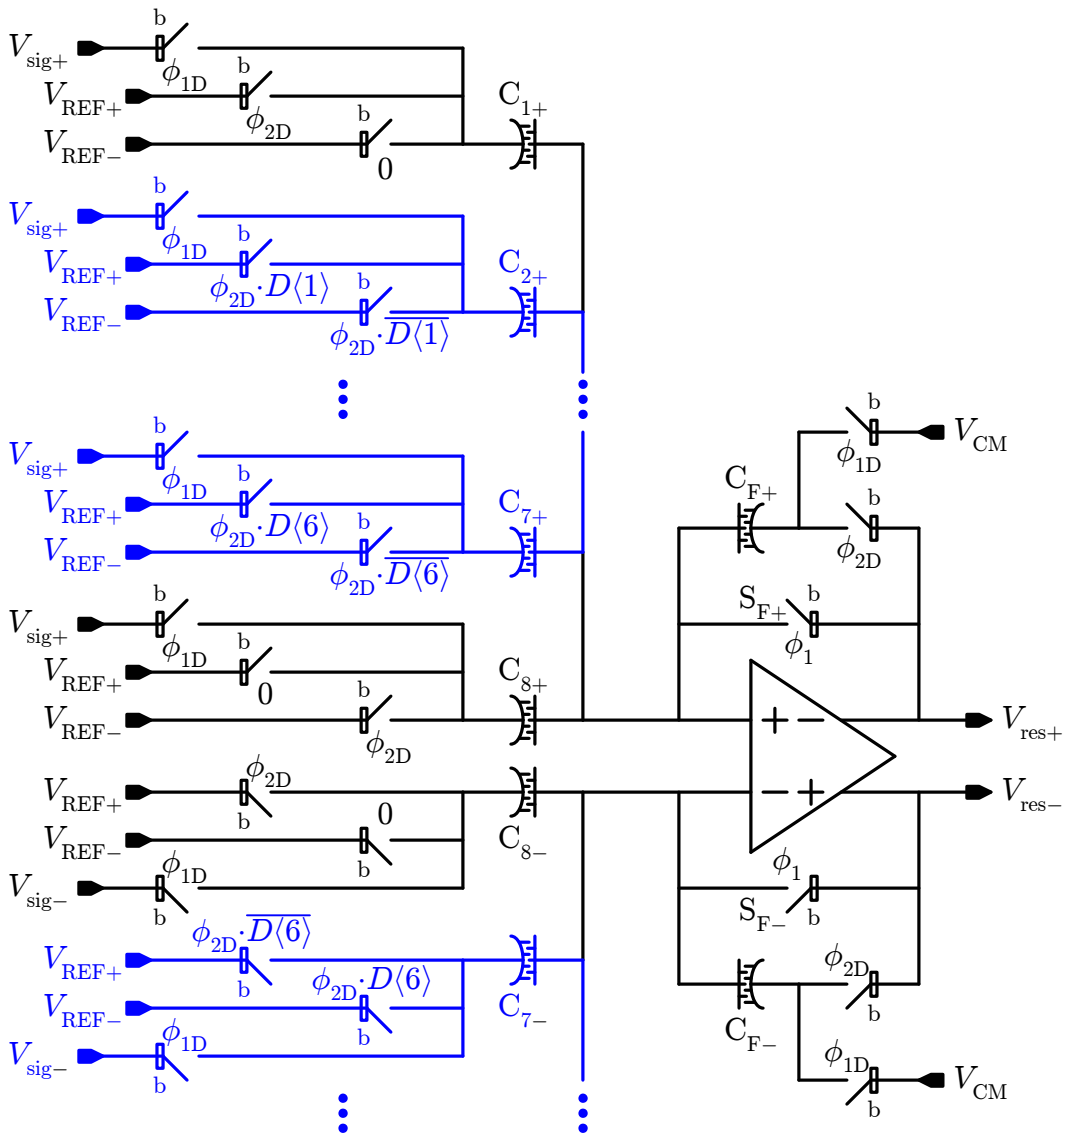

Figure 5.21: Uncalibrated stage residue amplifier. Sample capacitors driven by sub-ADC decisions shown in blue, set decision sample capacitors shown in black. $C_{1+} = \ldots = C_{8+} = 62.5$ fF, $C_{F+} = 125$ fF; negative branch capacitors same as positive branch equivalents. For operational amplifier circuit see Figure 5.30.

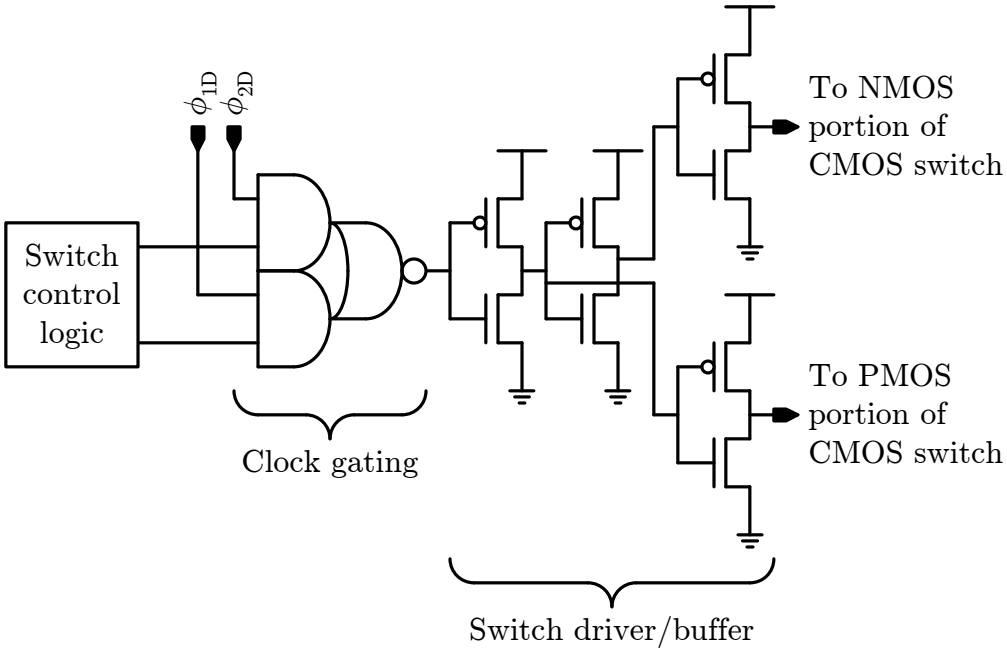

Figure 5.22: Gating circuit for driving switches of sample capacitors in both uncalibrated and calibrated residue amplifiers. Device sizing differs depending on the switch type being driven.

sample capacitor are the same, reducing dynamic mismatch between these branches. The switch driver circuit itself is shown in Figure 5.22. Note that the switch control logic is symmetrically gated by either ϕ_{1D} or ϕ_{2D} as late as possible. This practice 1) reduces dephasing between ϕ_{1D} and ϕ_{2D}, and 2) eliminates possible clock arrival differences due to different delays through the switch control logic. The latter proves especially important in the calibrated stages where the switch control logic may differ greatly between branches depending on the calibration logic implementation.

The residue amplifier of Figure 5.21 adopts many practices described in the track-and-hold amplifier, including: 1) early-release, feedback-embedded sampling switches S_{F+} and S_{F-}—here driven by the early clock ϕ_1—to provide clean sampling; 2) offset-cancellation to maintain noise performance down to 100 Hz; 3) self-resetting circuitry to efficiently promote radiation single-event upset tolerance; and 4) deliberately oriented enclosed terminal layout on switches to prevent radiation-induced I_{LEAK}.

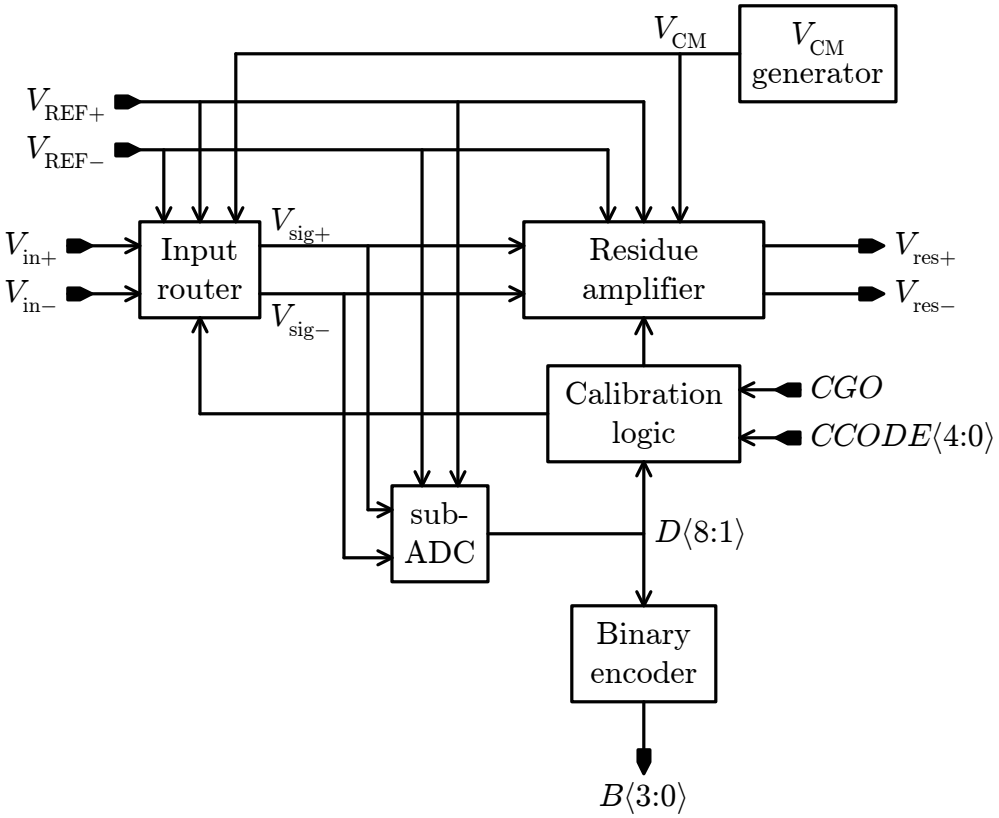

Figure 5.23: Calibrated stage block diagram. Transfer function (3.1-bit) as per Figure 4.29. For stage timing see Figure 5.20, for residue amplifier see Figure 5.24.

5.4.2.2 Calibrated Stage

To produce a calibrated stage, the stage of Figure 5.19 is modified as per Figure 5.23. In particular, a calibration logic block is added to reroute the residue amplifier during calibration. An input router block is also added as an aid for radiation testing (as explained later in this section). The calibrated stages are 3.1-bit stages with transfer functions as per Figure 4.29.

The modified residue amplifier is shown in Figure 5.24. In normal operation the calibrated stage residue amplifier operates as that of the uncalibrated stage with C_{3+} (C_{3-}) through C_{10+} (C_{10-}) driven by the sub-ADC decisions $D\langle 1\rangle$ through $D\langle 8\rangle$, respectively, and the remaining sample capacitors defaulted to set decisions. Formally,

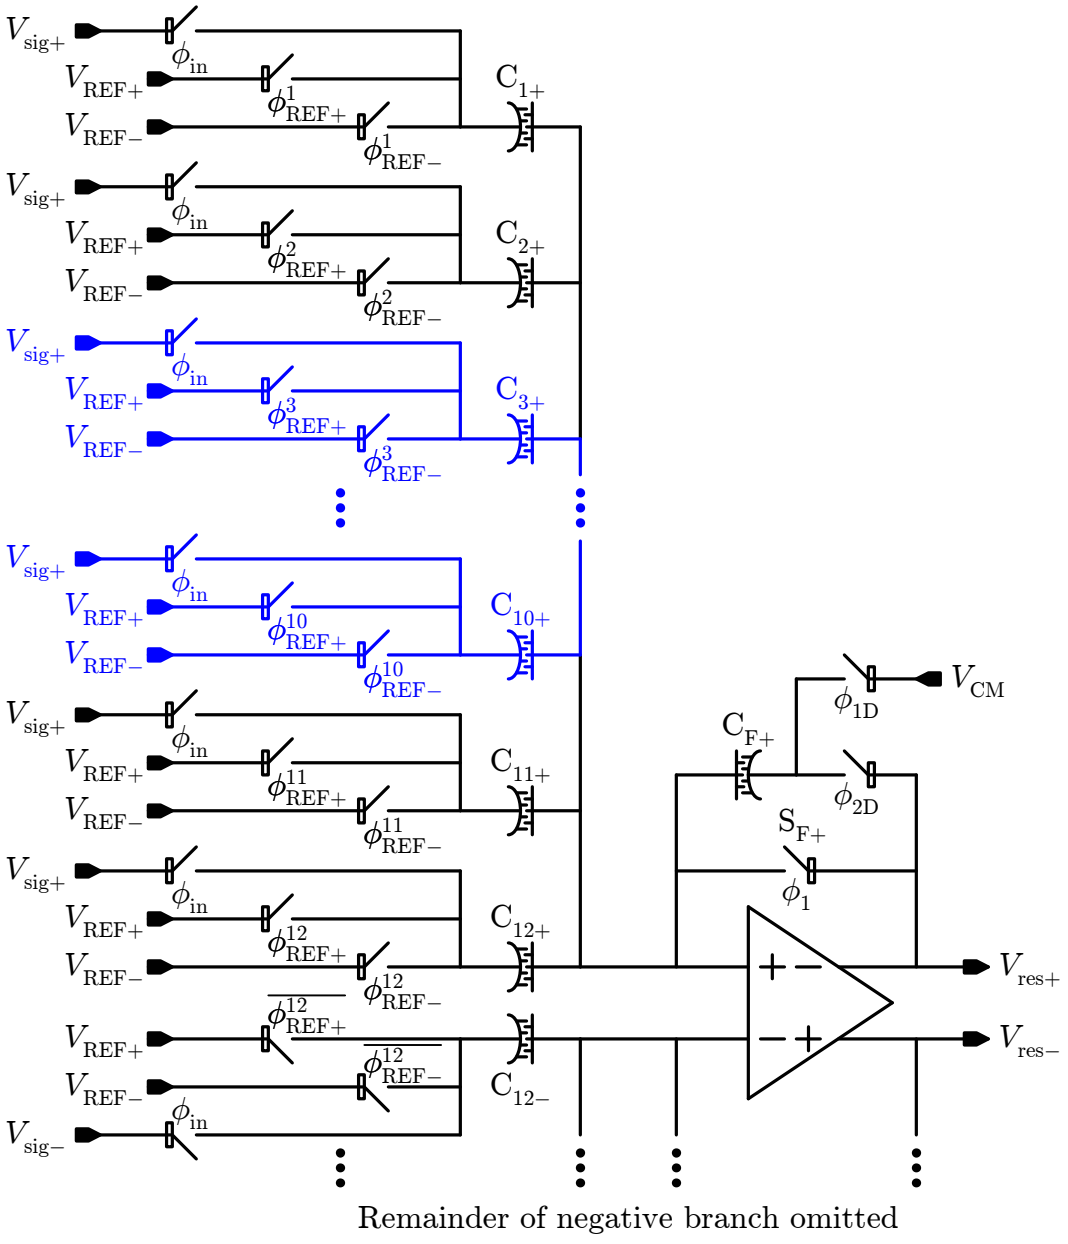

Figure 5.24: Calibrated stage residue amplifier. Sample capacitors driven by sub-ADC decisions shown in blue, set decision sample capacitors shown in black. For stage 1: $C_{3+} = \ldots = C_{10+} = 93.75$ fF, $C_{1+} = C_{2+} = C_{11+} = C_{12+} = 62.5$ fF, and $C_{F+} = 250$ fF. For stages 2 and 3: $C_{3+} = \ldots = C_{10+} = 46.875$ fF, $C_{1+} = C_{2+} = C_{11+} = C_{12+} = 31.25$ fF, and $C_{F+} = 125$ fF. Negative branch capacitors same as positive branch equivalents. All switches type "c" for stage 1, type "b" for stages 2 and 3. For operational amplifier circuit see Figure 5.29.

$\phi_{\text{in}} = \phi_{1\text{D}}$ and:

$$\phi^r_{\text{REF}+} = \begin{cases} \phi_{2\text{D}} & , \quad r = 1, 2 \\ \phi_{2\text{D}} D\langle r-2 \rangle & , \quad r = 3, \ldots, 10 \\ 0 & , \quad r = 11, 12 \end{cases} \tag{5.5}$$

while:

$$\phi^r_{\text{REF}-} = \begin{cases} 0 & , \quad r = 1, 2 \\ \phi_{2\text{D}} \overline{D\langle r-2 \rangle} & , \quad r = 3, \ldots 10 \\ \phi_{2\text{D}} & , \quad r = 11, 12 \end{cases} \tag{5.6}$$

Calibration is invoked by an active *CGO* signal and the resulting stage configuration dictated by *CCODE*. Both digital control signals are flopped internally within the stage to align them with the local $\phi_{1\text{D}}$ and $\phi_{2\text{D}}$. During calibration, the calibration logic preempts the sub-ADC decision and reroutes the residue amplifier into autozero and DAC-difference configurations. For the autozero configuration, the logic sets $\phi_{\text{in}} = 0$ and:

$$\phi^r_{\text{REF}+} = \begin{cases} \phi_{1\text{D}} + \phi_{2\text{D}} & , \quad r = 1, \ldots 6 \\ 0 & , \quad r = 7, \ldots 12 \end{cases} \tag{5.7}$$

while:

$$\phi^r_{\text{REF}-} = \begin{cases} 0 & , \quad r = 1, \ldots 6 \\ \phi_{1\text{D}} + \phi_{2\text{D}} & , \quad r = 7, \ldots 12 \end{cases} \tag{5.8}$$

For the DAC-difference configurations, the connection of one sample capacitor is changed between $\phi_{2\text{D}}$ and $\phi_{1\text{D}}$. Assuming the capacitor C_{m+} (C_{m-}) is thus singled out, the logic again sets $\phi_{\text{in}} = 0$ but:

$$\phi^r_{\text{REF}+} = \begin{cases} \phi_{1\text{D}} + \phi_{2\text{D}} & , \quad r < m \\ \phi_{1\text{D}} & , \quad r = m \\ 0 & , \quad r > m \end{cases} \tag{5.9}$$

while:

$$\phi^r_{\text{REF}-} = \begin{cases} 0 & , \quad r < m \\ \phi_{2\text{D}} & , \quad r = m \\ \phi_{1\text{D}} + \phi_{2\text{D}} & , \quad r > m \end{cases} \tag{5.10}$$

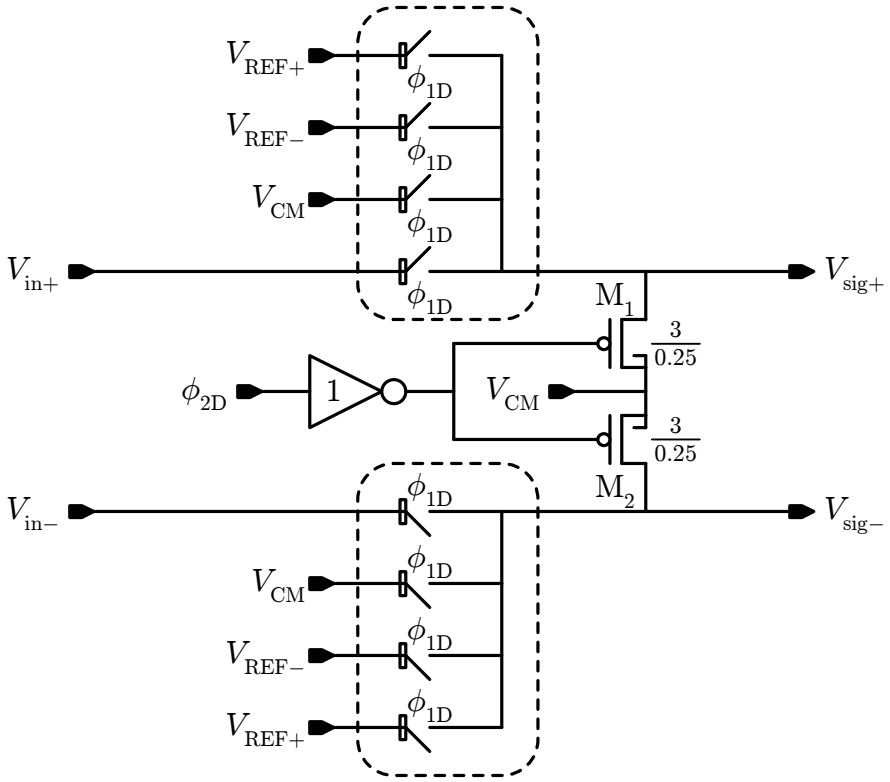

Figure 5.25: Input router for calibrated stages. At most one switch in each dashed box is active during $\phi_{1\mathrm{D}}$, others remain disconnected during $\phi_{1\mathrm{D}}$. CMOS switches type "e" for stage 1, type "d" for stages 2 and 3. All supplies tied to $V_{\mathrm{DD,CLK}}$, all device sizings in μm.

Although technically only the configurations isolating C_{3+} (C_{3-}) through C_{10+} (C_{10-}) are needed for self-calibration, configurations isolating all sample capacitors are included for debugging purposes and as an aid in interpreting residue amplifier radiation response. For this reason, C_{1+} (C_{1-}) and C_{2+} (C_{2-}) are implemented as separate capacitors: had these capacitors been combined, the resulting capacitor would have overranged the backend ADC during DAC differencing (in essence, the capacitor splitting solution of Section 4.3.2.4 is applied to these capacitors). The same is true of C_{11+} (C_{11-}) and C_{12+} (C_{12-}).

Also as an aid for assessing radiation response, an input router is included in the calibrated stages. The router schematic is shown in Figure 5.25. Also controlled by *CGO* and assigned *CCODE* values, it routes combinations of $V_{\mathrm{REF+}}$, $V_{\mathrm{REF-}}$, and

V_{CM} into the stage input.[24] The stage is then operated normally (i.e., as described in Equations (5.5) and (5.6)). Such test signal injection enhances stage-by-stage visibility. Additional switches M_1 and M_2 prevent the V_{SIG+} and V_{SIG-} nodes from floating during ϕ_{2D}, reducing upset sensitivity.[25]

Finally, the robust design practices of the uncalibrated stage—as discussed in Section 5.4.2.1—are employed in the calibrated stage as well.

5.4.3 Analog Reference Generation

The SVADC-1 also includes circuits for reference voltage and bias current generation.

Generation of the reference voltages V_{REF+} and V_{REF-} is shown in Figure 5.26. The reference voltages can be generated on chip via a buffered resistor ladder (by externally connecting V_{REF+} and V_{SENSE+}, and V_{REF-} and V_{SENSE-}, to close the negative feedback loops around the operational amplifiers), or injected from off-chip sources (by driving V_{REF+} and V_{REF-} externally).[26] Notably, the on-chip resistor ladder is given its own supply $V_{DD,REF}$ to reduce noise coupling. Finally, large off-chip capacitors placed near the chip act as charge reservoirs to rapidly supply V_{REF+} and V_{REF-} as needed.

All analog circuitry in the SVADC-1 is biased from master current I_{master}, generated as shown in Figure 5.27. Again, provision is made for either internal biasing from an on-chip resistor, or external biasing from an off-chip element, such as a resistor or current source. I_{master} is then mirrored and distributed to each stage. Each stage subsequently derives all internal biasing from this single received current. The use of PMOS devices here is intentional: as $|V_{T,P}|$ only rises with dose (versus $|V_{T,N}|$ which may rise or fall) I_{master}—and hence the analog power consumption—if anything reduces with dose. From a system perspective, such behavior is more desirable than

[24] Specifically, the following combinations are enabled (given in the form of the resulting V_{SIG+}/V_{SIG-}): V_{in+}/V_{in-} (of course), V_{REF+}/V_{CM}, V_{CM}/V_{REF+}, V_{REF-}/V_{CM}, V_{CM}/V_{REF-}, V_{CM}/V_{CM}, V_{REFP+}/V_{REF-}, and V_{REF-}/V_{REF+}. Note that these combinations exercise a range of both differential and common mode signals.

[25] As the signal range for these switches need not be large, and their speed need not be great, M_1 and M_2 are implemented by PMOS devices since their I_{LEAK} does not increase with dose.

[26] The latter course is the original intent and is adopted during performance assessment. However, the former often proves useful in testing environments where it is complicated to inject high quality off-chip references; for examples, see Sections B.2.1 and K.3.1.

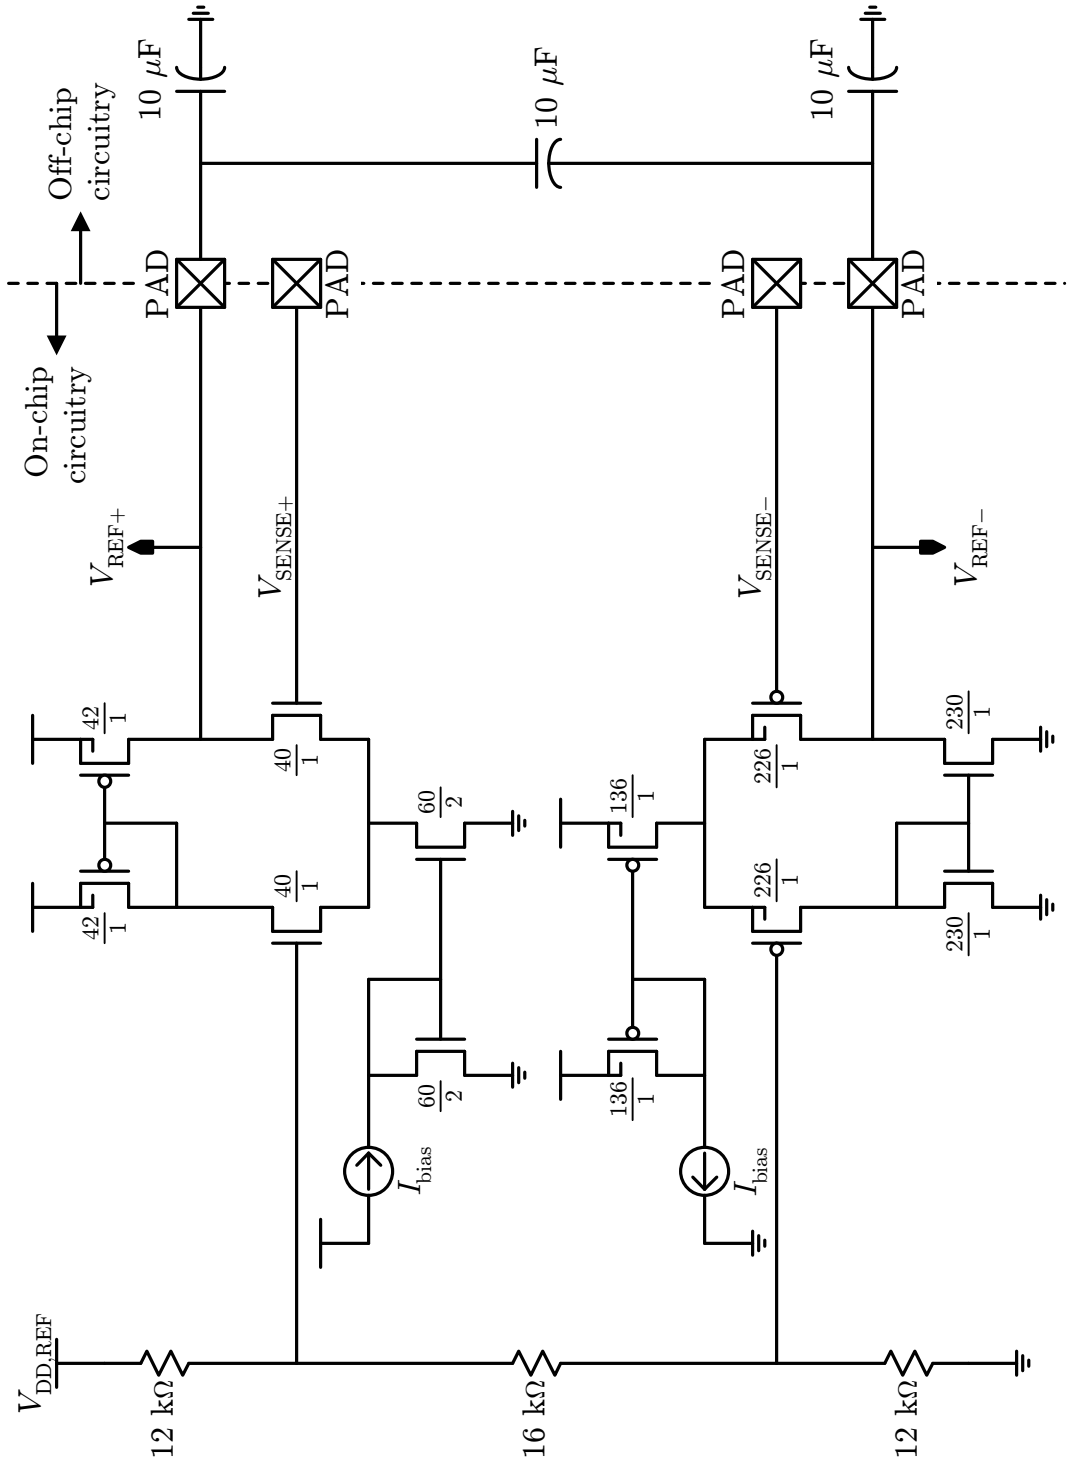

Figure 5.26: Reference voltage $V_{\mathrm{REF+}}$ and $V_{\mathrm{REF-}}$ generation. $I_{\mathrm{bias}} = 62.5$ μA. All supplies tied to $V_{\mathrm{DD,A}}$ unless otherwise indicated, all device sizings in μm.

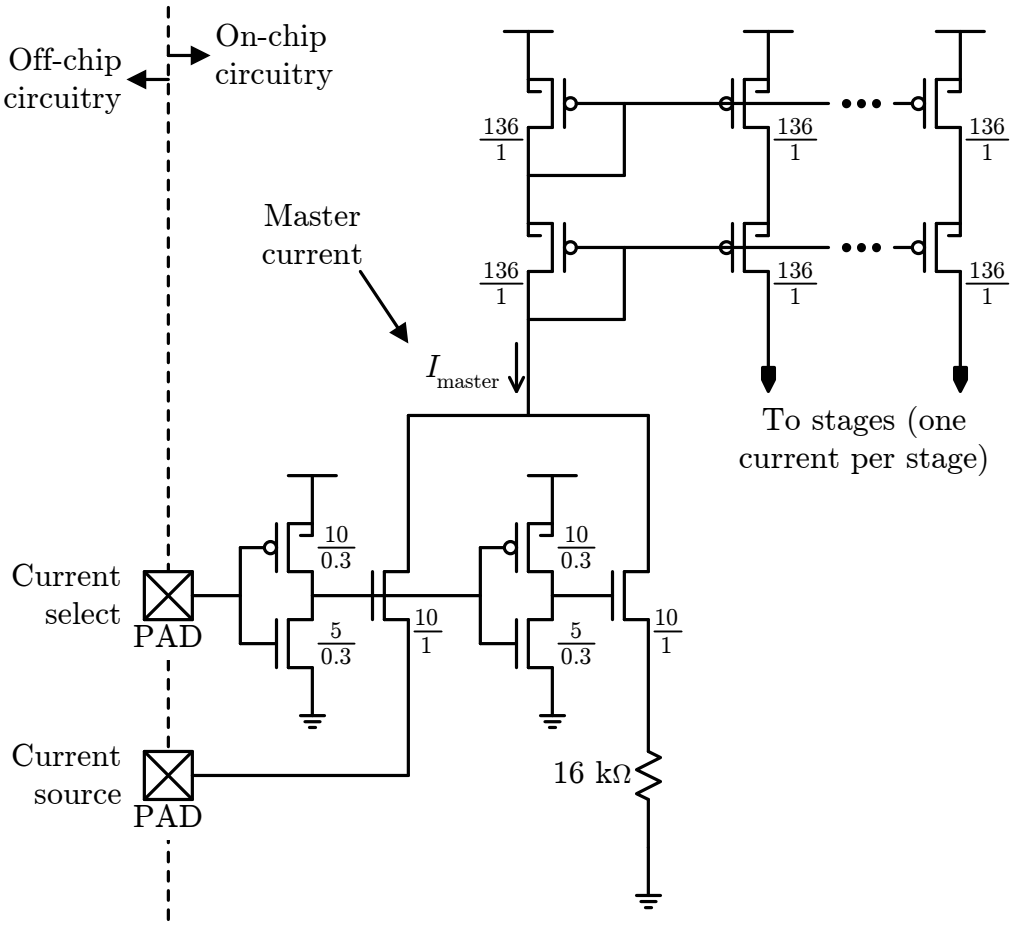

Figure 5.27: Reference current I_{master} generation. All supplies tied to $V_{\text{DD,A}}$, all device sizings in μm.

its opposite. This philosophy extends to the stages themselves: all bias branches throughout the SVADC-1 include PMOS devices to guard against rising bias current with dose.

Generally, throughout the SVADC-1 analog biasing is performed by standard current mirroring with cascoding if possible. In radiation environments, current mirroring can be complicated by the bias dependence of radiation-induced V_{T} shifts: as the biases of the mirroring devices differ, their radiation-induced shifts can differ, compromising mirror accuracy. It is possible to safeguard against this imbalance by adopting elaborate biasing schemes, such as that advocated by *Edwards et al.* [1999].

In the SVADC-1, though, such sophistications are deemed unnecessary. First, given the gate oxide thinness, though V_T shifts of differently biased devices differ, overall they ought to be small enough so as not to seriously endanger accuracy. Second, circuits are generally overdesigned so that change in the mirrored current consumes margin, not performance.

5.5 Building Blocks

There is ample opportunity for design reuse of many basic building blocks throughout the SVADC-1. The circuit implementation of these blocks is addressed here.

5.5.1 Operational Amplifiers

The schematics for the operational amplifiers of the track-and-hold amplifier, calibrated stage residue amplifier, and uncalibrated stage residue amplifier are shown in Figures 5.28, 5.29, and 5.30, respectively. The schematics include biasing circuits. All the operational amplifiers use the switched-capacitor common mode feedback circuit of Figure 5.31.

The operational amplifiers are PMOS-input folded cascode amplifiers. Throughout, the input devices M_1 and M_2 are sized relatively small to reduce the input parasitic capacitance (the accompanying undesired increases in offset and flicker noise are mitigated by offset cancellation). Generally, thinner device sizings are used along the signal path (to increase bandwidth), and longer are used in the current sources (to increase R_{out}). Finally, note that the input current (i.e., through M_1 and M_2) is less than the output current (i.e., through M_3 and M_4). This allotment prevents M_3 and M_4 from turning off during slewing (as can occur at the beginning of the amplify phase) and hastens recovery from extremum input conditions (as can occur during radiation-induced single-event upsets).

The simulated performance of the operational amplifiers is summarized in Table 5.2. Several specifications are overdesigned, including gain linearity and unity gain frequency to compensate for radiation-induced g_m loss. The linearity is computed as

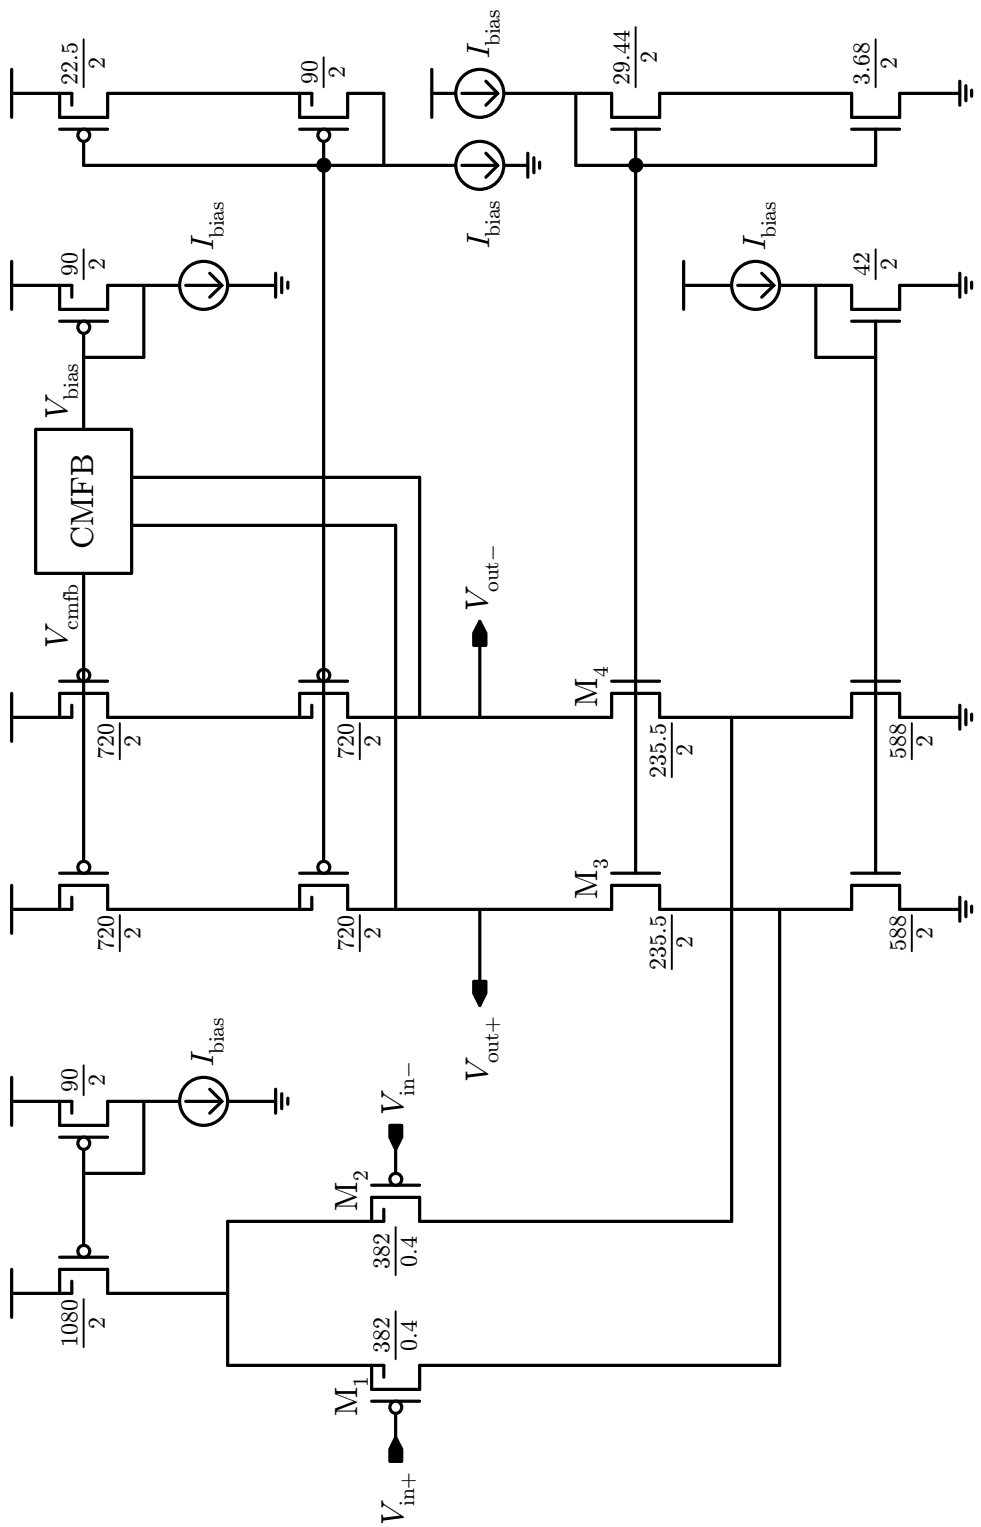

Figure 5.28: Operational amplifier for the track-and-hold amplifier of Figure 5.16. $I_{\text{bias}} = 62.5\ \mu\text{A}$. For common mode feedback (CMFB) circuit see Figure 5.31. All supplies tied to $V_{\text{DD,A}}$, all device sizings in μm.

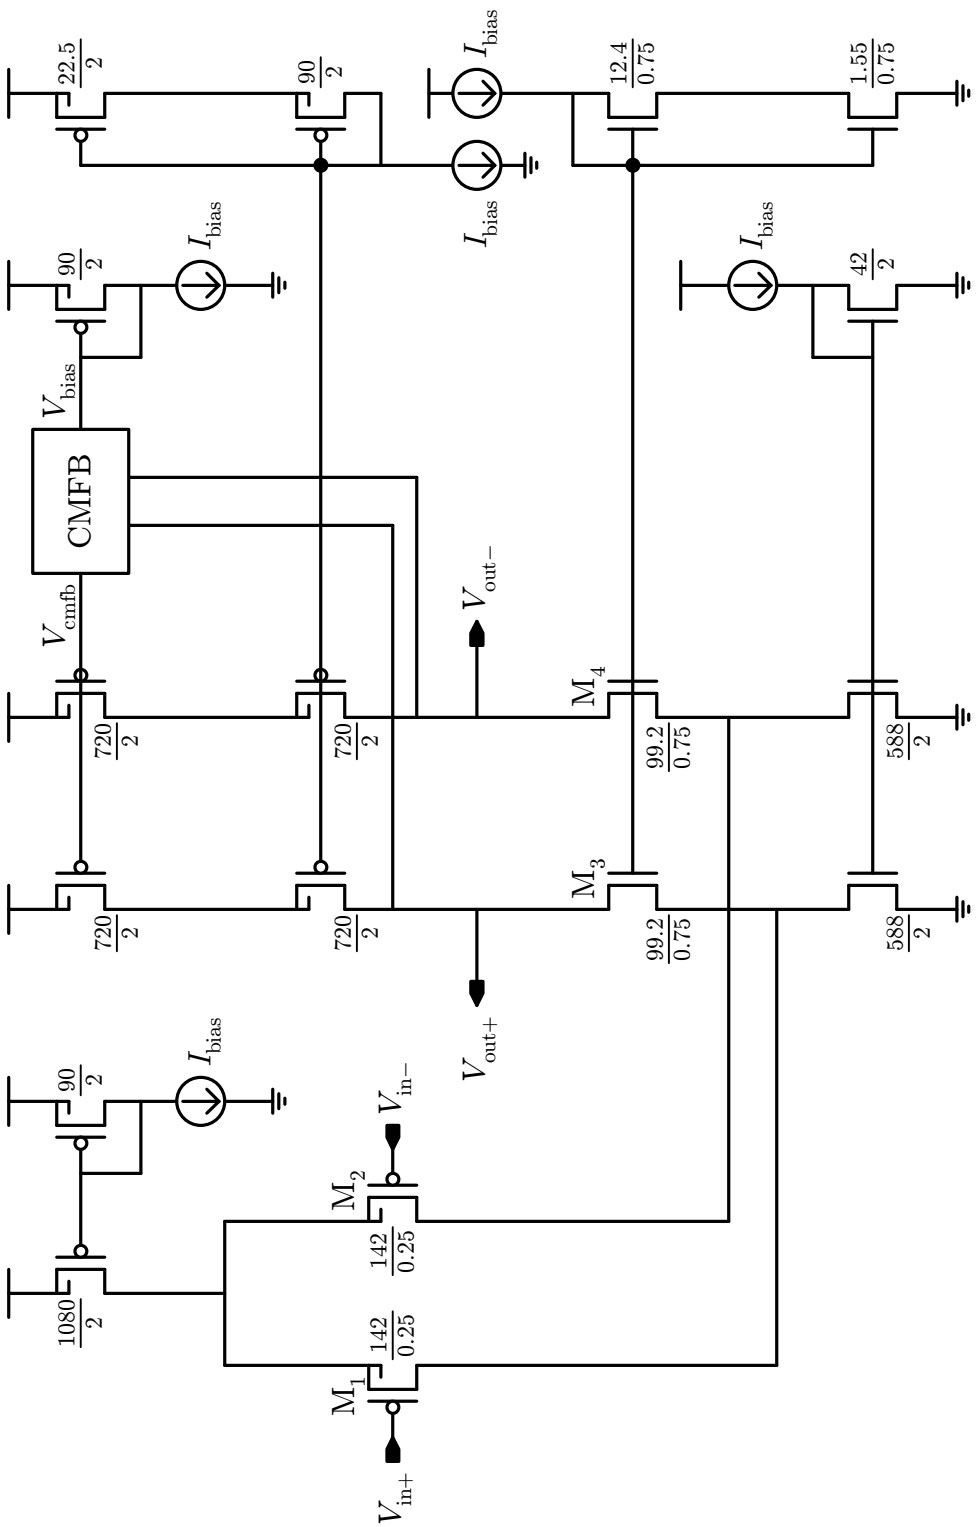

Figure 5.29: Operational amplifier for the calibrated stage residue amplifier of Figure 5.24. $I_{\text{bias}} = 62.5\ \mu\text{A}$. For common mode feedback (CMFB) circuit see Figure 5.31. All supplies tied to $V_{\text{DD,A}}$, all device sizings in μm.

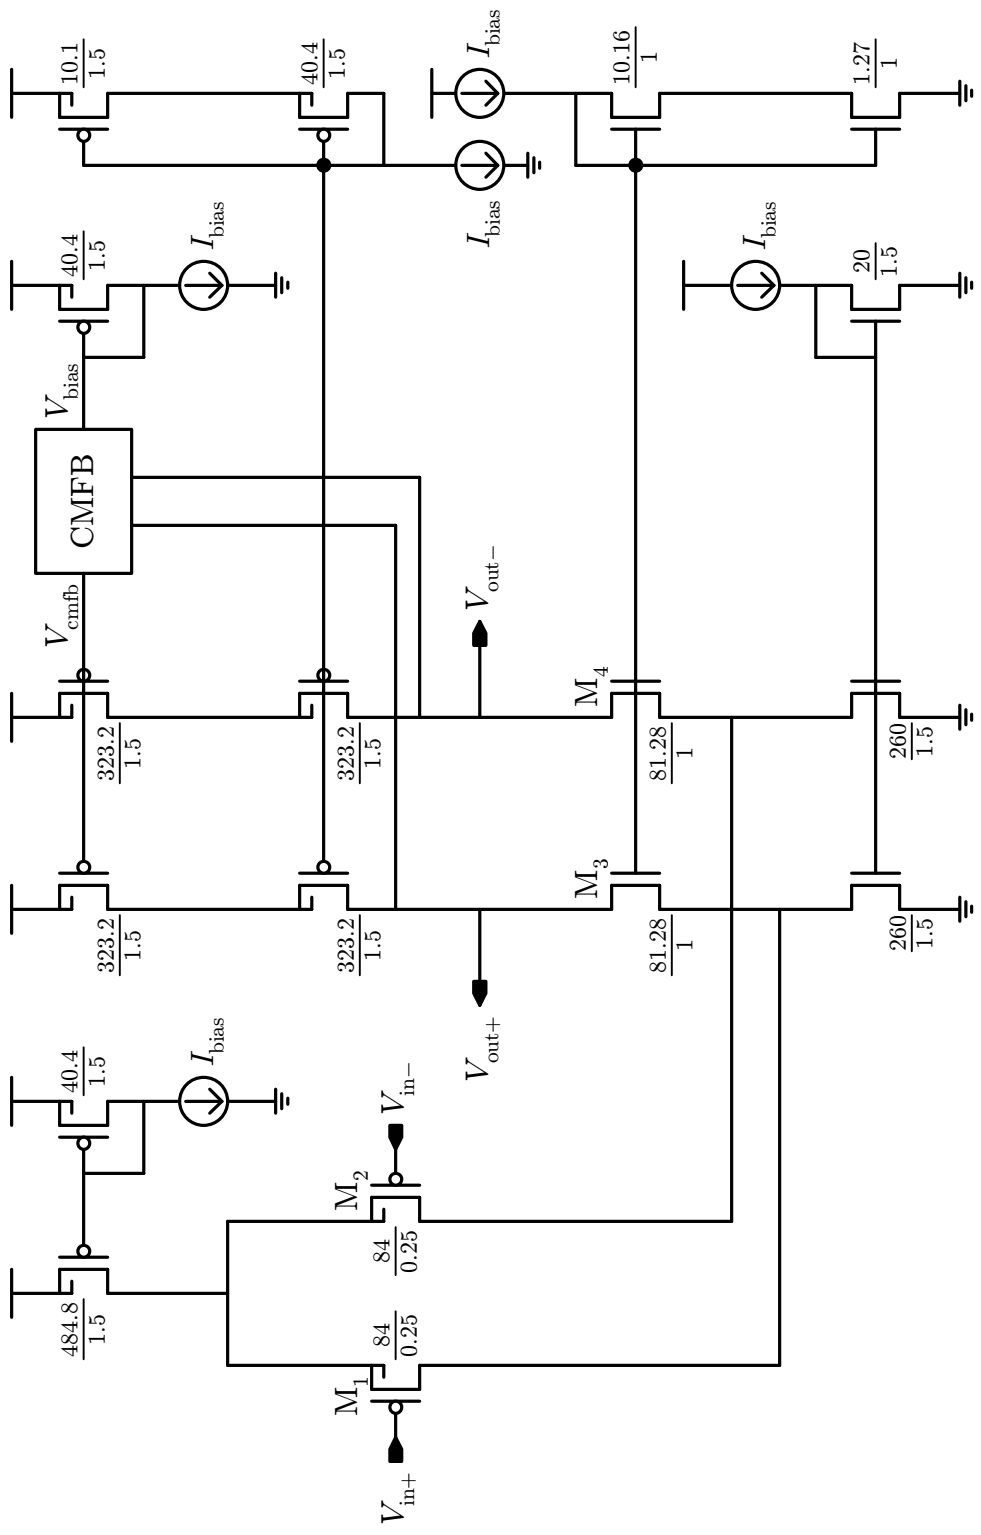

Figure 5.30: Operational amplifier for the uncalibrated stage residue amplifier of Figure 5.21. $I_{\text{bias}} = 37.5\ \mu\text{A}$. For common mode feedback (CMFB) circuit see Figure 5.31. All supplies tied to $V_{\text{DD,A}}$, all device sizings in μm.

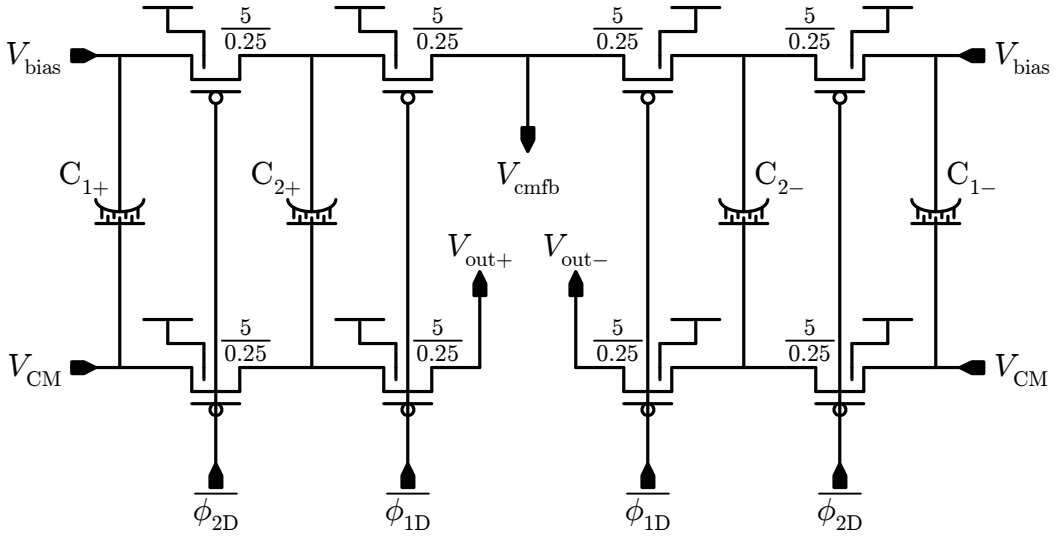

Figure 5.31: Common mode feedback circuit for all operational amplifiers. $C_{1+} = C_{1-} = C_{2+} = C_{2-} = 400$ fF. All supplies tied to $V_{\text{DD,A}}$, all device sizings in μm.

per Section 4.2.4, although the differential output at 1.2 V (instead of at 1.0 V) is used for A_2 to accommodate possible output range shifts due to process or radiation variations.

Finally, a note on the common mode feedback circuit. The use of PMOS switches here is intentional: as these devices do not require enclosed terminal layout, simulation can employ accurate device models. It has since been pointed out, though, that during ϕ_{2D} the common mode feedback is but weakly functional: as C_{2+} and C_{2-} are disconnected from the operational amplifier, the common mode is maintained by the parasitic capacitance at V_{cmfb}. Future revisions should thus consider replacing the circuit of Figure 5.31 with a continuous common mode feedback circuit.

5.5.2 Comparator

Each sub-ADC (and the terminating ADC) is composed of a series of parallel comparators. Figure 5.32 shows the comparator architecture and timing diagram. The comparator adopts a conventional architecture of a preamplifier (which samples V_{in} and subtracts from it a ratio of V_{REF} to implement the comparator transition

Characteristic	Track-and-hold stage	Calibrated stage			Uncalibrated stage	
Stage	0	1	2	3	4	5
Power						
Active	4.375 mW	4.375 mW	4.375 mW	4.375 mW	2.625 mW	2.625 mW
Biasing	781 μW	781 μW	781 μW	781 μW	468 μW	468 μW
Gain						
$V_{\text{out}} = 0$ V	81.7 dB	74.1 dB	74.1 dB	74.1 dB	73.6 dB	73.6 dB
$V_{\text{out}} = 1.2$ V	80.0 dB	72.5 dB	72.5 dB	72.5 dB	72.1 dB	72.1 dB
Linearity	16.4 bits	14.9 bits	14.4 bits	14.4 bits	14.8 bits	14.8 bits
Transconductance	5.00 mS	4.64 mS	4.64 mS	4.64 mS	2.77 mS	2.77 mS
Input capacitance	685 fF	169 fF	169 fF	169 fF	100 fF	100 fF
Frequency response #1						
Load capacitance	4.69 pF	1.42 pF	794 fF	794 fF	725 fF	725 fF
Unity gain frequency	135 MHz	304 MHz	401 MHz	401 MHz	361 MHz	361 MHz
Phase margin	57°	63°	56°	56°	55°	55°
Frequency response #2						
Load capacitance	3.20 pF	1.62 pF	1.66 pF	1.40 pF	1.40 pF	2.08 pF
Unity gain frequency	176 MHz	269 MHz	278 MHz	307 MHz	237 MHz	177 MHz
Phase margin	48°	66°	65°	63°	66°	72°

Table 5.2: Operational amplifier performance summary. Input capacitance and load capacitance refer to single-ended values. Frequency response #1 is indicative of performance during sampling phase in given stage, frequency response #2 indicative of same during amplifying phase (while there are only three operational amplifier designs, performance may vary depending on their use in a given stage).

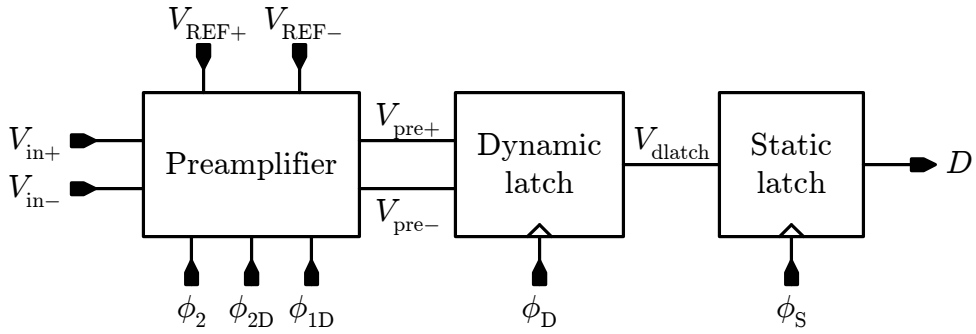

(a) Comparator architecture.

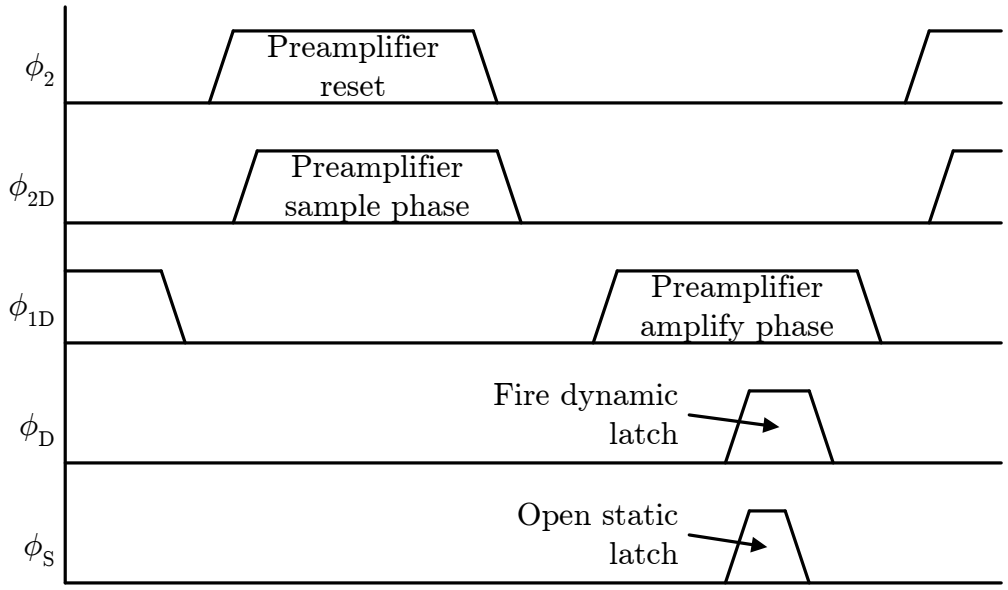

(b) Comparator timing diagram.

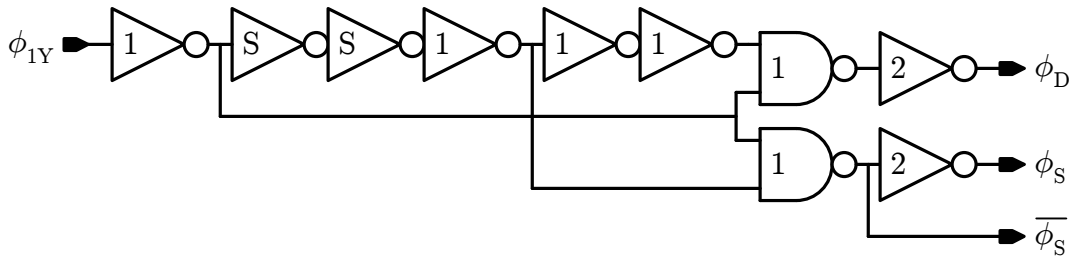

(c) Dynamic and static latch clock generation.

Figure 5.32: Comparator implementation, including (a) architecture, (b) timing diagram, and (c) generation of ϕ_{D} and ϕ_{S} for dynamic and static latches. Regarding (c): $\phi_{1\mathrm{Y}}$ is a delayed version of ϕ_1, and all supplies are tied to $V_{\mathrm{DD,D}}$.

level) followed by a dynamic latch (which rapidly gains the result through positive feedback) and terminated by a static latch (which stores the result).

The preamplifier circuit is shown in Figure 5.33. This preamplifier is adopted from *Limotyrakis et al.* [2005] but introduces an explicit amplifier to reduce kickback noise from the early dynamic latch firing. Following the philosophy described in Section 5.4.1, the preamplifier is a self-resetting, switched-capacitor amplifier with deliberately oriented enclosed terminal switches. Assuming an amplifier gain A and input-referred offset $V_{\rm OFF}$, and assuming capacitor matching ($C_{1+}=C_{1-}=C_1$ and $C_{2+}=C_{2-}=C_2$), the preamplifier output $V_{\rm pre}$ at the end of $\phi_{\rm 1D}$ is:

$$V_{\rm pre} = -A\left(V_{\rm in} - \frac{C_1 - C_2}{C_1 + C_2}V_{\rm REF}\right) - \frac{A}{1+A}V_{\rm OFF} \tag{5.11}$$

The value $[(C_1-C_2)/(C_1+C_2)]V_{\rm REF}$ determines when $V_{\rm pre}=0$ V and sets the comparator transition point. The table of capacitor values is given in Table 5.3.[27] Simulations show the gain A to be about 4.8.

The dynamic latch is shown in Figure 5.34. It is based upon $V_{\rm pre+}$ and $V_{\rm pre-}$ creating unequal resistances in M_1 and M_2 and favoring a node in the race of V_1 and V_2 from their precharged states, similar to the principle used by *Cho and Gray* [1995]. After buffering, one dynamic latch output is fed to the static latch. Shown in Figure 5.35, the static latch is based on the dual interlocked storage cell (DICE), first presented by *Calin et al.* [1996]. This cell stores the desired value twice—in uninverted form at V_1 and V_3, and in inverted form at V_2 and V_4—and uses dual feedback networks so that, should one storage node be upset, the two adjacent storage nodes act to restore it. Hence single-bit upsets to the static latch are but temporary, reducing not only the upset rate of the pipeline signal path but also the probability of railing in the residue amplifiers due to upset DAC codes.

As seen in the timing diagram of Figure 5.32(b), pulse waveforms are used to time the dynamic and static latches. Figure 5.32(c) shows the circuit for generating these pulses from $\phi_{\rm 1Y}$, a delayed version of ϕ_1. The circuit of Figure 5.32(c) is

[27] The indicated ratios are implemented by dividing a master interdigital capacitor of 64 fingers by the fractions of Table 5.3. For example, for the capacitors of the −21/32 transition level of the calibrated stages, the master capacitor is divided into capacitors of 11 and 53 fingers.

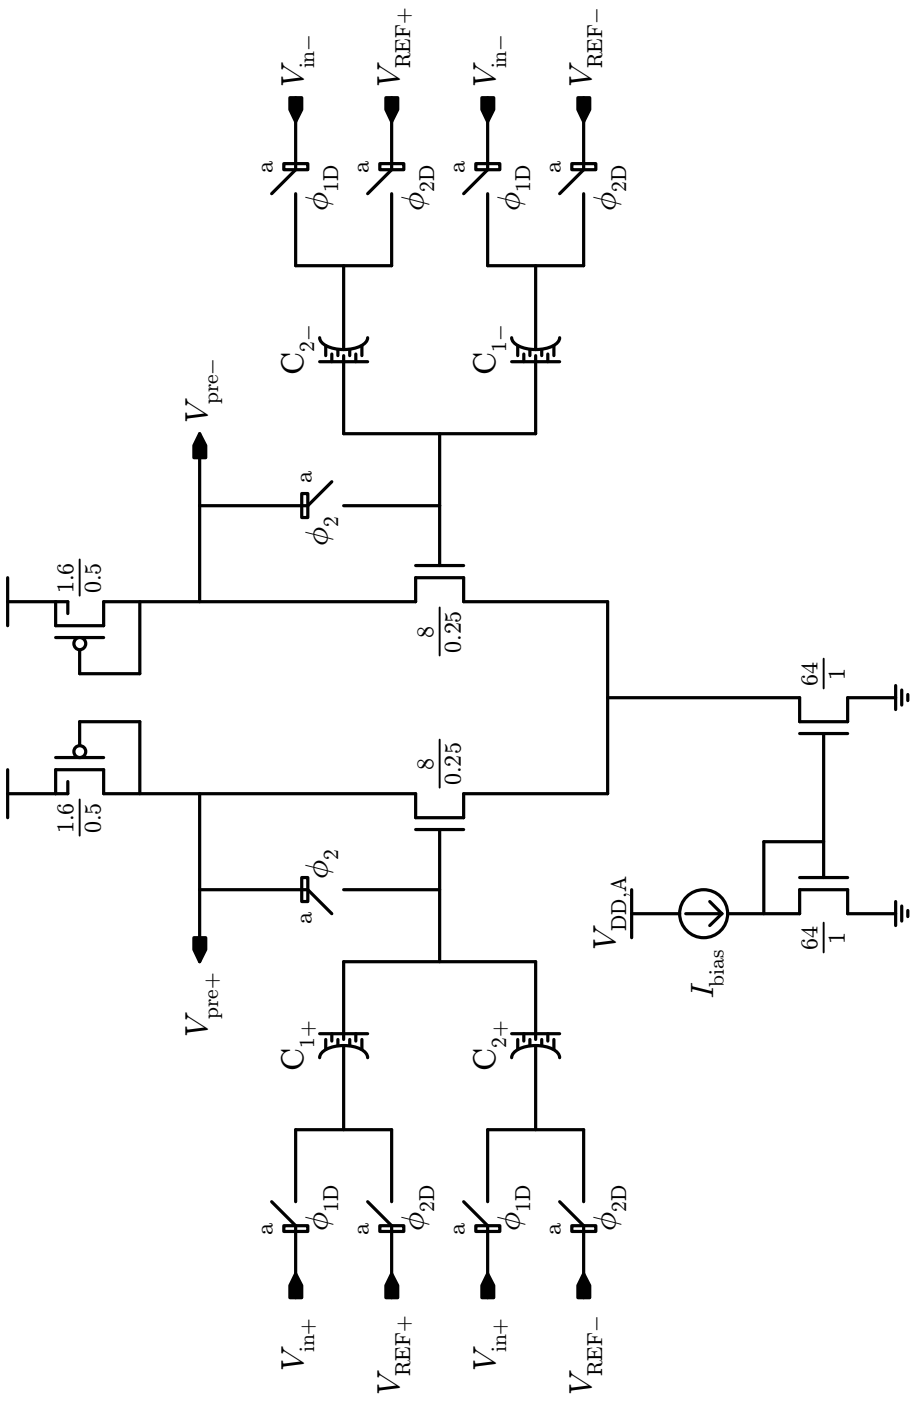

Figure 5.33: Comparator preamplifier. $I_{\text{bias}} = 46.875\ \mu\text{A}$. All supplies tied to $V_{\text{DD,D}}$ except where explicitly noted otherwise or for switches (which are tied to $V_{\text{DD,CLK}}$). All device sizings in μm.

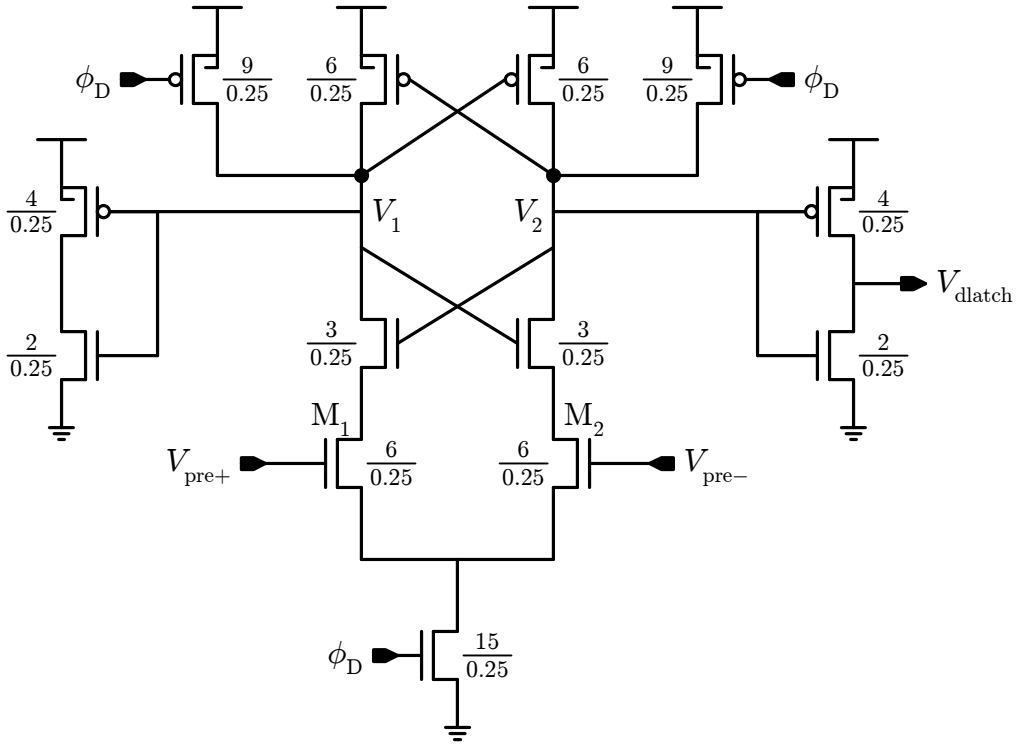

Figure 5.34: Comparator dynamic latch. All supplies tied to $V_{DD,D}$, all device sizings in μm.

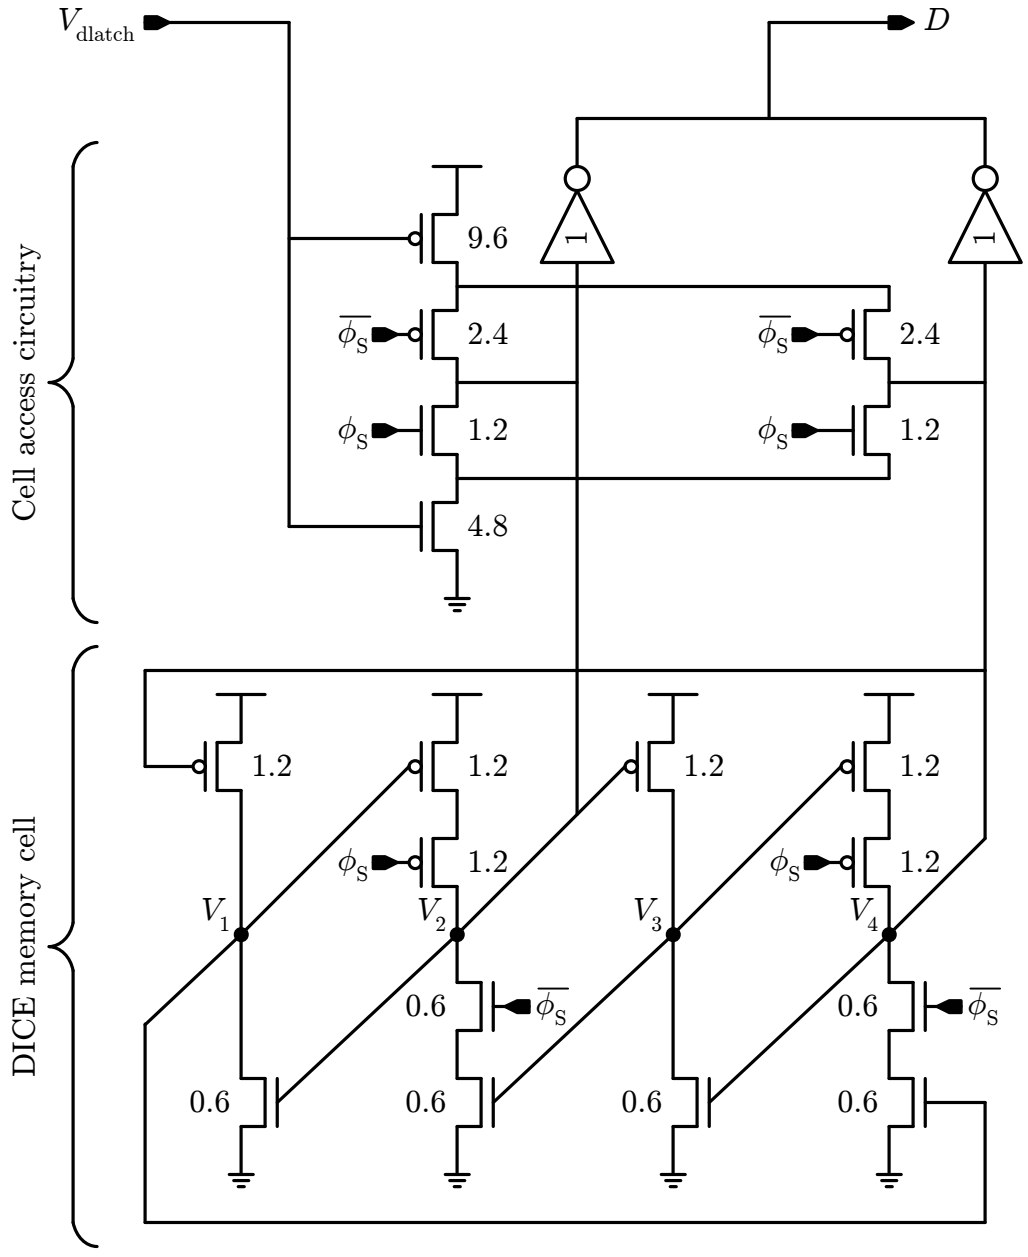

Figure 5.35: Comparator static latch. Only device widths, in μm, given: all devices of 0.24-μm length. All PMOS bulks tied to supply, all supplies to $V_{\text{DD,D}}$.

Stage	1, 2, 3		4, 5		6	
Number of comparators	8		6		15	
	C_1/C_2	$T[k]$	C_1/C_2	$T[k]$	C_1/C_2	$T[k]$
	11/53	$-21/32$	12/52	$-5/8$	4/60	$-7/8$
	17/47	$-15/32$	20/44	$-3/8$	8/56	$-6/8$
	23/41	$-9/32$	28/36	$-1/8$	12/52	$-5/8$
	29/35	$-3/32$	36/28	1/8	16/48	$-4/8$
	35/29	3/32	44/20	3/8	20/44	$-3/8$
	⋮	⋮	52/12	5/8	24/40	$-2/8$
	53/11	21/32			28/36	$-1/8$
					32/32	0
					36/28	1/8
					⋮	⋮
					60/4	7/8

Table 5.3: Comparator preamplifier capacitor values. Capacitor ratio C_1/C_2 and consequent $T[k]$ (as a fraction of V_{REF}) given; for all comparators, $C_1 + C_2 = 132$ fF.

included locally at each comparator to prevent dephasing of the pulses and reduce inter-comparator interference.

Simulations show that the total input-referred offset of the comparator is on the order of 10 mV (standard deviation). This estimate incorporates capacitor and transistor mismatch in both the preamplifier and dynamic latch. While it ignores the offset incurred by early firing of the dynamic latch (i.e., that incurred by the dynamic latch firing before V_{in}—and hence V_{pre}—have completely settled), nonetheless the total offset should be well within the 156.25 mV and 125 mV allowed under digital redundancy for the 3.1-bit and 2.8-bit stages, respectively.

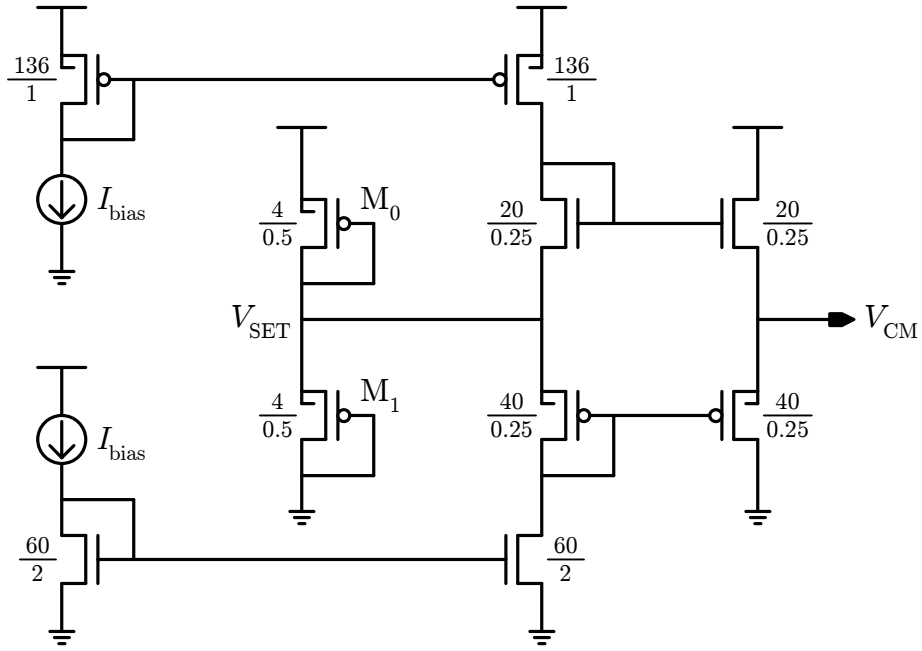

Figure 5.36: V_{CM} generator. All supplies tied to $V_{\text{DD,A}}$, all device sizings in μm.

5.5.3 Common Mode Generation

The common mode voltage V_{CM} is generated locally in each stage. The circuit is shown in Figure 5.36. V_{CM} is set to $V_{\text{DD,A}}/2$ by the matched, series PMOS devices M_0 and M_1. PMOS devices are used as their bulks can be source-tied, eliminating body effect differences between M_0 and M_1. V_{SET} then establishes the bias point of NMOS and PMOS mirrors, which feed an output stage. Given the orientation of NMOS mirror over PMOS mirror, the class-AB output stage nominally maintains low power consumption, but can rapidly source or sink transient currents much greater than the bias current when necessary, an especially beneficial characteristic when feeding switched-capacitor circuitry.

5.5.4 Digital Logic

All digital logic is implemented by a custom standard cell library built from scratch for the SVADC-1. In particular, in each logic cell the NMOS and PMOS components are enclosed by a diffusion guard ring for latchup robustness. Furthermore, this guard

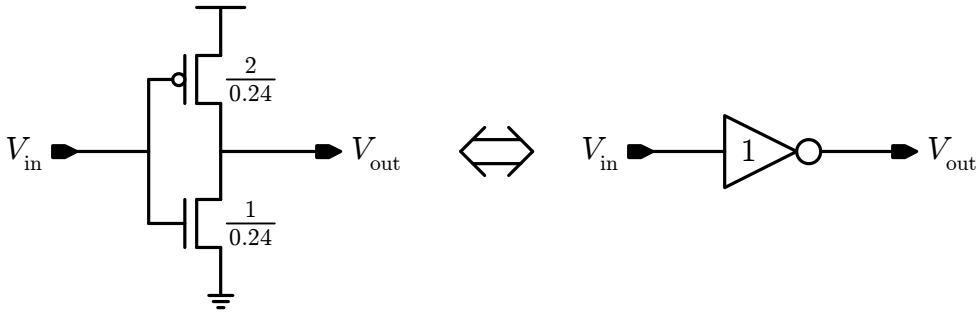

Figure 5.37: A 1X inverter: circuit (left) and corresponding symbol (right). All device sizings in μm.

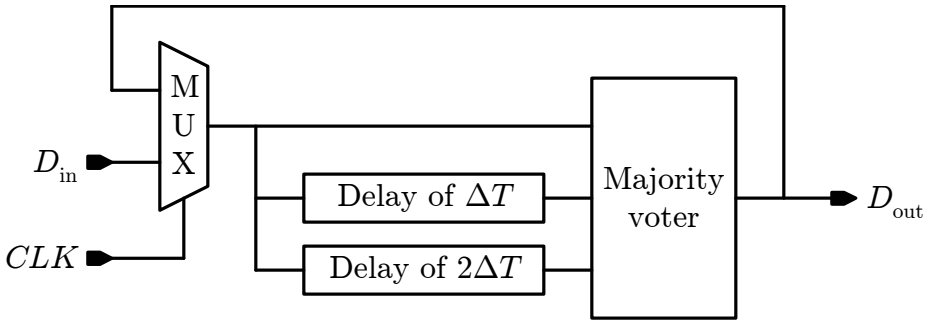

Figure 5.38: Temporal-sampling latch architecture. Multiplexer denoted as MUX. All elements implemented by active CMOS logic.

ring is metallized on three sides: the remaining side—that between the NMOS and PMOS sections—is left unmetallized to ease routing. An example is shown in Figure 5.3. As noted in Section 5.2.1, such extensive guard ringing does incur a significant area (and by extension, power) penalty.

For combinational logic, active CMOS logic is used instead of pass gate or dynamic logic. Combinational logic is based on a unit 1X inverter, the symbol and schematic of which is given in Figure 5.37. Other logic cells—including NAND, AND-OR-INVERT, and INVERT-OR-AND cells among others—are sized to achieve the same drive as this unit inverter. All are then scaled to drive strengths of 2X, 4X, and beyond as needed: in schematics, the drive strength is indicated by the number inside the gate symbol.

For sequential logic, the SVADC-1 adopts an upset tolerant flip flop proposed by *Mavis and Eaton* [2002]. This flip flop uses a temporal-sampling latch, illustrated in Figure 5.38. It can be shown that the temporal-sampling latch is immune to upsets of temporal widths as great as ΔT both in its internal circuitry and on its clock input.[28,29] Two such latches in series—with the second multiplexed by the inverse of CLK—create a flip flop. For the SVADC-1, this inversion is implemented in the multiplexer itself as shown in Figure 5.39 and the latches oriented to generate rising-edge triggered flip flops.[30] Notably, such flip flops consume significantly more area than traditional flip flops. In the SVADC-1, though, this area penalty is acceptable as the flip flops consume but a small portion of the overall converter area.[31]

Enclosed terminal layout is not used in the core digital standard cell library. Thus under radiation testing these cells do demonstrate increased I_{LEAK} with dose: while this increase is not sufficient to compromise performance, it does increase power consumption (see Section 6.3.1.3). Future revisions should consider redrawing the standard cell library to support enclosed terminal layout.

5.5.5 Pads and I/O

All pads of the SVADC-1 are protected by electrostatic-discharge (ESD) protection cells provided by National Semiconductor. These cells include ESD protection diodes to both the supply and ground. For pads processing analog signals, the ESD cells are

[28] An upset in the multiplexer or delay elements does not propagate past the majority voter so long as it manifests on at most one majority voter input at any time. In general, this condition holds so long as the upset width is less than ΔT. (For example, consider an upset on the multiplexer output: while each of the majority voter inputs eventually manifests the upset, so long as the upset width is less than ΔT only one of the inputs is affected at a time.) Of course, an upset in the majority voter output causes an upset at the latch output. However, again so long as the upset width is less than ΔT, such an upset is ultimately rejected as the majority voter inputs drive its output back to the proper value.

[29] An upset on the clock input temporally switches the input of the multiplexer. Thus, it is essentially the same as an upset in the multiplexer: by the reasoning of footnote 28 of this chapter, so long as the upset is of temporal width less than ΔT, it is ultimately rejected by the majority voter.

[30] In particular, a latch with the multiplexer implementation of Figure 5.39(a) is followed by a latch with the multiplexer implementation of Figure 5.39(b).

[31] However, in larger digital designs—for example, an implementation of the digital portion of the SVADC-1—the large flip flop area may be of greater concern.

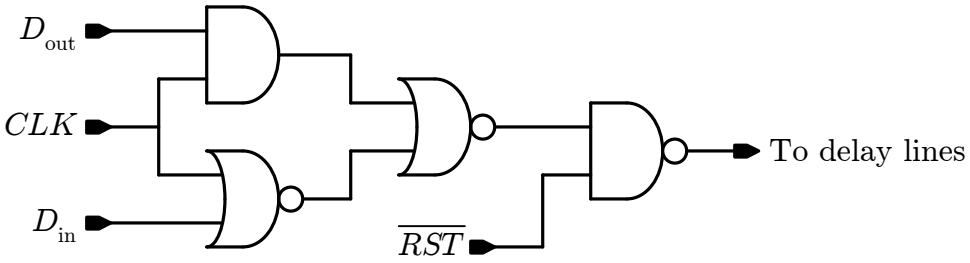

(a) Clock-low transparent multiplexer.

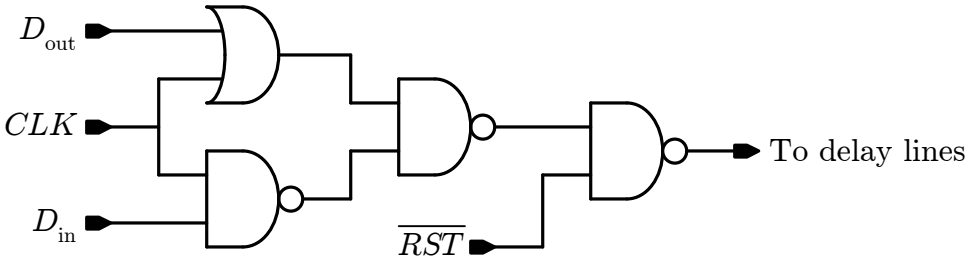

(b) Clock-high transparent multiplexer.

Figure 5.39: Implementation of multiplexer in temporal-sampling latch of Figure 5.38. (a) Multiplexer output is $\overline{D_{in}}$ when clock is low (output is D_{out} when clock is high). (b) Multiplexer output is $\overline{D_{in}}$ when clock is high (output is D_{out} when clock is low). Active-low reset signal $\overline{RST}$ incorporated into both multiplexers.

supplied by $V_{DD,A}$, while for pads processing digital signals, an independent $V_{DD,IO}$ supply is used. All ESD cells include an NMOS clamping circuit to protect against sudden rises in the supply. The clamping circuit is triggered by an RC circuit, one of which is included in each of the analog and digital pad rings.[32] To improve reliability, HV versions of the ESD cells, rated for up to 3.3 V signaling, are used. Due to standard electrical latchup concerns these cells are heavily guard-ringed: no additional

[32]The clamp circuit essentially follows the same principles as the clamp described by *Ker* [1999]. Briefly, when an input pin is driven sufficiently high—such as by an ESD event—it can overwhelm the ESD protection diode and increase V_{DD}. Unchecked, this supply rise can lead to sudden current flow in the I/O buffer or even core circuitry, causing transistor damage (especially in the case of core circuitry). To prevent such damage, the ESD cells include a large NMOS clamping device that shorts the ESD supplies V_{DD} and V_{SS}, shunting the ESD current before it propagates deeper into the chip. To identify rapid ESD-based rises in V_{DD} (versus slower rises due to chip power-up, for example), the NMOS clamp is driven by an RC trigger circuit: when the RC circuit cannot react quickly enough, the clamp is activated. Hence at least one trigger per ESD V_{DD}-V_{SS} pair is needed.

latchup protection is added.

The digital pads also synchronously buffer I/O signals. Specifically, digital signals input to the chip are buffered, flopped (on the rising edge of ϕ_{IO} of Figure 5.13), and then buffered again for propagation to core. The same structure, reversed in direction, is also used for signals output from the chip. The buffer directly contacting the pad is supplied by $V_{DD,IO}$ and constructed from HV devices. As the oxides of such devices are generally thicker, the NMOS devices of the buffer are drawn in enclosed terminal fashion. The remaining I/O logic is supplied by $V_{DD,D}$, constructed from LV devices, and drawn in standard fashion. All pad flops are triggered in phase with the rising edge of the input clock ϕ_{IO} to provide a uniform digital interface.[33] As the converter does not produce digital outputs at high speed, CMOS inverters are used for the buffers (both input and output).[34]

5.6 The SVADC-1 Chip

A photograph of the fabricated SVADC-1 chip is shown in Figure 5.40. The chip is fabricated in a SiGe BiCMOS manufacturing process (BiCMOS8iED) generously provided by National Semiconductor Corporation under a collaboration agreement with Stanford University. While fabricated in a BiCMOS process, only the CMOS layers are used in the SVADC-1: the process is thus equivalent to a single-well, 1P5M, 0.25-μm CMOS process fabricated in bulk Si on a non-epitaxial substrate with interdevice separation achieved by STI oxides. In total, the chip includes about 28,000 transistors and 800 interdigital capacitors.[35]

The SVADC-1 chip includes both the converter proper and an assembly of test circuits. The converter proper includes all converter stages and the reference (both voltage and current) generation circuitry (labeled "analog bias" in Figure 5.40).[36]

[33] The exception to the rule is the reset signal for the flops: being asynchronous, it is only buffered and not flopped.

[34] For faster digital I/O signaling, open drain logic—such as that used by *Thompson* [2001, pp. 88–89]—can be used instead.

[35] For matching reasons, many interdigital capacitors are implemented as several individual interdigital capacitors wired in parallel, yielding the elevated count.

[36] In addition to the circuitry of Figures 5.26 and 5.27, analog bias also includes an extra V_{CM}

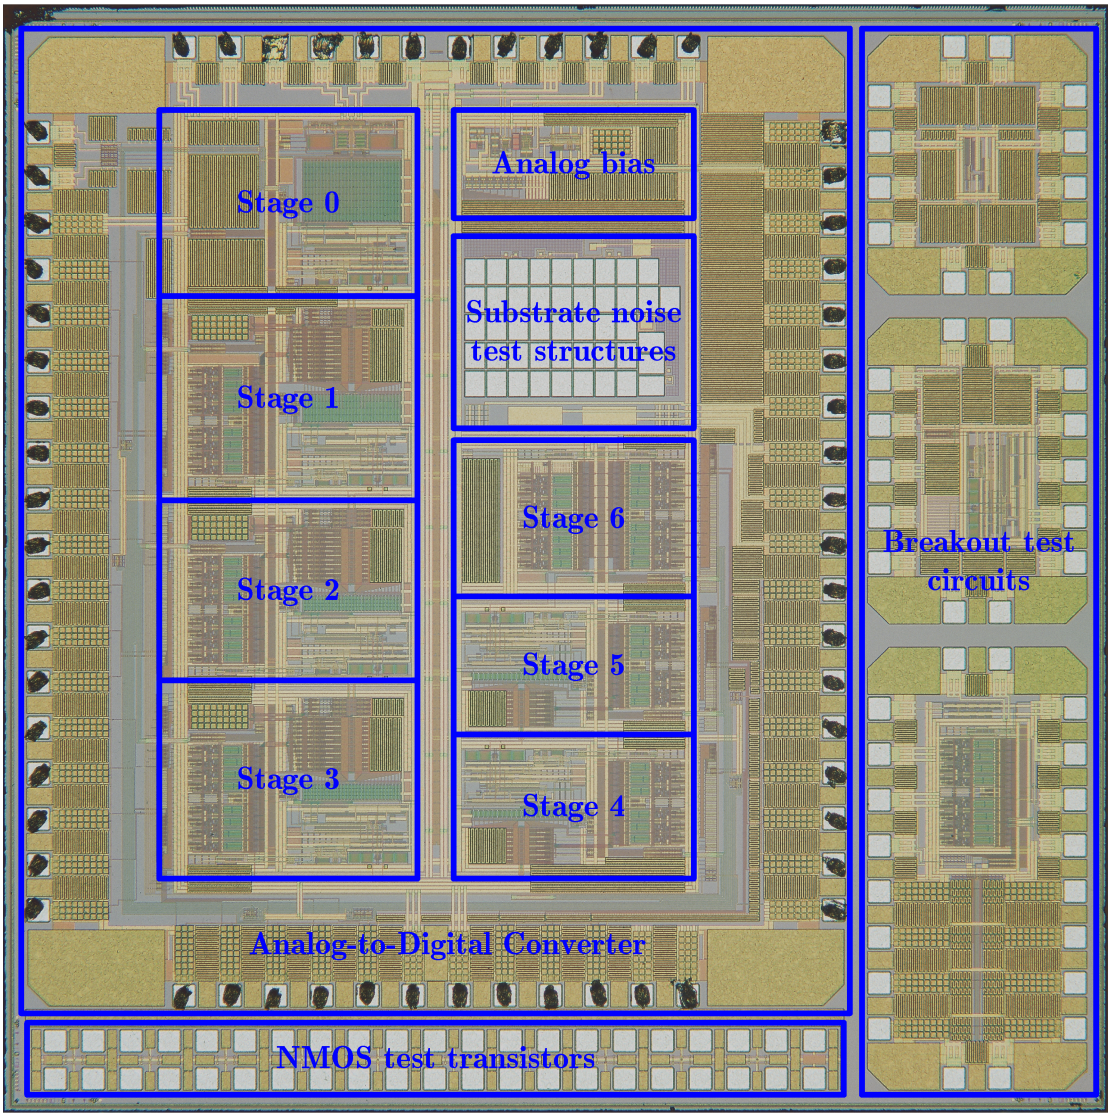

Figure 5.40: Chip micrograph of the SVADC-1. Total silicon dimensions: 3.5 mm by 3.5 mm.

Note that the converter stages are laid out in a "J", incurring a large "bend" in the signal path between stages 3 and 4. If desired, future layouts can opt to straighten the converter into a straight line (a configuration often easier to integrate with other designs, for example). Care is taken to place the track-and-hold stage (i.e., stage 0) as close as possible to the differential analog input and clock pads in the upper left portion of the chip to minimize parasitics on these sensitive lines. In addition, these signals are shielded by ground pads.

The test circuitry on the SVADC-1 encompasses three parts:

1. *NMOS test transistors*

   A series of NMOS devices intended for device-level radiation characterization. These test devices are composed of five pairs of NMOS devices, each pair the same device laid out once in standard and once in enclosed terminal fashion. Each NMOS device is independent, with its own gate, source, drain, and bulk pads. All the transistor-level radiation data from BiCMOS8iED (see Sections 2.3.3 and 5.2.2) is culled from ^{60}Co total-dose testing of these devices: see Appendix F for further details.

2. *Breakout test circuits*

   A set of breakout circuits—each with its own pad ring and supplies—intended for circuit-level radiation characterization. These circuits are, from top to bottom: a biasing circuit, an operational amplifier, and a 3.1-bit (i.e., 8-comparator) flash ADC. Each is taken directly from circuitry in the converter proper, and is made fully independent with its own pad ring and supplies. Given the strong radiation performance of the converter proper, these circuits—while confirmed operational—were not rigorously tested, radiation or otherwise. Details of their design, though, are included for completeness in Appendix J.

3. *Substrate noise test structures*

   Intended for use with a probe station, this array of pads leads to test structures that enable extraction of parameters for substrate noise modeling. Designed by

generator as per Figure 5.36. The V_{CM} of this extra circuit is driven directly to an output pad to serve as a simple check that the analog circuitry is "alive".

Hai Lan and Jae Wook Kim (from Professor Robert Dutton's group) at Stanford University, unfortunately time and resource constraints precluded their use.[37]

Finally, small top-metal pads (albeit passivated) are also included at the output of each stage of the converter proper: if necessary, it is possible to access these pads with a probe station to assess internal signal conditions.

The complete SVADC-1 chip occupies 12.25 mm^2 of total area. However, the converter proper only occupies 9.0 mm^2 of this area, and the converter core (i.e., excluding pads and I/O circuitry) only occupies 5.2 mm^2.

5.7 Conclusion

This chapter discussed the circuit implementation of the SVADC-1 in a commercial 0.25-μm CMOS process. In addition to incorporating the novel DAC-differencing self-calibration technique of Chapter 4, the SVADC-1 also adopts a philosophy of radiation-hardness by design to ensure performance under radiation. While specific techniques—such as guard rings to prevent latchup and enclosed terminal devices to prevent radiation-induced I_{LEAK}—are identified, it should be clear that the philosophy of radiation-hardness by design is not limited to these particular techniques. Rather, radiation concerns informed—and in many cases drove—the design of much of the underlying circuitry, from the architecture of the residue amplifiers to the transconductance of the operational amplifiers. The ultimate test of the efficacy of the design techniques described in this chapter is the measured performance of the SVADC-1, which is the subject of the next chapter.

[37] Their presence should be noted, and use considered, in any future attempts at integration of the SVADC-1 with the LNA/AAF design of Benjamin J. Mossawir [*Mossawir et al.*, 2006; *Mossawir*, in preparation].

Chapter 6

Measured Results

This chapter presents the measured performance of the SVADC-1, including performance both before radiation (referred to as baseline performance) and after radiation (referred to as radiation performance). For reference, recall that the target specifications of the SVADC-1 are given in Table 3.1.

The chapter begins by describing the experimental setup used to acquire the data presented here. It then presents the measured baseline performance and demonstrates the efficacy of the self-calibration technique for linearizing the conversion and achieving high SFDR. Measured radiation performance—including both total-dose (via high-energy proton irradiation) and single-event (via heavy-ion irradiation and pulsed laser) results—follows, and shows the effectiveness of the radiation-hardness-by-design approach employed throughout the design. Finally, the SVADC-1 has been incorporated into a number of satellite instrument designs, and results from those efforts—in particular, the resource usage of a field-programmable gate array (FPGA) implementation of the digital portion, and outcomes from burn-in testing—are presented.

Additional details of many of the testing experiments presented here can be found in the appendices. In particular, Appendix A gives a general introduction to radiation testing, while Appendix B (for the total-dose testing) and Appendices C, D, and E (for the single-event testing) address the specifics of particular radiation testing experiments. Appendix K discusses the burn-in testing.

6.1 Experimental Setup

The experimental setup used to assess the SVADC-1 is shown in Figure 6.1, where the SVADC-1 is referred to as the device-under-test (DUT). The DUT is mounted on a custom test board (designed and fabricated by the author) that provides interface circuitry between it and and the test equipment. This experimental setup is used for baseline DUT performance characterization and, with small amendments, for DUT radiation performance characterization as well.

6.1.1 Device Under Test

The SVADC-1 chip is manufactured in the BiCMOS8iED manufacturing process generously provided by National Semiconductor Corporation under a collaborative agreement with Stanford University. For testing, the chip is housed in a 84-pin, J-lead, ceramic, quad flat-pack package. The chip is affixed to the metallized package cavity with conductive epoxy; indeed, the package cavity is used as a local ground plane. The package comes with a lid to cover the cavity: this lid is left removable to allow direct exposure of the chip to radiation sources.

As the DUT is removed and replaced several times during radiation testing, zero-insertion force (ZIF) socketing is important for reliability. The test board uses a clam shell socket. A large hole in the center of the clam shell lid allows exposure of the chip even when the shell is closed.[1]

6.1.2 Input Signal Generation

To assess performance, sinusoids of varying frequencies and amplitudes are input to the DUT. For the SVADC-1, sinusoid generation is challenging since the sinusoids should have at least 90-dB SFDR, and ideally much higher. To accomplish this linearity while maintaining flexibility in varying the sinusoid amplitude and frequency,

[1] As the clam shell chosen is intended to socket a slightly thicker package, a shim of 50-mil Nylon 6/6 plastic is placed atop the DUT package to increase its height. Naturally, the center of the shim is also open so that the chip may be exposed. Given its material, radiation backscatter from the shim should be negligible.

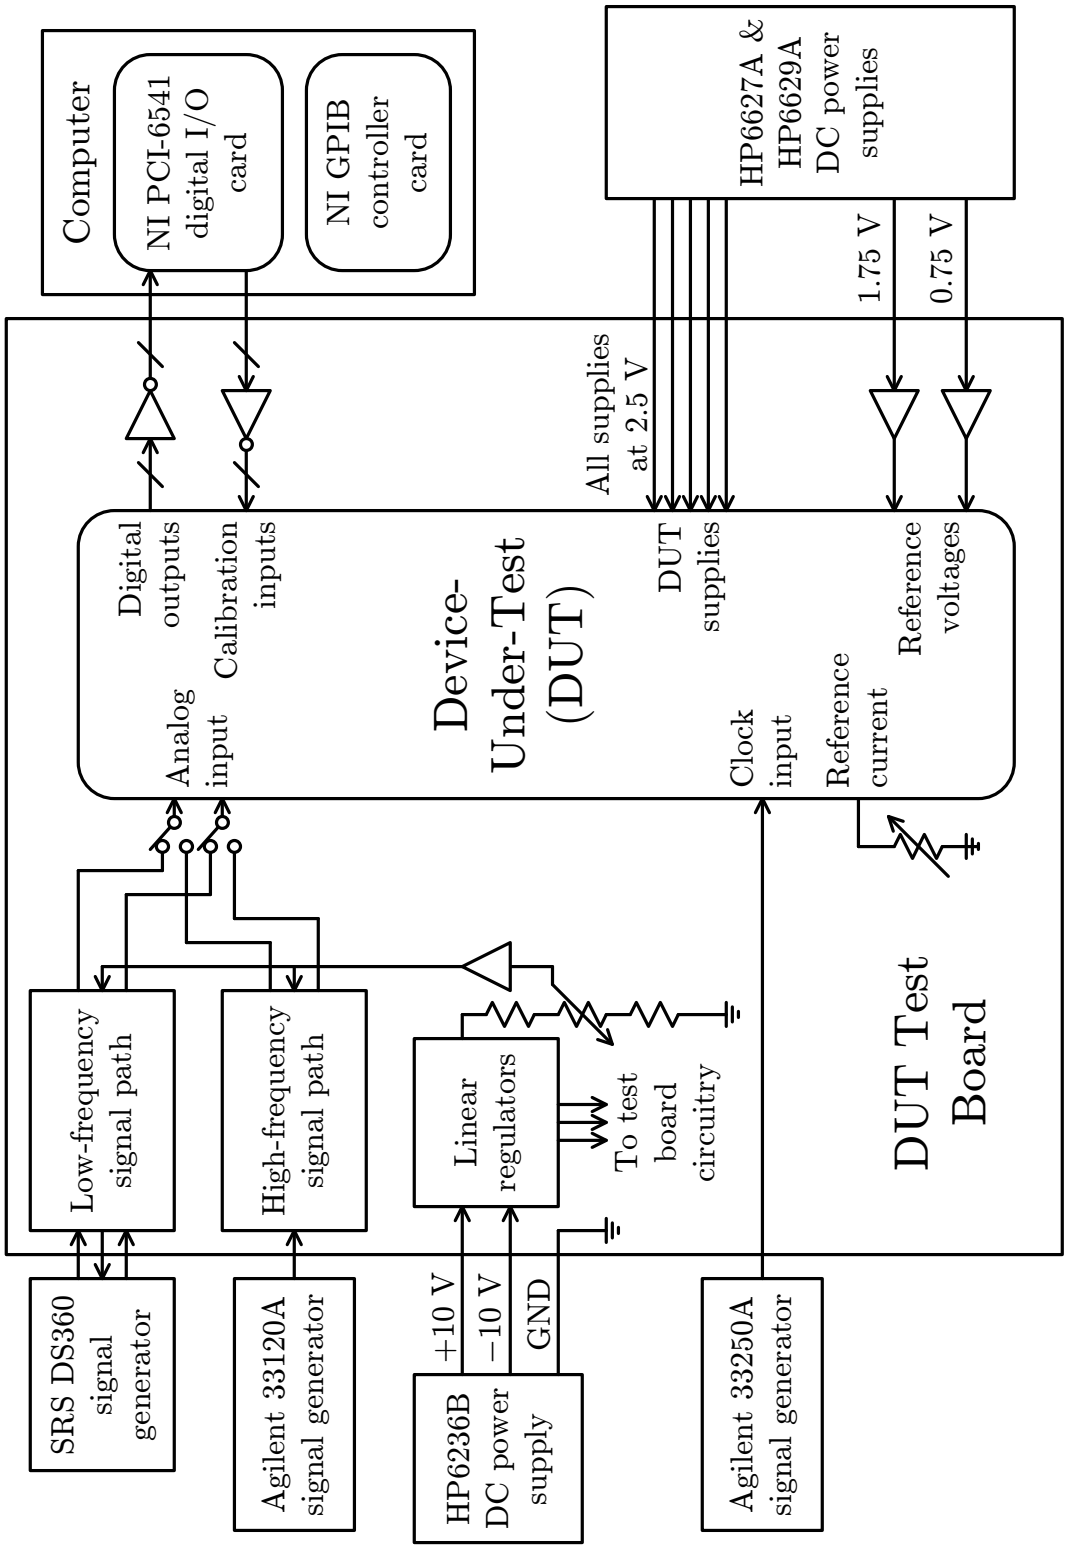

Figure 6.1: Experimental setup for assessing SVADC-1 performance. All test equipment shown (except HP6236B) controlled by an NI GPIB controller card (GPIB connections omitted for clarity).

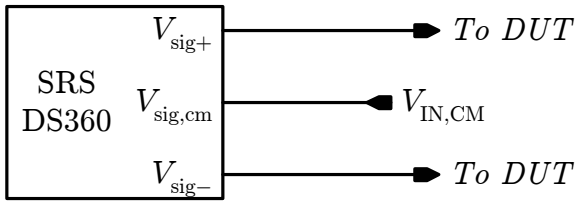

(a) Direct drive configuration.

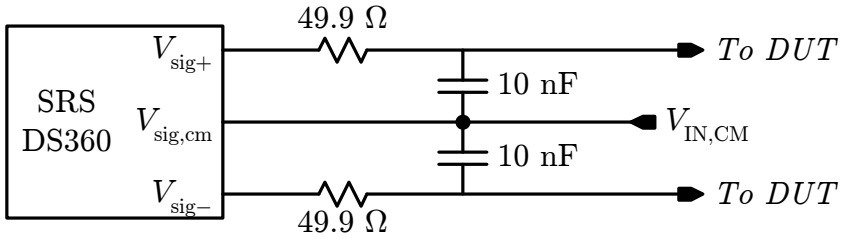

(b) Configuration with RC filter.

Figure 6.2: Low-frequency input signal path, including (a) direct drive configuration and (b) RC-filtered configuration as used in harmonic and nonharmonic performance assessment, respectively. Capacitors implemented by polystyrene capacitors for better linearity. See Figure 6.4 for $V_{IN,CM}$ generation.

two separate signal paths—corresponding to low-frequency sinusoid generation and high-frequency sinusoid generation—are used. For low-frequency signals (<40 kHz), the DUT differential input is provided by a Stanford Research Systems (SRS) DS360 signal generator, while for high-frequency signals (>100 kHz), the differential input is provided by an Agilent 33120A.

Signal conditioning on both paths ensures the required linearity. The low-frequency path is shown in Figure 6.2. While the SRS DS360 generates fully differential sinusoids of the required linearity up to 40 kHz, it demonstrates increased nonharmonic noise in the hundreds of kiloHertz range and above. An on-board differential RC filter quells this noise. However, though polystyrene capacitors are used for their better linearity [*Rabii*, 1998, p. 163], the filter nonetheless incurs harmonic distortion measurable by the SVADC-1. Hence, the filter is excluded and the configuration of Figure 6.2(a) adopted to assess the harmonic performance of the DUT, whereas the filter is included and the configuration of Figure 6.2(b) adopted

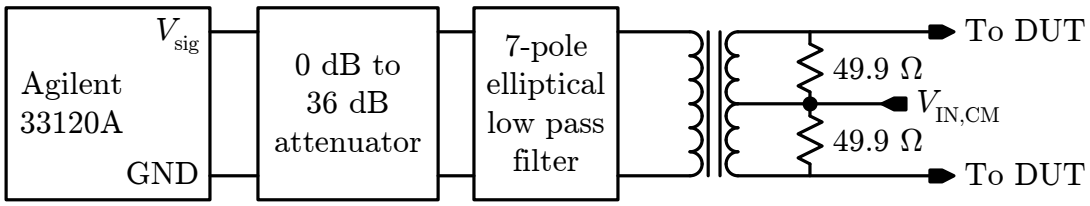

Figure 6.3: High-frequency input signal path. Attenuation implemented by cascade of Mini-Circuits HAT-6, Mini-Circuits HAT-10, and Mini-Circuits HAT-20 attenuators that provide 6-dB, 10-dB, and 20-dB attenuation, respectively. Low-pass filter implemented by Coilcraft P7LP-604L. Transformer implemented by Coilcraft AS8456-A. See Figure 6.4 for $V_{\mathrm{IN,CM}}$ generation.

to assess the nonharmonic performance of the DUT.

The high-frequency path is shown in Figure 6.3. The Agilent 33120A single-ended output is converted to differential via a transformer. The common mode is set by driving the secondary center tap. To reduce the harmonic distortion of the Agilent 33120A output, a 7-pole elliptical low-pass filter suppresses second and higher harmonics. The filter is terminated on the transformer secondary. In contrast to the low-frequency path which can inject a wide range of frequencies, the filter confines the high-frequency input signal to a narrow frequency: 366.007 kHz is chosen.[2] Also in contrast to the low frequency path, the Agilent 33120A cannot create as small signals[3] and hence attenuators are added on the primary side to test the DUT performance at lower input powers. Up to 36-dB attenuation can be added. For each used attenuation setting, the DUT input signal amplitude is independently measured so that results under different attenuations may be properly combined.

In both the low-frequency and high-frequency cases the differential signal common mode is set to 1.25 V by an on-board reference generated from a resistor ladder buffered by a unity gain operational amplifier circuit, as shown in Figure 6.4.

[2]The low-pass filter thus provides attenuations of 28 dB, 46 dB, 69 dB, and 66 dB for the 2nd, 3rd, 4th, and 5th harmonics of the high-frequency input signal, respectively [*Coilcraft, Incorporated*, P7LP-604L].

[3]The Agilent 33120A can create signals as small as 50 $\mathrm{mV_{PP}}$ (into a 50 Ω load), whereas the SRS DS360 can create signals as small as 20 $\mu\mathrm{V_{PP}}$ (into a high impedance, balanced load).

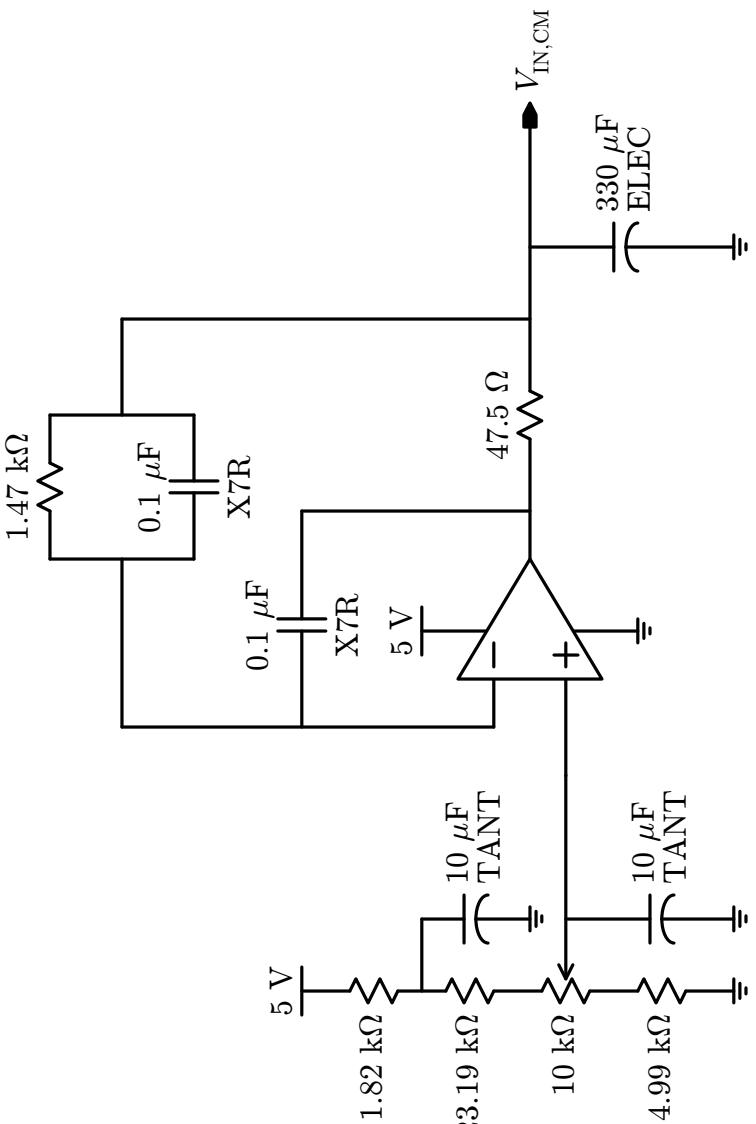

Figure 6.4: Input signal common mode voltage $V_{\mathrm{IN,CM}}$ generator. Operational amplifier implemented by Maxim MAX4252. TANT refers to tantalum capacitor, ELEC to electrolytic capacitor.

6.1.3 Clock Signal Generation

A 5 MHz clock is provided by an Agilent 33250A. To minimize unwanted termination effects, a sinusoidal clock is used. As the original chip is designed for a square wave external clock, the sinusoid offset and amplitude need to be trimmed for performance. This trimming is parasitic-dependent: depopulating and repopulating the SVADC-1 alters the optimal offset and amplitude by a few tens of milliVolts. Nonetheless, this trimming technique is sufficient to demonstrate the performance of the SVADC-1.[4] Finally, note that the clock and input signal generators are not synchronized to each other: the resulting frequency smearing is handled by subsequent signal processing (see Section 6.2.1.1).

6.1.4 Supplies and References

An HP6627A/HP6629A DC supply bank provides all five DUT supplies ($V_{\mathrm{DD,A}}$, $V_{\mathrm{DD,D}}$, $V_{\mathrm{DD,CLK}}$, $V_{\mathrm{DD,IO}}$, and $V_{\mathrm{DD,REF}}$). The HP6627A and HP6629A also measure the average current of each supply. Multiple decoupling capacitors—both ceramic and tantalum—are included near each DUT pin.

The DUT $V_{\mathrm{REF+}}$ and $V_{\mathrm{REF-}}$ reference voltages of 1.75 V and 0.75 V, respectively, are also set by the HP6627A/HP6628A supply bank. However, these voltages are subsequently buffered on board by unity gain operational amplifier circuits as shown

[4]Future chip revisions may wish to consider a more robust on-chip clock receiver design. Commercial parts often circumvent such clock transmission issues by adopting on-chip clocking by internal oscillators. However, a free-running internal sample clock is not recommended for the SVADC-1. The time of the sampling instant is important in many scientific analyses that rely on phase measurements. These investigations often compare the measured phase between different receiver channels, or between different receivers. Naturally, changing the sampling instant alters the phase shifts. While it is possible to correct these phase shifts if the difference in the sampling instants between different ADCs is known, conducting such in-flight calibration is better avoided if possible (note that ground-calibration is insufficient here as ADCs may power cycle during flight). Hence in many receivers the ADC sampling clocks are synchronized to a GPS clock signal, often the edge of the 1 pulse-per-second (PPS) signal as this edge can be guaranteed to an accuracy of several nanoseconds. For example, a commercial, terrestrial GPS receiver from i-Lotus Corporation achieves 1-PPS edge accuracies of <2 ns in a 1-sigma, and <6 ns in a 6-sigma, sense [*i-Lotus Corporation*, M12M] (accuracy measured by deviation from Coordinated Universal Time (UTC)). Ideally, then, the SVADC-1 should contain an on-chip PLL that synchronizes with an external clock to generate the internal sampling edge.

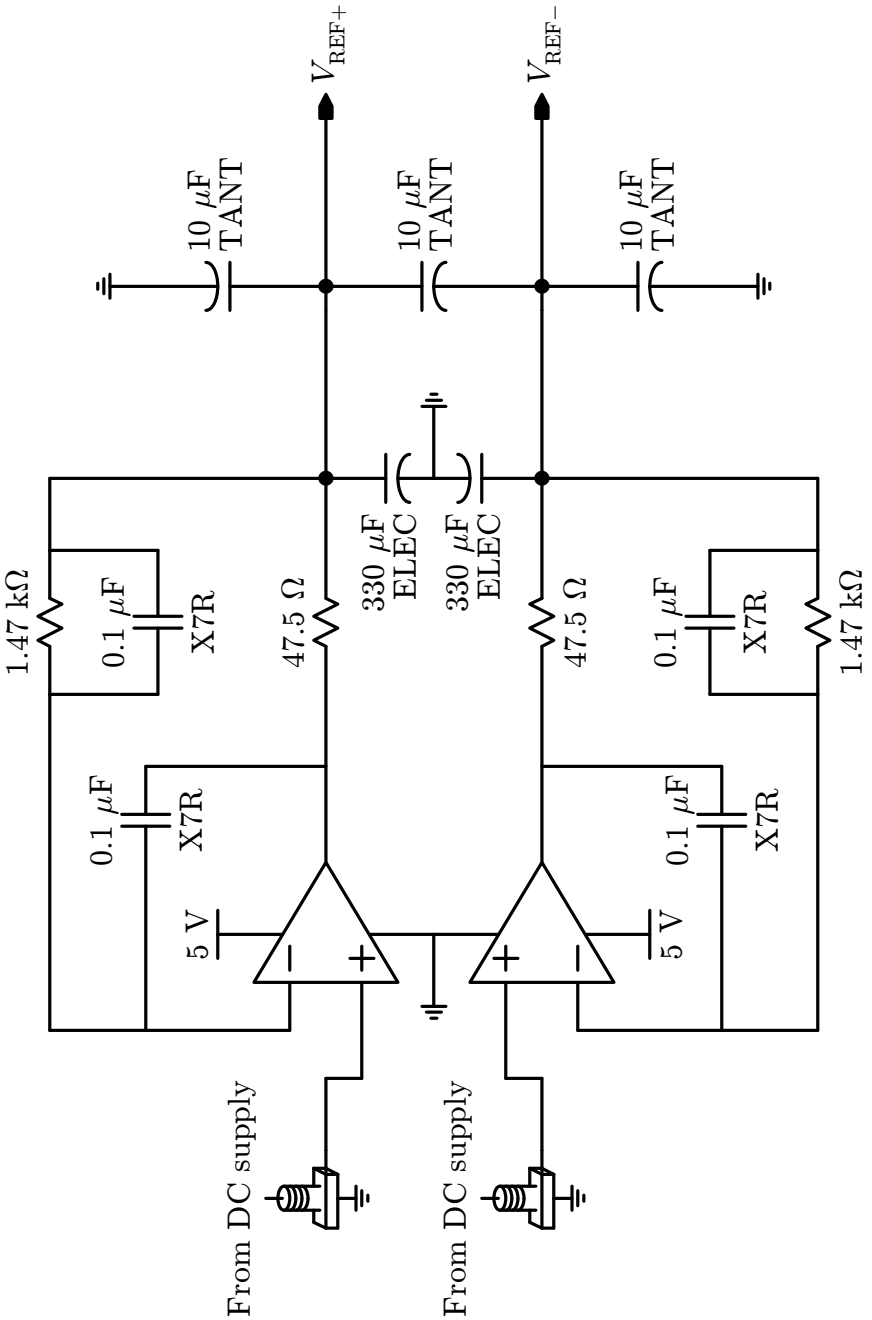

Figure 6.5: Reference voltages $V_{\mathrm{REF+}}$ and $V_{\mathrm{REF-}}$ generator. Operational amplifiers implemented by a single Maxim MAX4252 dual operational amplifier part. TANT refers to tantalum capacitor, ELEC to electrolytic capacitor.

in Figure 6.5. A single MAX4252 dual operational amplifier chip buffers both $V_{\mathrm{REF}+}$ and $V_{\mathrm{REF}-}$ to reduce differential noise between the two references. The buffered voltages are decoupled with large 10 μF tantalum capacitors (both between each other and to ground) near to the DUT reference voltage pins.

As shown in Figure 6.1, the DUT reference current is provided externally by an adjustable resistor to ground and the current trimmed to 62.5 μA (see Figure 5.27 for more detail).

All on-board circuits—such as operational amplifiers and resistor ladders—are powered separately from the DUT by $\pm$10 V supplies from an HP6236B power supply. The $\pm$10 V are linearly regulated on board to appropriate voltages.

6.1.5 Digital I/O and Control

All digital signals input to and output from the DUT are buffered by on-board CMOS inverters. A 22-Ω resistor along each driven line reduces ringing in the CMOS-to-CMOS signaling. The inverters are operated at 2.5 V.

A National Instruments (NI) PCI-6541 digital I/O card mounted in a dedicated computer handles the digital bus, including both control of the SVADC-1 calibration inputs and acquisition of the SVADC-1 digital output from each stage. Acquired data is buffered on-card, then written to hard drive.[5] In addition, the same computer includes an NI GPIB controller card to control all the test equipment shown in Figure 6.1 (except the HP6236B, which has no GPIB interface).

A custom C program (written by the author) manages both the NI PCI-6541 digital I/O card and the NI GPIB controller card. Thus this program directs, and ultimately automates, all testing: the latter is especially important in time-constrained radiation testing.[6]

[5]Continuous acquisition length is limited by the on-card memory of 8 Mbit per channel. Note that, at a SVADC-1 clock frequency of 5 MS/s, the memory is sufficient for gathering at least 1 second of continuous data, which in turn enables construction of 1-Hz/bin spectra, important for assessing SVADC-1 performance as detailed in Section 6.2.1.1.

[6]The time-limited nature of radiation testing is described in more detail in Appendix A. In particular, for total-dose testing, characterization of the DUT must occur within a 2-hour time window following irradiation to minimize annealing effects. For single-event testing, access to the radiation source itself is typically heavily time-constrained.

6.2 Baseline Performance

This section presents the baseline performance characterization of the SVADC-1, including both spectral and transition level performance as assessed by the metrics of Sections 3.2.2 and 3.2.1, respectively.

Performance is presented both before and after self-calibration. For the before self-calibration results, the digital reconstruction defaults all parameters to their values assuming ideal transfer functions and interprets the pipeline at 12 bits (i.e., the 2 LSBs of the terminating ADC—that is, of stage 6—are ignored). For the after self-calibration results, the SVADC-1 is self-calibrated as per the procedure of Section 4.4.2. During self-calibration the SVADC-1 pipeline operates with 2 extra bits (that is, the full 4-bit resolution of the terminating ADC is used) and the backend ADC output is averaged over 2048 samples for all stages. After self-calibration, the SVADC-1 pipeline is again interpreted at 12 bits (i.e., the 2 LSBs of the terminating ADC are again ignored). Notably, the self-calibration algorithm creates lookup table entries wider than the backend ADC bitwidth: for the results presented in this chapter, lookup table entries and all arithmetic calculations are performed in 32-bit double precision floating point.[7] The final output word of the digital reconstruction, though, is rounded to a 16-bit unsigned integer, chosen since 16 bits is a common bitwidth used in many digital computers and digital signal processors (DSPs). Notably, though the output word is 16 bits, only a small subset of the codes are active (exactly 2^{12} before self-calibration, and $\sim 2^{12}$ after).[8]

6.2.1 Spectral Performance

The SVADC-1 spectral performance is measured by testing with single-frequency input sinusoids. As per Table 3.1, spectral metrics are assessed over an evaluation bandwidth of 100 Hz to 1 MHz. An exception is made for the high-frequency,

[7] Naturally, the lookup table entries and arithmetic computations can alternately be performed in fixed-point precision. Indeed, if implemented in custom digital logic—such as in the FPGA implementation described in Section 6.4.1—the digital reconstruction most probably adopts fixed-point arithmetic, with bitwidths throughout the digital reconstruction "pipeline" scaling so as to maintain sufficient precision while minimizing hardware cost and power.

[8] In the terminology of Section 3.1.2, $M_{\text{out}} > M_{\text{min}}$.

366.007-kHz input: for this case, the evaluation bandwidth is increased to 100 Hz to 1.2 MHz to include the 3rd harmonic.

6.2.1.1 Procedure

Spectral metrics are computed by constructing output spectra $Y'[k]$ of the SVADC-1 output sequence $y'[n]$. In particular, spectra of both 100-Hz/bin and 1-Hz/bin spectral resolution are constructed as per:

- *100-Hz/bin spectral resolution*
  Construction of the 100-Hz/bin spectrum—corresponding to a 50,000-point FFT at 5 MS/s—is shown in Figure 6.6(a). Several steps are taken to reduce frequency smearing. First, while the DC component is removed directly, since the DC and 100-Hz bins are adjacent at this spectral resolution, spectral components near DC can readily smear into the evaluation bandwidth. To remove these components, the signal is high-pass filtered by a 1,000,000-tap finite impulse response (FIR) filter with a 10 Hz cutoff. Second, the signal is windowed by a 4-term cosine window with continuous third derivative of the form [*Nuttall*, 1981]:
  $$w[n] = \frac{1}{N}\sum_{k=0}^{3}\alpha_k \cos\left[\frac{2\pi kn}{N}\right] \tag{6.1}$$
  where [*Nuttall*, 1981]:
  $$\alpha_0 = 0.338946\ ,\ \alpha_1 = -0.481973\ ,\ \alpha_2 = 0.161054\ ,\ \alpha_3 = -0.018027 \tag{6.2}$$
  The window length N is set equal to the FFT length. A 50,000-point FFT on the windowed signal generates the final 100-Hz/bin spectrum.

- *1-Hz/bin spectral resolution*
  Construction of the 1-Hz/bin spectrum—corresponding to a 5,000,000-point FFT at 5 MS/s—is shown in Figure 6.6(b). Windowing—in this case by a 5,000,000-point Chebyshev window with 130-dB sidelobe attenuation [*Harris*, 1978]—is employed to reduce frequency smearing. Notably, given the fine

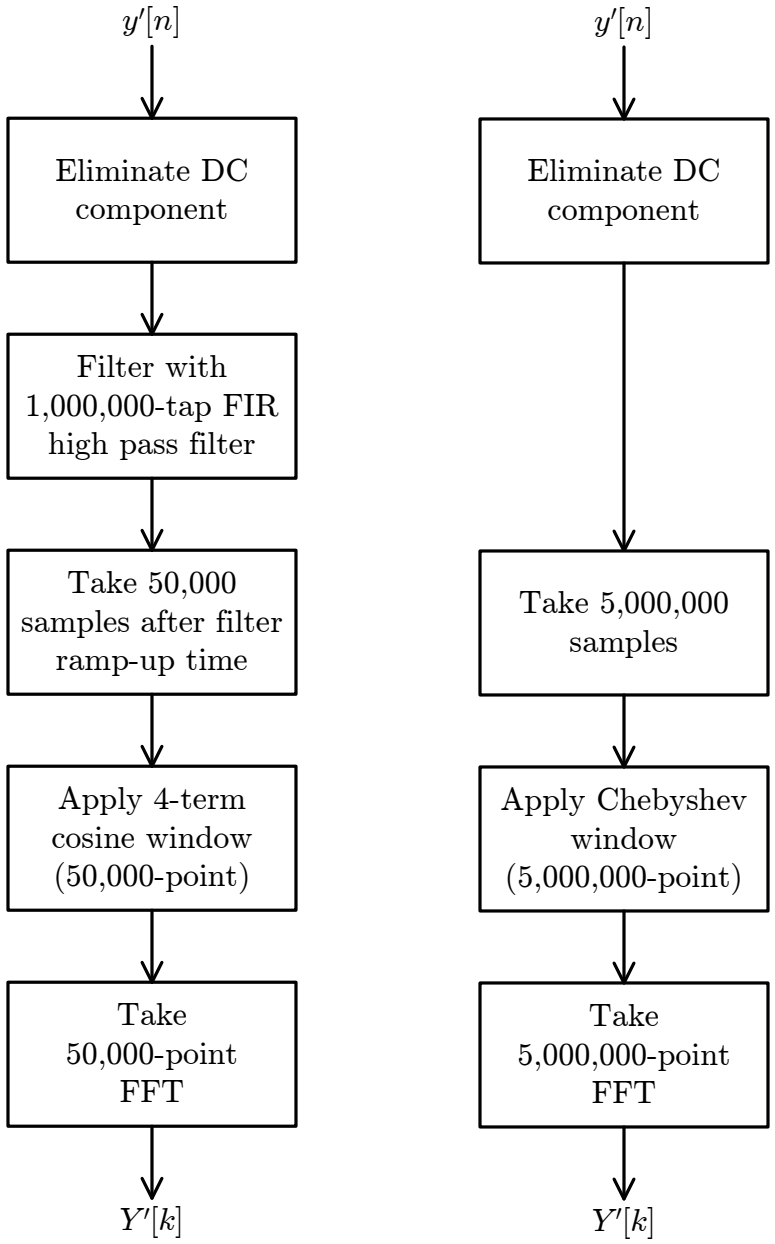

(a) 100-Hz/bin spectrum.

(b) 1-Hz/bin spectrum.

Figure 6.6: Construction of the output spectra $Y'[k]$ from the SVADC-1 output sequence $y'[n]$ at spectral resolutions of (a) 100 Hz/bin and (b) 1 Hz/bin. Recall that $y'[n]$ is given at a sampling rate of 5 MS/s.

spectral resolution, frequency smearing of near DC frequency components should not significantly leak into the evaluation bandwidth and hence no high-pass filtering is needed. A 5,000,000-point FFT on the windowed signal generates the final 1-Hz/bin spectrum.

Spectra of both 100-Hz/bin and 1-Hz/bin spectral resolutions are used to compute the SFDR and SNDR. Specifically:

- *SFDR computation*
  The SFDR is computed as per Equation (3.31), which decomposes SFDR into harmonic (hSFDR) and nonharmonic (nhSFDR) components. The ultimate SFDR is then the minimum of the hSFDR and nhSFDR at each input power.

- *hSFDR computation*
  For the hSFDR, the harmonic components are assessed at a 1-Hz/bin spectral resolution: especially after self-calibration, at 100-Hz/bin the harmonics are small enough that the noise floor interferes with accurate harmonic power estimation.[9] All harmonics of the fundamental over the entire evaluation bandwidth are considered.[10] Furthermore, for low-frequency inputs, the harmonic input signal path of Figure 6.2(a) is used.

- *nhSFDR computation*
  For the nhSFDR, which is FFT length dependent (see Section 3.3.3), 100-Hz/bin spectra are used. To isolate the nonharmonic portion, islands of the bandwidth of the window main lobe centered at the fundamental and at each harmonic frequency are removed. The largest value of the remaining spectrum is then determined. Since the noise floor part of the nonharmonic portion is a

[9] More technically: after windowing, the vast majority of the power of coherent signals is confined to a peak of a few bins centered around the coherent signal frequency. To compute the power of the coherent signal, the powers over the entire peak are summed. So long as the peak is well-defined above the noise floor, such a calculation is a sufficiently accurate approximation of the ultimate coherent signal power. Unfortunately, at 100 Hz/bin the noise floor is sufficiently large that it interferes with this calculation.

[10] Thus for a 9.9909-kHz low-frequency input, say, 100 harmonics are considered. Note that for the 366.007-kHz high-frequency input, the evaluation bandwidth is expanded to 100 Hz to 1.2 MHz so that 3 harmonics can be considered.

probabilistic quantity, the nhSFDR is computed over 16 separate spectra and the result averaged.[11] Furthermore, for low-frequency inputs, the nonharmonic input signal path of Figure 6.2(b) is used.

- *SNDR computation*
  The SNDR requires computing the total non-fundamental power. Hence for this metric the 100-Hz/bin spectra are used directly. Again, as it involves a probabilistic quantity, the SNDR is computed over 16 separate spectra and the resulting SNDRs averaged to derive the final SNDR at each input signal power.[12] Furthermore, for low-frequency inputs, the nonharmonic input signal path of Figure 6.2(b) is used.

6.2.1.2 Results

Sample spectra of the SVADC-1 output while processing a near full-scale 366.007-kHz high-frequency input are shown in Figures 6.7 and 6.8 at spectral resolutions of 100 Hz/bin and 1 Hz/bin, respectively. The spectra are presented in units of dBFS, or decibels-relative-to-full-scale, and both spectra before (top) and after (bottom) self-calibration are included for comparison. Two observations can be drawn. First, the noise floor is essentially flat with frequency. This characteristic is especially apparent in Figure 6.8, whose logarithmic frequency scale highlights the low frequency components. Clearly, any low-frequency flicker noise is greatly suppressed, confirming the efficacy of the residue amplifier offset cancellation scheme. Second, the success of the self-calibration technique is clear: whereas the uncalibrated spectra show strong distortion, after self-calibration this distortion is dramatically reduced, resulting in much quieter spectra and higher SFDR.

The SFDR improvement is formalized in Figure 6.9, which shows the measured SFDR as a function of input signal power both before and after self-calibration. A low-frequency input case (9.99093-kHz) is shown at left, and the high-frequency

[11] Notably, it is the nhSFDR values which are averaged, and not the underlying spectra themselves: the latter would cause averaging of the noise floor, reducing it and artificially elevating the nhSFDR.

[12] Again, note that it is the SNDR values which are averaged, and not the underlying spectra.

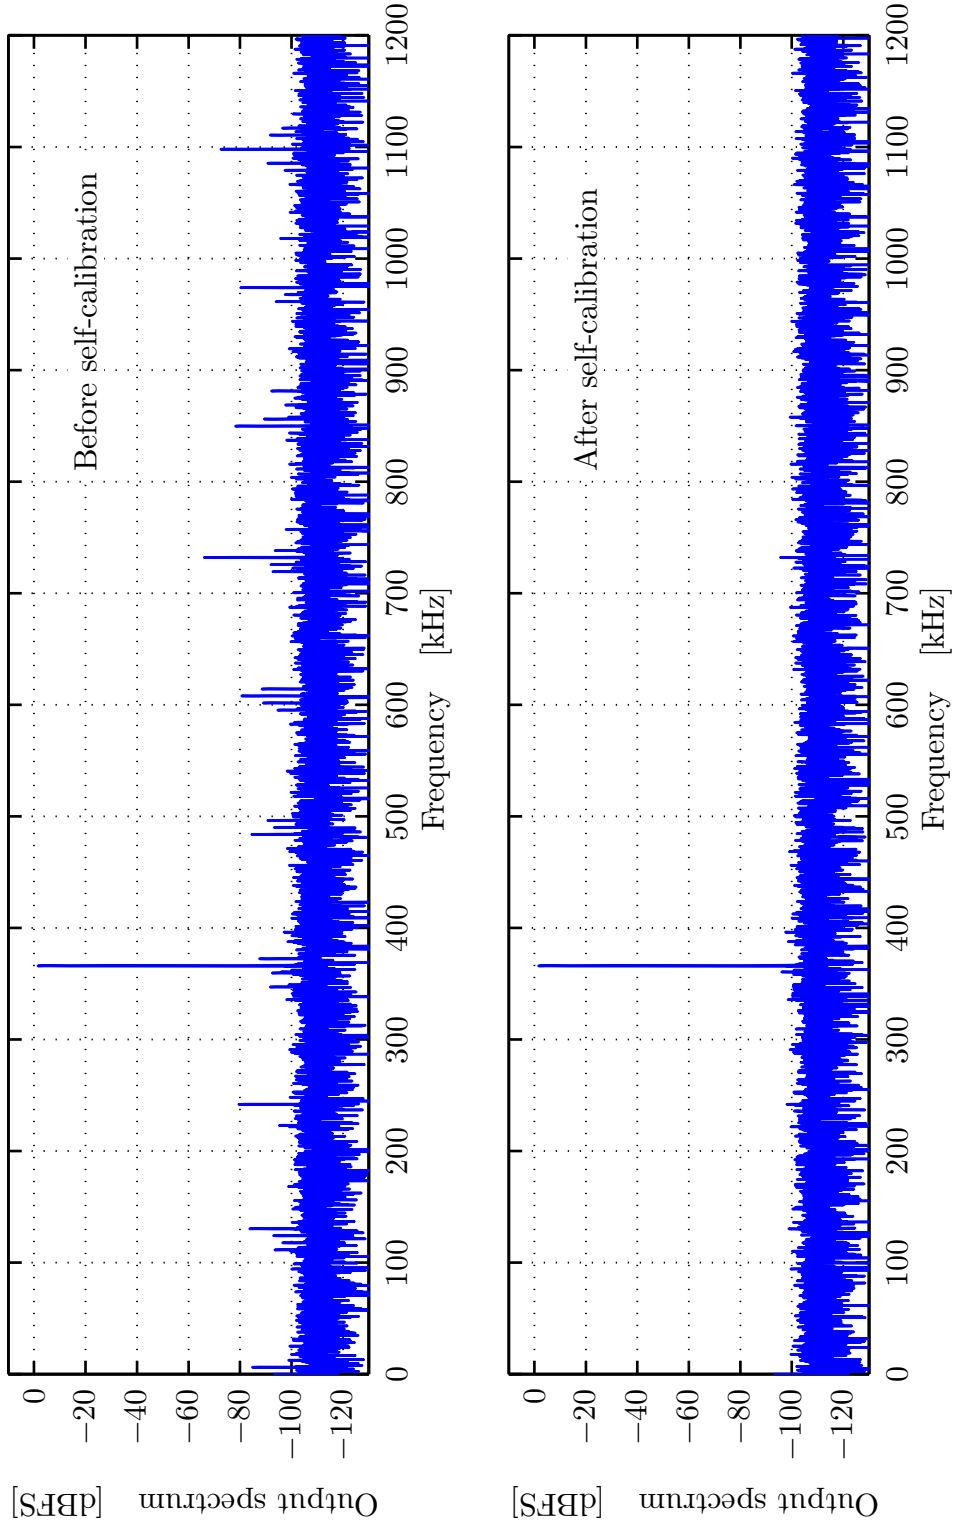

Figure 6.7: Measured 100-Hz/bin SVADC-1 output spectrum in response to a 366.007-kHz, high-frequency input sinusoid at −1.20 dBFS input power. Before self-calibration (top) and after self-calibration (bottom) spectra shown. Only frequency components up to 1.2 MHz (upper limit of evaluation bandwidth for this high-frequency signal) plotted. Spectra compiled as per Figure 6.6(a).

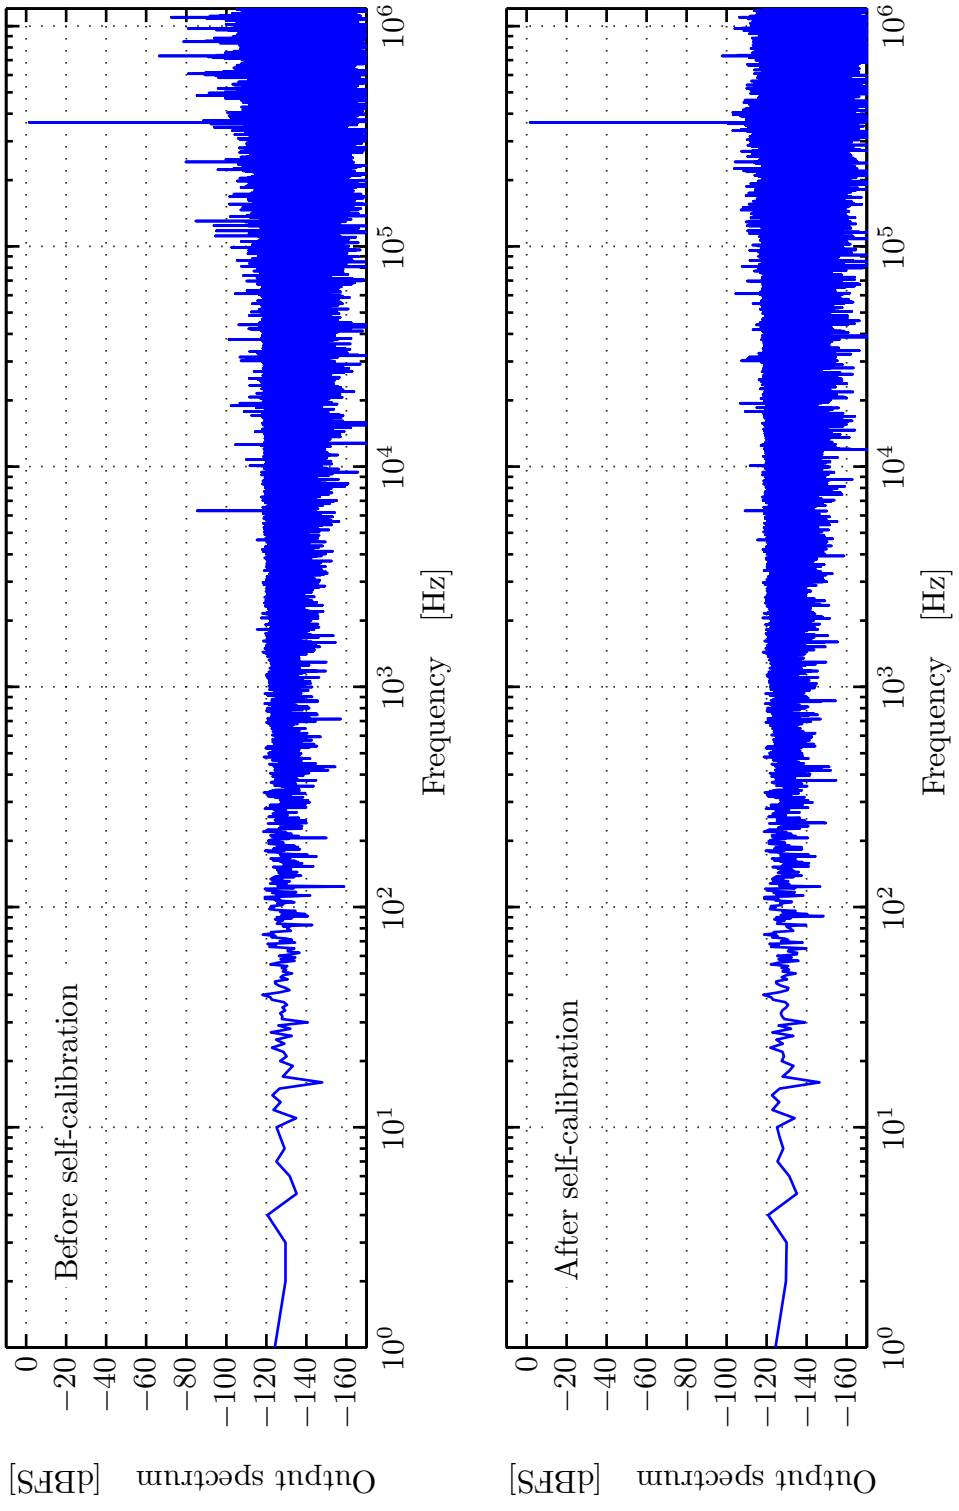

Figure 6.8: Measured 1-Hz/bin SVADC-1 output spectrum in response to same signal as in Figure 6.7. Before self-calibration (top) and after self-calibration (bottom) spectra shown. Spectrum plotted against logarithmic frequency to better highlight low-frequency performance. Spectra compiled as per Figure 6.6(b): note that (except for the DC bin), this signal processing preserves low-frequency performance undisturbed.

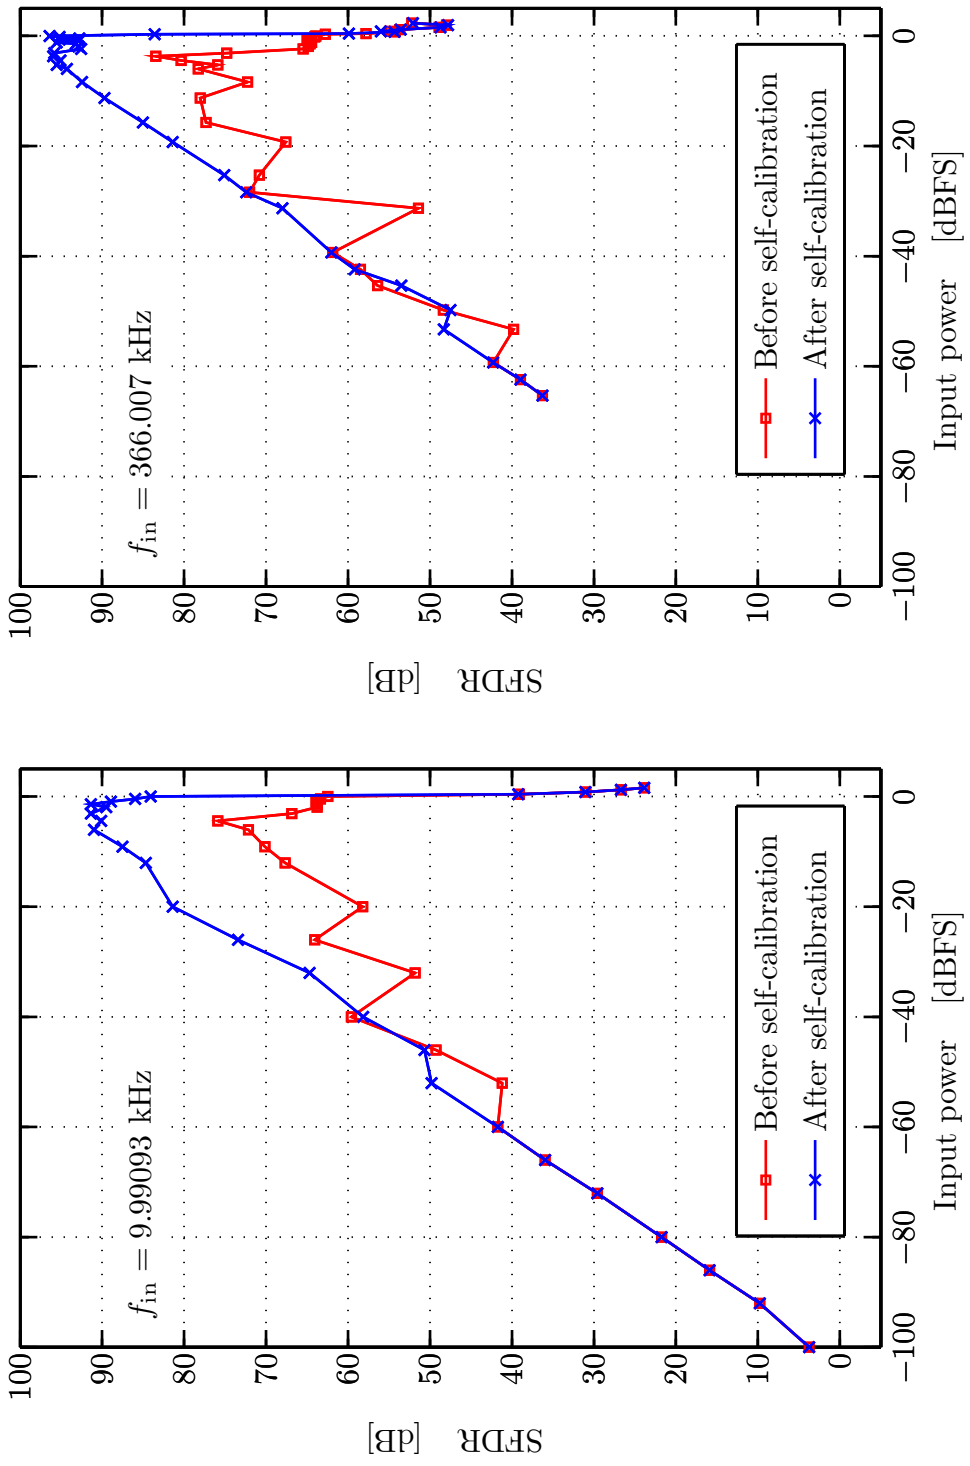

Figure 6.9: Measured SFDR performance of the SVADC-1 with both low (left) and high (right) frequency inputs. Performance both before (in red) and after (in blue) self-calibration shown. Markers indicate measured data points.

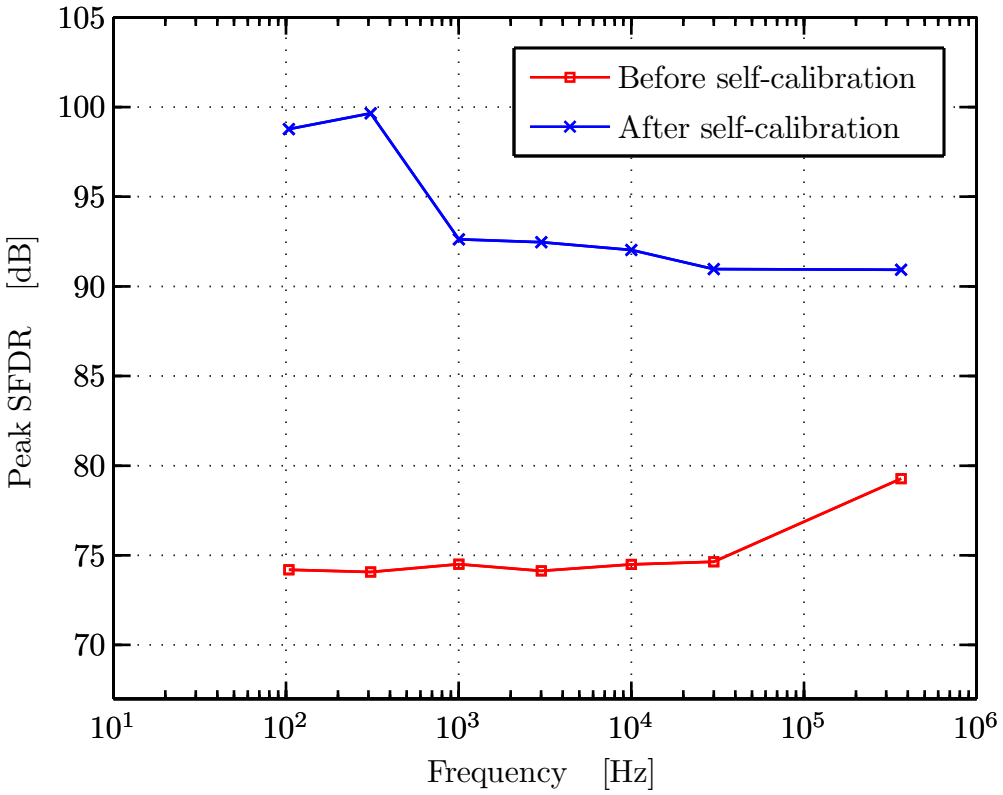

Figure 6.10: Measured peak SFDR performance of the SVADC-1 versus input frequency. Performance both before (in red) and after (in blue) self-calibration shown. Same calibration values used at all frequencies. For 103.712-Hz and 308.752-Hz inputs, peak hSFDR substituted for peak SFDR. All data except highest frequency input acquired with low-frequency input path; highest frequency acquired with high-frequency input path. Markers indicate measured data points.

input case (366.007-kHz) is shown at right. In both cases, correction of analog-digital mismatch through self-calibration greatly reduces distortion and significantly increases the SFDR: self-calibration achieves an improvement in the peak SFDR by 15.5 dB from 75.9 dB to 91.4 dB in the low-frequency case, and by 13.0 dB from 83.4 dB to 96.4 dB in the high frequency case.

The wideband efficacy of the self-calibration technique is further confirmed in Figure 6.10. Here, one set of calibration values is used to self-calibrate the SVADC-1. The SVADC-1 is then assessed for peak SFDR over a broad range of input frequencies. (Notably, for the 103.712-Hz and 308.752-Hz input frequencies, the peak hSFDR is

substituted for the peak SFDR due to unavoidable frequency smearing in assessing the nhSFDR of the 100-Hz/bin spectra.) Clearly, the self-calibration technique provides wideband SFDR improvement, producing at least 90.9 dB peak SFDR over all the tested input frequencies.

Self-calibration also improves the SNDR as shown in Figure 6.11, which plots the measured SNDR in the same fashion as in Figure 6.9. For both input frequencies, at lower input powers the SNDR is dominated by the noise floor and the SNDRs before and after self-calibration coincide. However, at large input powers harmonic distortion is no longer negligible and the before self-calibration SNDR is limited to peak values of 64.8 dB in the low-frequency, and 66.6 dB in the high-frequency, cases. Self-calibration eliminates much of this distortion, though, allowing the SVADC-1 to achieve a peak SNDR of 69.9 dB in the low-frequency, and 70.8 dB in the high-frequency, case after self-calibration. Additionally, the SNDR improvement holds over a wide range of input frequencies: similar to Figure 6.10, in Figure 6.12 one set of calibration values is used and the SVADC-1 peak SNDR assessed at multiple input frequencies.[13] As can be seen, after self-calibration the SVADC-1 maintains a wideband peak SNDR of at least 69.1 dB.

Finally, to compare the SVADC-1 performance with the SNR specification of Table 3.1, the 1-Hz/bin spectrum of the 366.007-kHz high-frequency input is assessed over the full bandwidth of 0 Hz to 2.5 MHz. The resulting peak SNR after self-calibration is 65.6 dB, within the specified range and close to the design value of 66 dB (see Section 4.4.1).

6.2.2 Transition Level Performance

The SVADC-1 transition level performance is assessed by measuring the INL of the SVADC-1. There are multiple methods for computing INL. This dissertation uses the sinusoidal technique, where an input-overranging, low-frequency sinusoid

[13]Notably, for the 103.712-Hz and 308.752-Hz input frequencies, to ensure sufficient bandwidth for assessing the peak of the fundamental, the 1-Hz/bin spectra are used instead of the 100-Hz/bin, and no averaging is performed. Furthermore, for the 103.712-Hz input, the evaluation bandwidth is extended down to 90 Hz from the nominal 100 Hz.

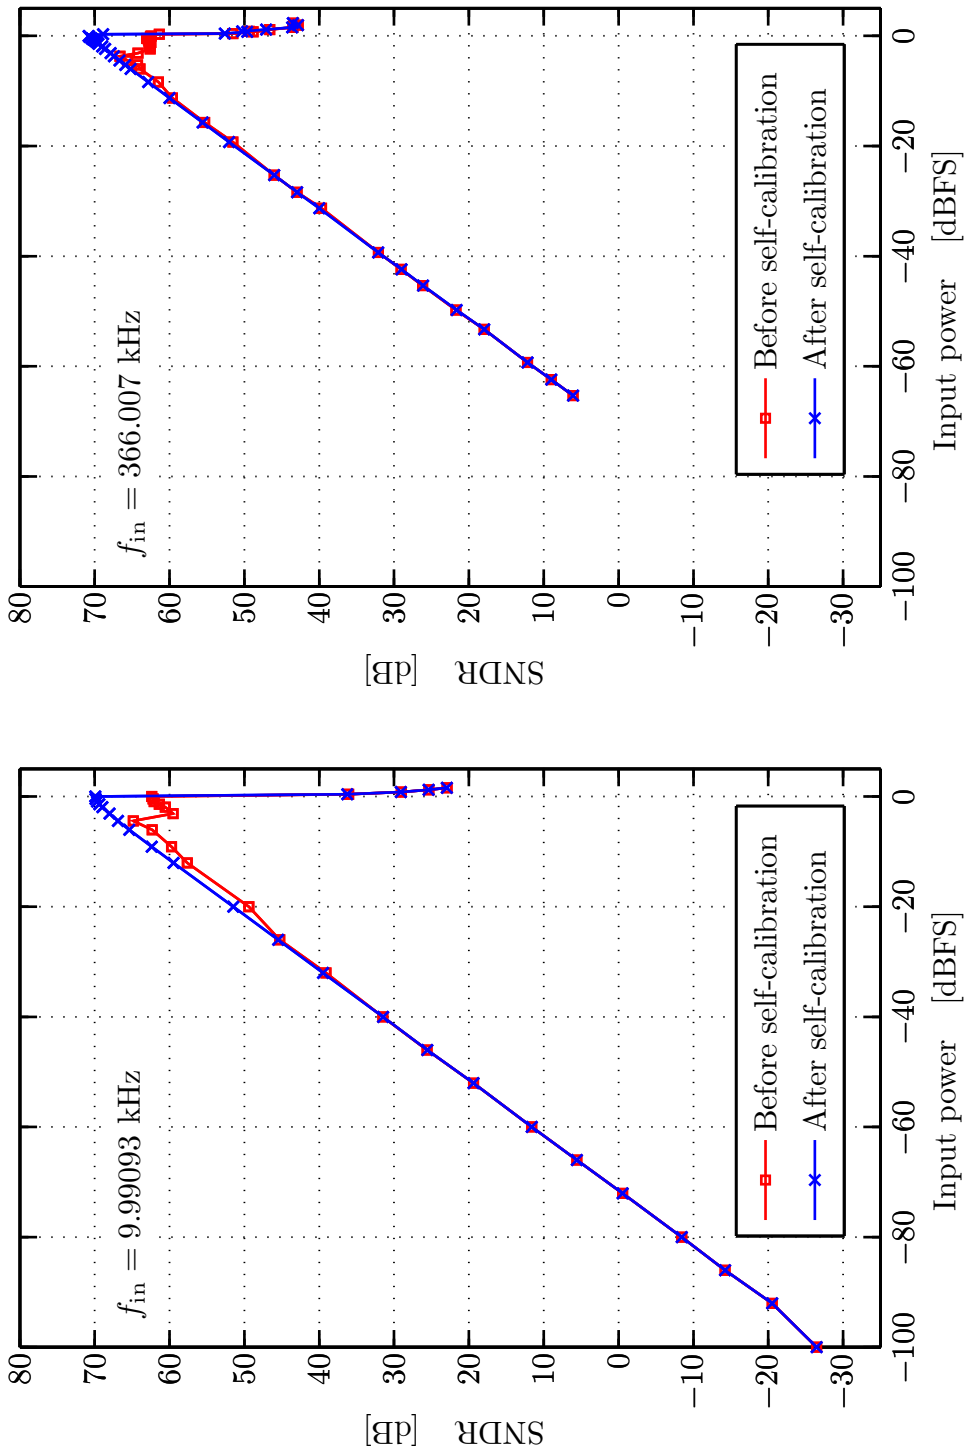

Figure 6.11: Measured SNDR performance of the SVADC-1 versus input power with both low frequency (left) and high (right) frequency inputs. Performance both before (in red) and after (in blue) self-calibration shown. Markers indicate measured data points.

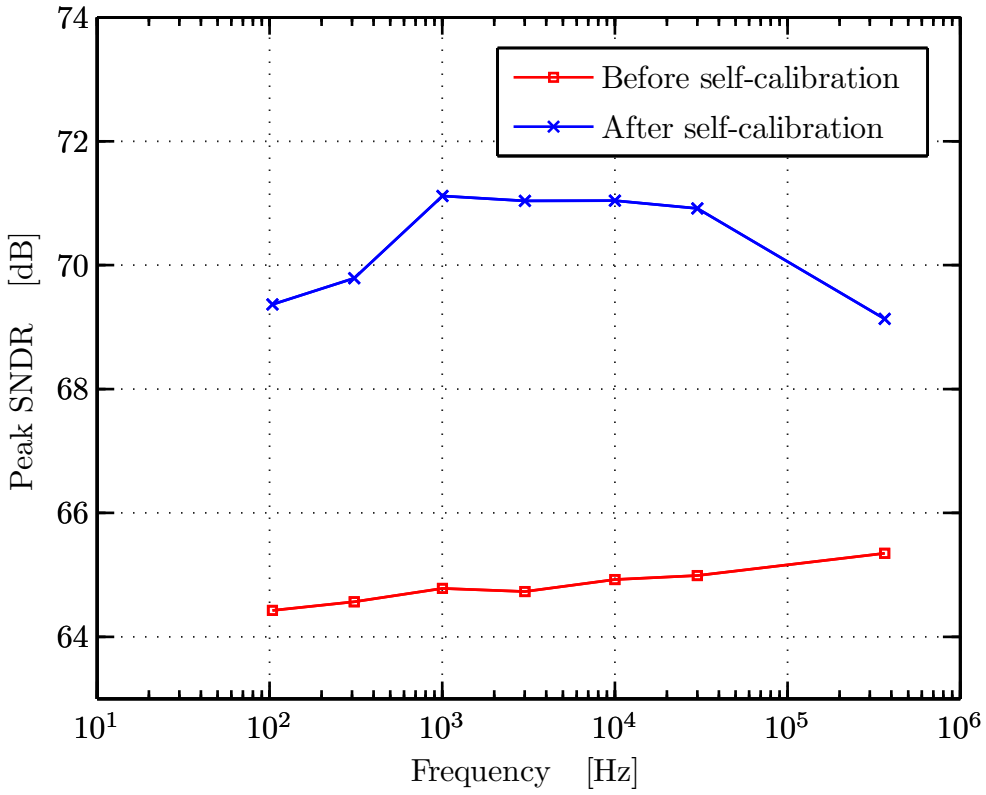

Figure 6.12: Measured peak SNDR performance of the SVADC-1 versus input frequency. Performance both before (in red) and after (in blue) self-calibration shown. Same calibration values used at all frequencies. For 103.712-Hz and 308.752-Hz inputs, 1-Hz/bin spectra used. All data except highest frequency input acquired with low-frequency input path; highest frequency acquired with high-frequency input path. Markers indicate measured data points.

is input to the ADC and the INL computed from statistical considerations of the output histogram [*Blair*, 1994; *IEEE Std 1241-2000*, 2000, Sect. 4.1.6.3].[14] For the

[14] In particular, it can be shown that, if $h[k]$ is a histogram of the ADC output codes (in unsigned integer form) in response to a slightly overranging sinusoidal input, and $h_C[k]$ the cumulative histogram:

$$h_C[k] = \sum_{l=0}^{k} h[l]$$

then [*Blair*, 1994; *Vanden Bossche et al.*, 1986]:

$$T_{\text{meas}}[m] = \cos\left[\frac{\pi h_C[k]}{S}\right]$$

SVADC-1, a 9.99093-kHz sinusoid is used, with the experimental setup configured for the low-frequency, nonharmonic measurement of Figure 6.2(b). The sinusoidal technique computes the transition levels to within a gain a and offset b, that is (recall Equation (3.17)):

$$a\,T_{\text{meas}}[m] + b + \epsilon[m] = T_{\text{ideal}}[m] \tag{6.3}$$

where the $T_{\text{meas}}[m]$ are the transition levels produced by the sinusoidal technique, the $T_{\text{ideal}}[m]$ are the ideal transition levels,[15] and $\epsilon[m]$ is the residual error after gain and offset correction. For this dissertation, a and b are chosen to minimize $\epsilon[m]$ over m in a mean-squared sense.[16]

The measured INL of the SVADC-1 before and after self-calibration is shown in Figure 6.13. Note that since the SVADC-1 output code is a 16-bit word, the INL is computed assuming $2^{16} = 65536$ total possible codes; however, since the conversion is nominally only 12 bits, the INL value itself is given in units of an LSB of $V_{\text{FS}}/2^{12} \simeq 488\ \mu\text{V}$. Before self-calibration, large discontinuities in the INL are seen. These discontinuities are indicative of analog-digital mismatch, especially in the first stage.

where S is the total number of samples used to assemble the histogram $h[k]$. Other techniques for measuring INL include the use of a DC input, and the use of a ramp input, to the ADC. In the former, a feedback loop adjusts the DC input such that 50% of the ADC output codes lie at code m or above, and 50% at code $m-1$ or below: the DC input then represents the transition level $T[m]$ between the codes m and $m-1$ [*IEEE Std 1241-2000*, 2000, Sect. 4.1.6.1]. In the latter, the input ramps through the full ADC input range: histograms of the ADC output codes are then used to estimate the transition levels [*IEEE Std 1241-2000*, 2000, Sect. 4.1.6.2]. Both techniques have disadvantages compared to the sinusoid input method: the DC input method entails a more sophisticated experimental setup, whereas the ramp input method requires generation of very linear ramp signals (which is more difficult than the generation of very linear sinusoidal signals).

[15] As a detail, the $T_{\text{ideal}}[m]$ are defined as a uniform spacing of transition levels between the ideal first and last transition levels, that is, between $V_{\text{REF}} - \Delta$ and $-V_{\text{REF}} + \Delta$. Notably, for the SVADC-1 the total number of transition levels is not necessarily $2^M - 1$, $M \in \mathbb{Z}$: assuming the SVADC-1 output is represented as a 16-bit unsigned integer, while the output code saturates at 0 on the lower end, on the higher end the output code commonly saturates before $2^{16} - 1 = 65535$. Before self-calibration, the early maximum code saturation occurs because of the extension of a 12-bit value to a 16-bit value: $2^4(2^{12} - 1) = 65520 < 65535$. After self-calibration, typically the lookup table entries conspire to cause early saturation (for example, see Figure 6.13). Assuming, then, that the minimum code is 0 and the maximum code C_{max}, $T_{\text{ideal}}[m]$ is generated by setting $T_{\text{ideal}}[1] = -V_{\text{REF}} + \Delta$ and $T_{\text{ideal}}[C_{\text{max}}] = V_{\text{REF}} - \Delta$ and equally spacing the remaining transition levels between these extrema.

[16] An alternate definition is to choose a and b such that the error at the extremum transition levels is 0, that is, such that $T_{\text{corr}}[0] = T_{\text{ideal}}[0]$ (i.e., $\epsilon[0] = 0$) and $T_{\text{corr}}[t_{\text{max}}] = T_{\text{ideal}}[t_{\text{max}}]$ (i.e., $\epsilon[t_{\text{max}}] = 0$), where t_{max} the highest transition level [*IEEE Std 1241-2000*, 2000, pp. 43–44].

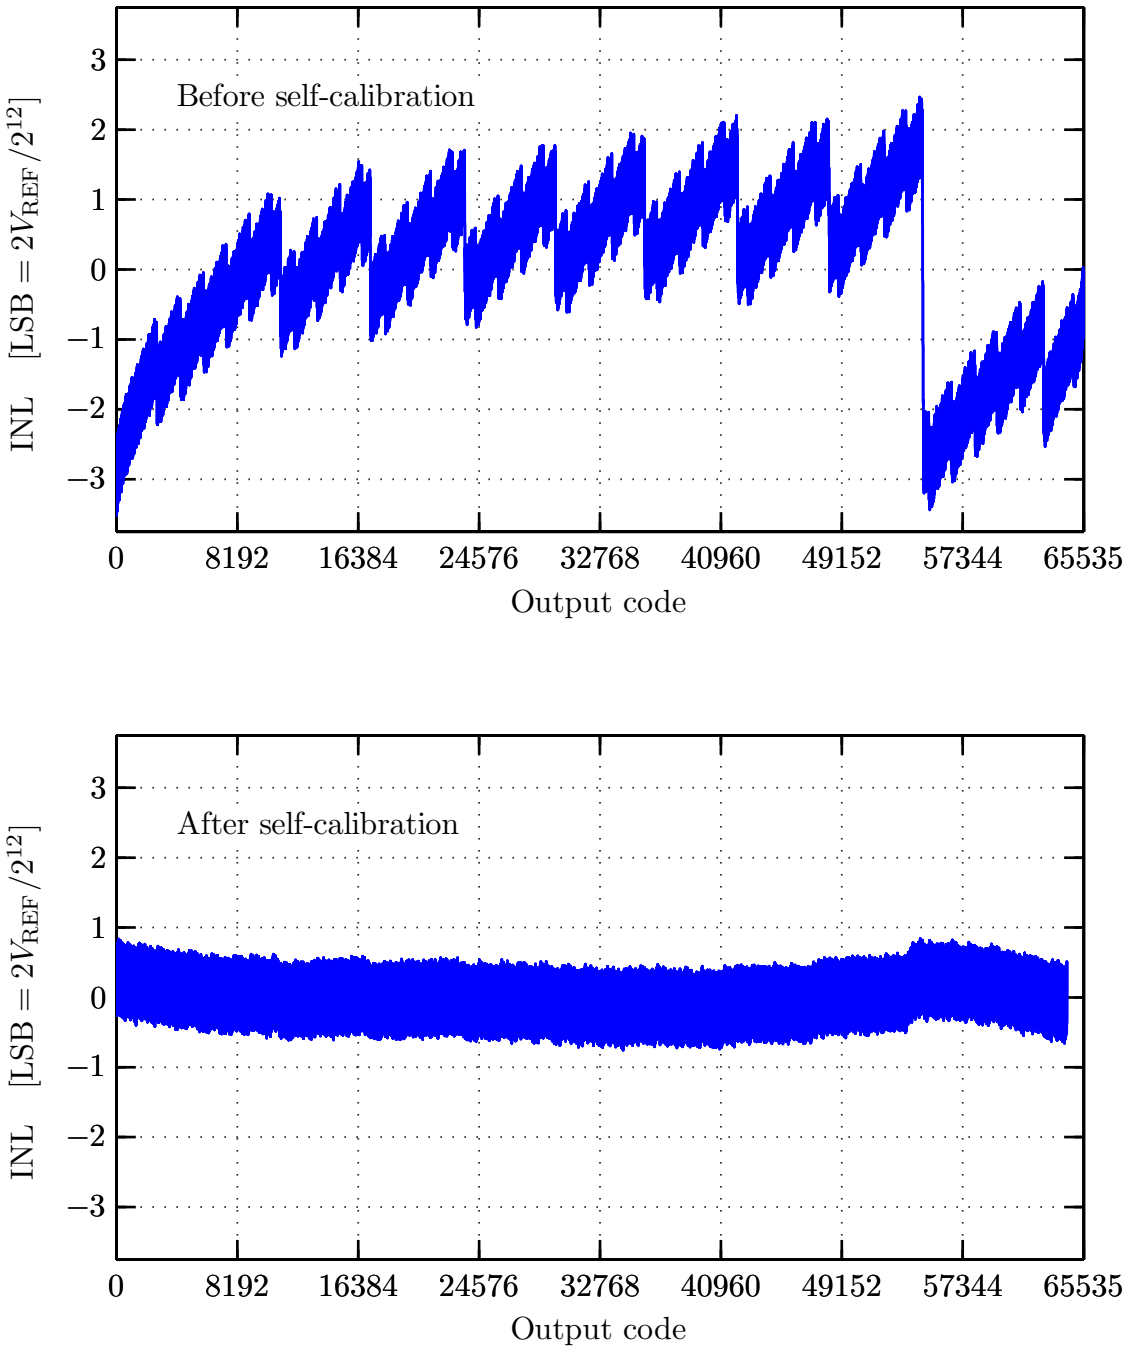

Figure 6.13: Measured INL performance of the SVADC-1, both before (above) and after (below) self-calibration. INL given in LSB assuming 12-bit conversion. 10,000,000 samples of SVADC-1 output in response to low frequency 9.99093-kHz input sinusoid at 0.99 dBFS used to compile INL.

After self-calibration, though, these discontinuities are removed, resulting in a much more uniform quantization and ultimately, better converter linearity. Quantitatively, the INL before self-calibration lies between −3.52 LSB and +2.47 LSB, improving to −0.76 LSB and +0.84 LSB after self-calibration.

6.3 Radiation Performance

Radiation testing[17] of the SVADC-1 includes both total-dose and single-event testing as summarized in Tables 6.1 and 6.2, respectively. These experiments were conducted in collaboration with The Aerospace Corporation, Lawrence Berkeley National Laboratory, National Semiconductor Corporation, and Texas A&M University. The discussions of this section concentrate on the testing results of the SVADC-1, specifically, the total-dose testing of the SVADC-1 by 50-MeV protons, and the single-event testing of the SVADC-1 by 10-MeV/nucleon heavy ions, 25-MeV/nucleon heavy ions, and pulsed laser. Details of these experiments are available in Appendix B for the 50-MeV proton testing, and Appendices C, D, E for the 10-MeV/nucleon heavy-ion, 25-MeV/nucleon heavy-ion, and pulsed-laser testing, respectively.[18]

6.3.1 Total-Dose Testing

For total-dose testing, the SVADC-1 is irradiated by 50-MeV protons in logarithmic dose steps up to a total dose of 2 Mrad(Si) as described in Appendix B. At each dose, the SVADC-1 is assessed for spectral, transition level, and power performance. Spectral and transition level performance is evaluated by 9.99093-kHz sinusoidal input testing, and power performance by 200-kHz sinusoidal input testing. For both cases,

[17]The basic concepts of radiation performance assessment are presented in Sections 2.2 and 2.3 for single-event and total-dose effects, respectively. While the presentations of those sections are sufficient for understanding the discussions here, a more complete introduction to radiation testing can be found in Appendix A.

[18]Regarding the remaining experiments: results from the ^{60}Co total-dose testing of the NMOS layouts are presented in Sections 2.3.3 and 5.2.2; the testing itself is described in Appendix F. Results from the ^{60}Co total-dose testing of the SVADC-1 are similar to those presented here for the 50-MeV proton total-dose testing. As proton testing is more rigorous than ^{60}Co—proton irradiation induces both displacement and ionization damage, whereas ^{60}Co irradiation induces only the latter—this dissertation elects to concentrate on the proton results.

DUT	Source	Maximum dose [krad(Si)]	Dates	Facility	Experimenters
SVADC-1	^{60}Co (γ-ray)	2000	28–29 Jun. 2006	The Aerospace Corporation, El Segundo, CA	Charles C. Wang[a], Steven C. Witczak[b], Mark Turpin[b], James L. Roeder[b]
NMOS layouts	^{60}Co (γ-ray)	2000	10–16 Oct. 2006	The Aerospace Corporation, El Segundo, CA	Charles C. Wang[a], Everett E. King[b], Mark Zakrzewski[b], Stephanie Brown[b], James L. Roeder[b]
SVADC-1	50-MeV protons	2000	8–11 Feb. 2008	88-inch Cyclotron at Lawrence Berkeley National Laboratory (LBNL), Cave 4A	Charles C. Wang[a], Jeffery S. George[b], Rocky Koga[b], Michael B. Johnson[c]

[a] Experimenter is with Stanford University.
[b] Experimenter is with The Aerospace Corporation.
[c] Experimenter is with Lawrence Berkeley National Laboratory.

Table 6.1: List of total-dose radiation testing performed in support of this dissertation.

DUT	Source	Maximum LET [MeV-cm^2/mg]	Dates	Facility	Experimenters
SVADC-1	Pulsed laser	[n/a]	30 Jun. 2006, 26–27 Oct. 2006	The Aerospace Corporation, El Segundo, CA	Charles C. Wang[a], Stephen D. LaLumondiere[b]
SVADC-1	Heavy ions (10 MeV/nuc.)	58.78	26–27 Sep. 2006	88-inch Cyclotron at Lawrence Berkeley National Laboratory (LBNL), Cave 4B	Charles C. Wang[a], Benjamin J. Mossawir[a], Jeffery S. George[b], Rocky Koga[b], Van T. Tran[b]
SVADC-1	Heavy ions (25 MeV/nuc.)	63	26–27 May 2008	Cyclotron at Texas A&M University (TAMU)	Kirby Kruckmeyer[c], Tom Santiago[c]

[a] Experimenter is with Stanford University.
[b] Experimenter is with The Aerospace Corporation.
[c] Experimenter is with National Semiconductor Corporation.

Table 6.2: List of single-event radiation experiments performed in support of this dissertation.

the low-frequency experimental setup described in Section 6.1 is used, albeit with some minor amendments.[19] In addition, the SVADC-1 performance is assessed before the onset of irradiation (referred to as pre-irradiation, or "Prerad" in the following figures) and after a 65-hour anneal in an unbiased state (referred to as "Anneal").

6.3.1.1 Spectral Performance

Spectral performance is assessed by the same methods as for baseline performance measurements (see Section 6.2.1). Figure 6.14 shows the resulting measured peak SFDR as a function of dose. For clarity, only the after self-calibration peak SFDR is shown. Little variation is seen in the peak SFDR: through the full 2 Mrad(Si) tested, the peak SFDR shows deviations of at most 2.9% dB from its pre-irradiation value. In particular, up to 1 Mrad(Si), the SVADC-1 maintains at least 90.1-dB peak SFDR, which is within the specifications of Table 3.1. At higher doses, the peak SFDR dips below 90 dB, but nonetheless remains high—at least 88.2 dB—up to the full 2 Mrad(Si).

The measured peak SNDR after self-calibration is shown as a function of dose in Figure 6.15. Again, little variation is seen with dose: the peak SNDR varies by at most 0.2% dB from its pre-irradiation value over the full 2 Mrad(Si) tested. Furthermore, the minimum peak SNDR is 70.6 dB over the full 2 Mrad(Si). This result affirms the radiation-hardness-by-design techniques of the SVADC-1. For example, if the I_{LEAK} of Figure 2.15 were present, gain loss from output droop throughout the stages would increase Q_{pipeline} (see Equation (4.15)), even after self-calibration. The consequent increase in quantization noise would manifest as SNDR loss. However, no such loss is seen here.

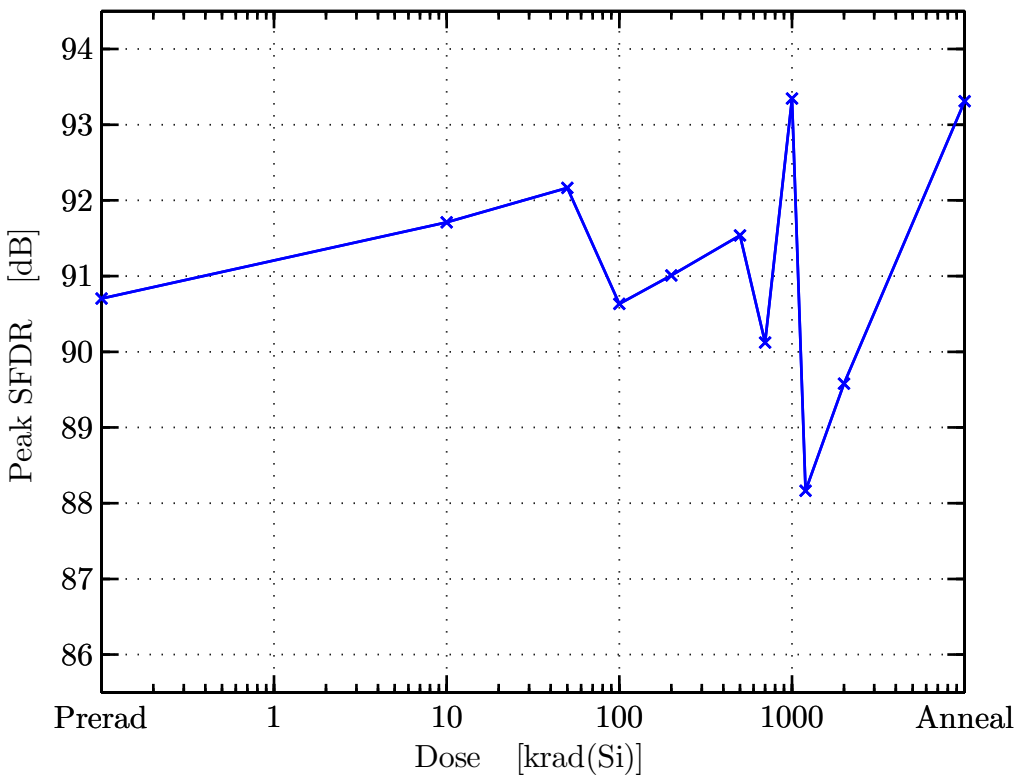

Figure 6.14: Measured peak SFDR performance of the SVADC-1 versus total dose. Peak SFDR measured in response to a 9.99093-kHz input sinusoid. Only after self-calibration peak SFDR shown. Markers indicate measured data points.

6.3.1.2 Transition Level Performance

Transition level performance of the SVADC-1 over dose is assessed by the same methods as for baseline performance measurements (see Section 6.2.2). The resulting measured maximum and minimum INL after self-calibration are shown in Figure 6.16. Small variations are seen with dose: the maximum INL bound changes by at most 11.1% LSB, and the minimum INL bound by at most 12.5% LSB, over the full 2 Mrad(Si) tested. However, in absolute terms the INL performance remains very

[19]Specifically, for nonharmonic performance assessment, instead of attaching an RC filter to the SRS DS360 output as shown in Figure 6.2(b), the SRS DS360 output is instead processed by the signal path of Figure 6.3 configured with 0-dB attenuation. This nonharmonic path (i.e., now Figure 6.3) is independently calibrated so that the signal at the input to the DUT is the same as that through the harmonic path (i.e., Figure 6.2(a)).

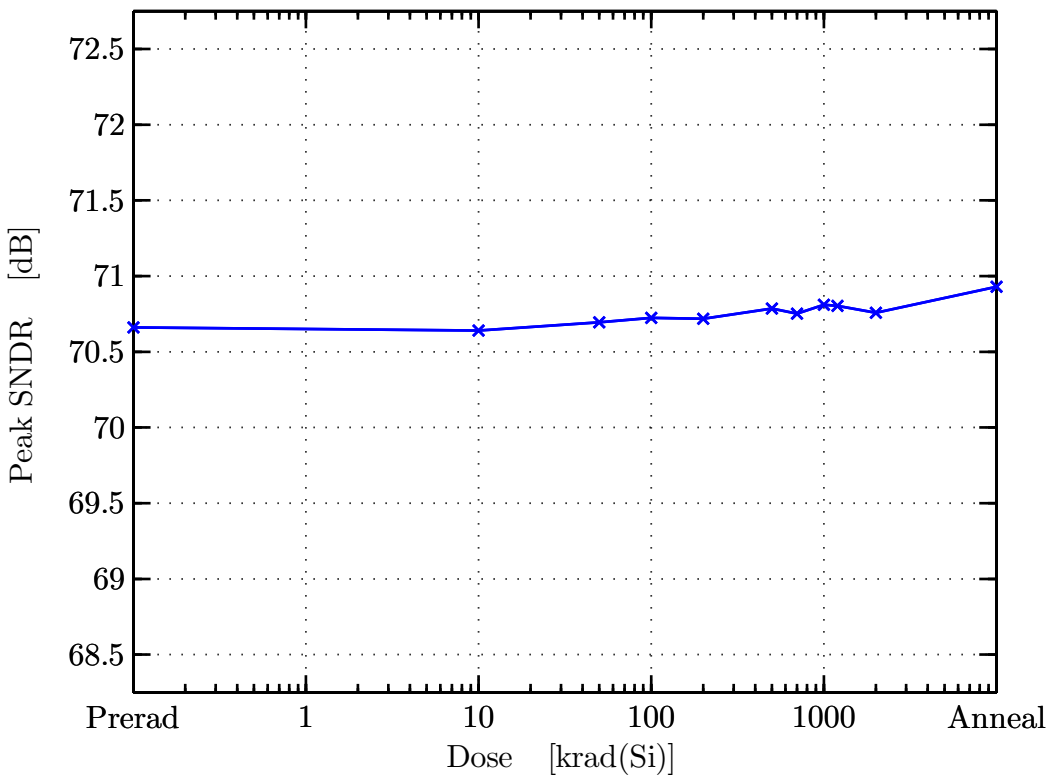

Figure 6.15: Measured peak SNDR performance of the SVADC-1 versus total dose. Peak SNDR measured in response to a 9.99093-kHz input sinusoid. Only after self-calibration peak SNDR shown. Markers indicate measured data points.

good, with worst case bounds of −0.76 LSB and +0.89 LSB over the full 2 Mrad(Si) tested.

In addition, the gain a and offset b of Equation (6.3) are shown as functions of dose in Figure 6.17. Given that $V_{\mathrm{REF}} = 1$ V for the SVADC-1, ideally $a = 1$ and $b = 0$ V. Hence the measured gain and offset are slightly nonideal. However, more importantly, they change little with dose, varying by just 3.7% and 1.5%, respectively, over the full 2 Mrad(Si) tested.

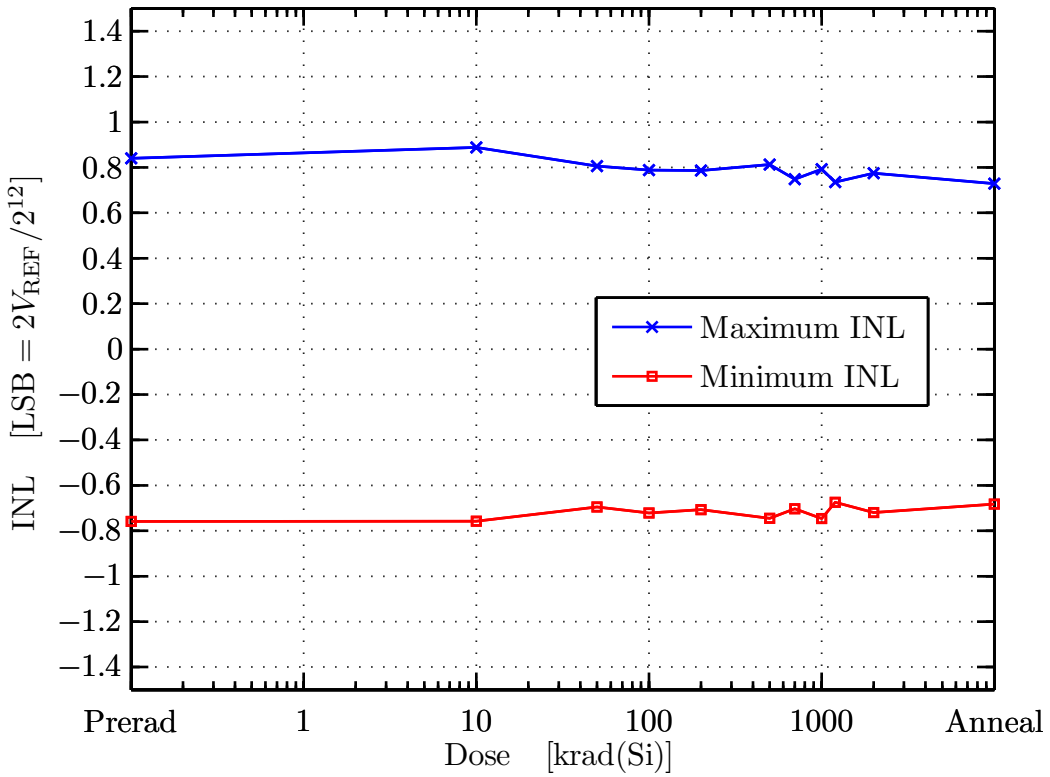

Figure 6.16: Measured INL performance of the SVADC-1 versus total dose. INL given in LSB assuming 12-bit conversion. To compile INL, 10,000,000 samples of SVADC-1 output in response to low frequency 9.99093-kHz input sinusoid at 0.99 dBFS used. Only after self-calibration INL shown. Markers indicate measured data points.

6.3.1.3 Power Performance

For power measurements, a 200-kHz sinusoid is input to the SVADC-1 to ensure the SVADC-1 sufficiently busy. Note that the SVADC-1 power consumption is input-amplitude dependent: for the results presented here, an input of −6.02 dBFS is used as this power was found to yield the highest power consumption at this input frequency.[20]

[20]The maximum total power consumption of the SVADC-1 does not occur at full scale or with overranging inputs. Of the five SVADC-1 supplies, it is the power consumption of the $V_{DD,IO}$ supply which varies the most with input signal amplitude. As the majority of the $V_{DD,IO}$ power is consumed by the digital pad I/O circuitry (see Section 5.5.5), the maximum SVADC-1 power consumption occurs at the input amplitude that creates the most activity in the output bits of the

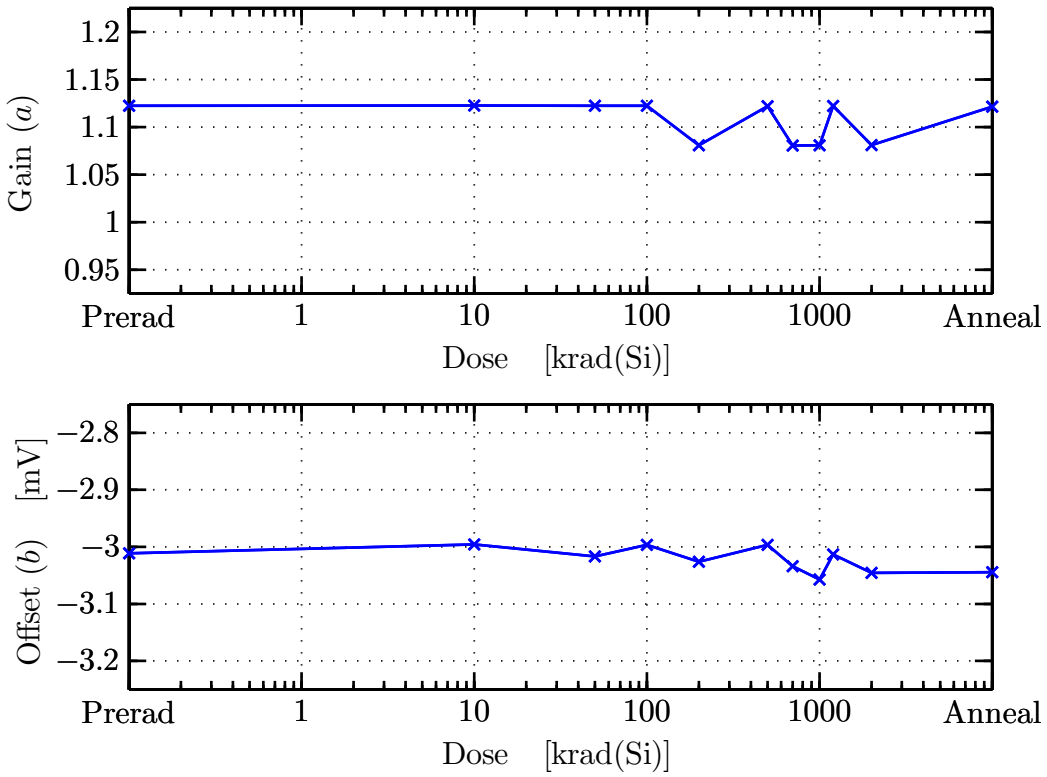

Figure 6.17: Measured gain a and offset b of the SVADC-1 versus total dose. Refer to Equation (6.3) for definitions of a and b. Only gain and offset after self-calibration shown. Markers indicate measured data points.

The measured total SVADC-1 power consumption as a function of dose is shown at top in Figure 6.18. While the power consumption remains steady around 50 mW up to 200 krad(Si), at higher doses there is an increase to roughly 60 mW. The reason is shown in the lower plots of Figure 6.18, which shows the power of each supply separately. The increased power consumption arises from the $V_{\mathrm{DD,D}}$ and $V_{\mathrm{DD,CLK}}$ supplies: over the 2 Mrad(Si) tested, these supplies demonstrate increases of as much as 24.3% and 89.2%, respectively, from their pre-irradiation values. Notably,

stages. Consider when the input amplitude is very small: then there is little variation in the output bits and the SVADC-1 power consumption is not at its maximum. On the other hand, consider when the input amplitude is overranging: then during the overrange period the outputs bits again display little variation (simply saturating at the extremum values) and the power consumption is not at its maximum. Therefore, the maximum SVADC-1 power consumption must occur at some large, but not overranging amplitude.

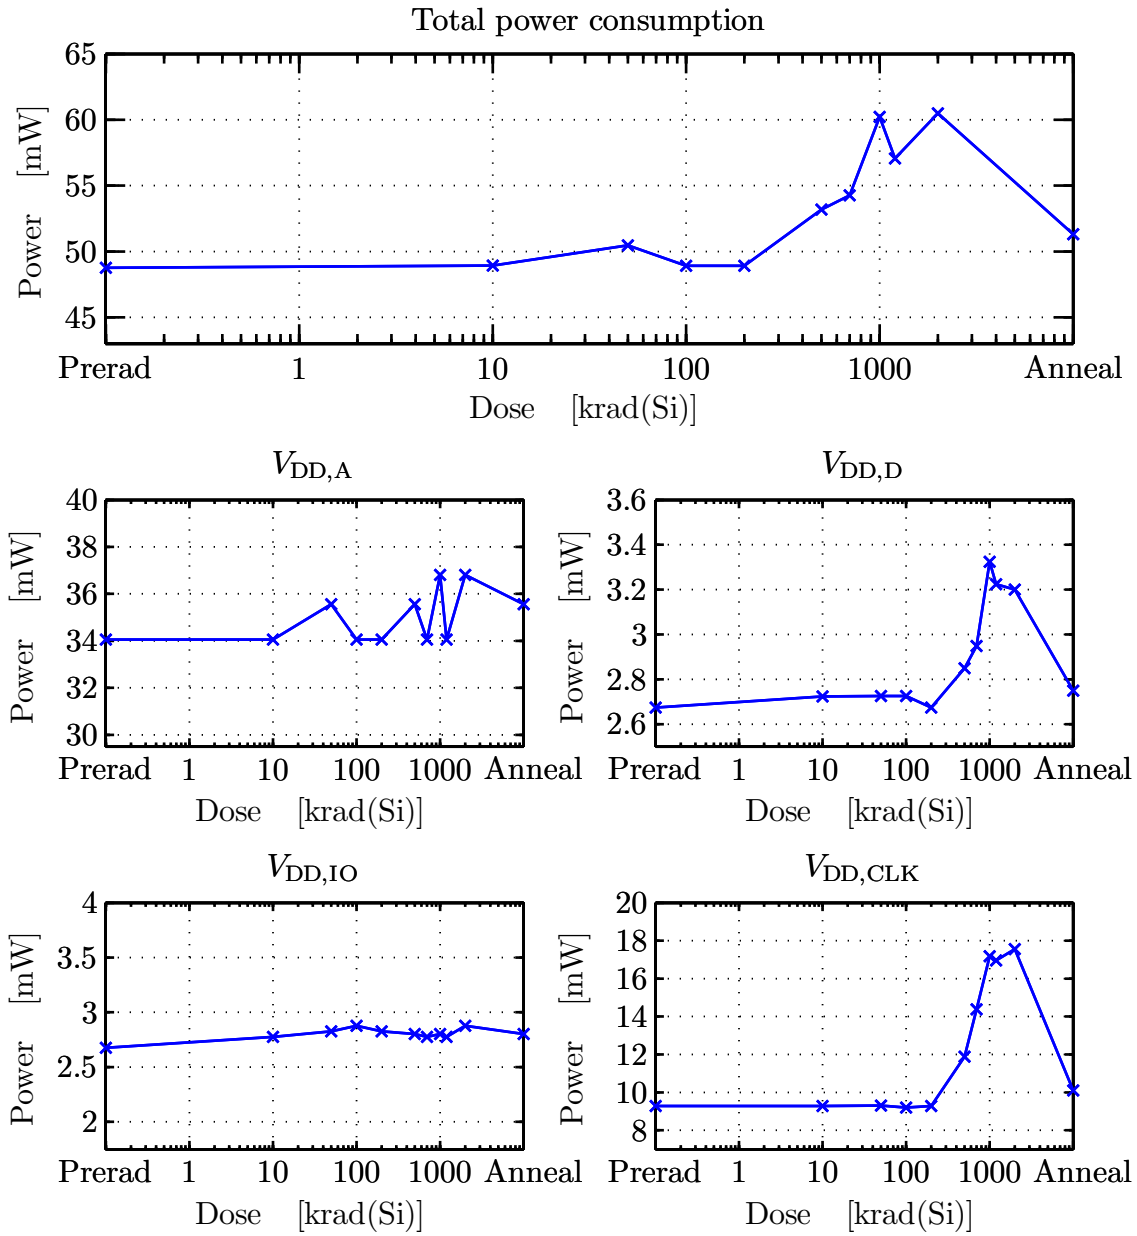

Figure 6.18: Measured power consumption of the SVADC-1 versus total dose. Total power consumption shown at top, power consumption of each supply shown below ($V_{DD,REF}$ is omitted as its contribution to total power is negligible compared to the other four supplies). Power measured in response to a 200-kHz input sinusoid of −6.02-dBFS power. Markers indicate measured data points.

the logic gates attached to these supplies do not generally use enclosed terminal layouts. In contrast, those attached to the $V_{DD,IO}$ supply do, and hence its power consumption changes much less—by just at most 7.5% over the 2 Mrad(Si) tested—with radiation.[21,22] The power evolution of these different supplies, then, further confirms the efficacy of the enclosed terminal layout technique.

6.3.2 Single-Event Testing

Multiple radiation testing experiments are undertaken to assess the single-event performance of the SVADC-1. These experiments include heavy-ion irradiation—via both a 10-MeV/nucleon heavy-ion beam at Lawrence Berkeley National Laboratory, and a 25-MeV/nucleon heavy-ion beam at Texas A&M University—and pulsed-laser testing. It is notable that all DUTs remained functional throughout the duration of testing: no hard errors (such as latchup, which is discussed in more detail below) were observed in any of the experiments.

The single-event testing of the SVADC-1 required the construction of experimental setups custom tailored to the particular constraints of each testing venue. Complete descriptions of the experiments can be found in Appendices C, D, and E for the 10-MeV/nucleon heavy-ion, 25-MeV/nucleon heavy-ion, and pulsed-laser testing, respectively.

6.3.2.1 Latchup

To assess latchup susceptibility, the current supplied to each of the five SVADC-1 supplies is monitored. An example is shown in Figure 6.19, which plots the measured current of each supply sampled over a 5-minute heavy-ion exposure. If latchup were to occur in an SVADC-1 supply, then the measured current would rapidly rise and

[21] Due to time constraints, only the digital cells attached to $V_{DD,IO}$ are drawn with enclosed terminal layout applied to the NMOS devices: the standard cell library attached to $V_{DD,D}$ and $V_{DD,CLK}$ is drawn in standard fashion (see Section 5.5.4). Future revisions of the SVADC-1 should evaluate whether the increase in power at high doses warrants redrawing of the digital logic attached to $V_{DD,D}$ and (especially) $V_{DD,CLK}$.

[22] For completeness, the $V_{DD,A}$ supply power consumption changes by at most 8.1% over the 2 Mrad(Si) tested.

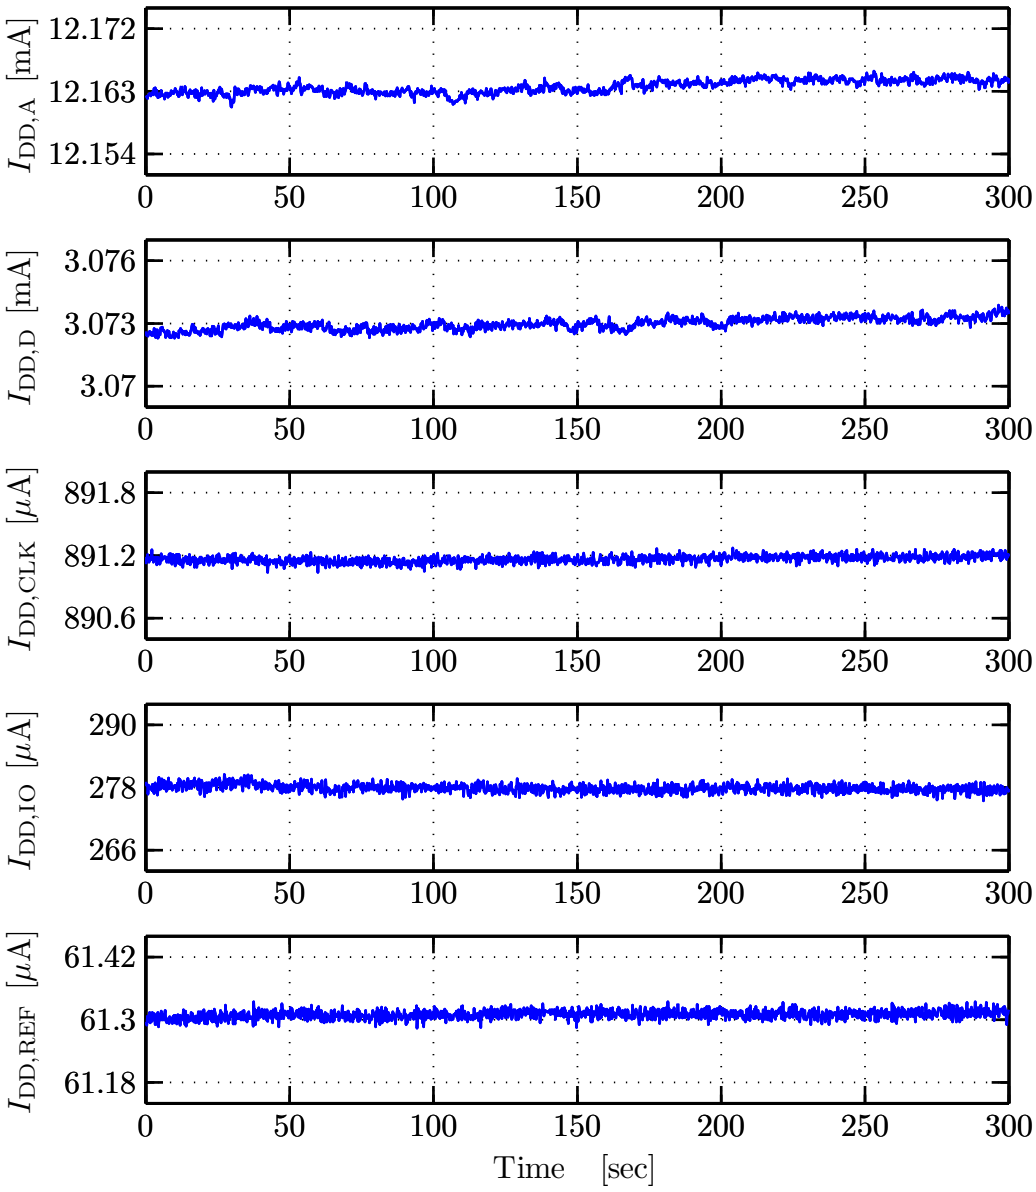

Figure 6.19: Example of current monitoring of all five DUT supplies for latchup assessment. Data taken from 10-MeV/nucleon heavy-ion testing with LET set at 58.72 MeV-cm^2/mg. 1-mV DC signal input to SVADC-1.

quickly saturate at the current limit of the source feeding the SVADC-1 supply (the current limit is 500 mA in the data of Figure 6.19). Since latchup is a stable condition, the current would remain at this elevated level until the latched SVADC-1 supply was power cycled or the SVADC-1 burned out. Figure 6.19, then, displays no such current rise in any SVADC-1 supply and hence no latchup.

Indeed, no latchup is seen in any SVADC-1 supply for the duration of all radiation testing experiments. These experiments include testing at standard supply voltage (i.e., all DUT supplies at 2.5 V) and room temperature, and also testing at elevated supply voltage (as high as 2.7 V) and elevated temperature (as high as 131°C). Specifically:

- *Pulsed laser*

  Frontside pulsed-laser testing of the SVADC-1 is performed to determine whether more expensive heavy-ion testing was justified: if a DUT experiences latchup under pulsed-laser testing, then it most certainly will experience latchup under heavy-ion exposure. In this testing, the SVADC-1 is probed by a 590-nm, actively mode-locked, cavity-dumped dye laser (as described in *Moss et al.* [1995]) with beam energies as high as 4.26 nJ/pulse at standard supply voltages and room temperature. Both DC and AC signals are input to the SVADC-1. No latchup is seen during testing. Notably, the pulsed-laser testing is not comprehensive: since the laser beam reflects off metallization, only select portions of the chip are probed.

- *Heavy ions, 10 MeV/nucleon*

  A more complete latchup test is undertaken via exposure to 10-MeV/nucleon heavy ions at Lawrence Berkeley National Laboratory. No latchup is observed during testing, which includes exposure up to a maximum LET of 58.72 MeV-cm^2/mg (the highest available) at standard supply voltage and room temperature. Two separate SVADC-1 DUTs are tested. At maximum LET, the first displays no latchup under a total beam fluence of 7.22×10^6 ions/cm^2 with a DC input. At maximum LET, the second displays no latchup under a total

beam fluence of 2.85×10^7 ions/cm^2 with a mixture of both DC and AC inputs.

- *Heavy ions, 25 MeV/nucleon*
  An more thorough latchup test is undertaken via exposure to 25-MeV/nucleon heavy ions at Texas A&M University. Since the SVADC-1 uses a non-epitaxial substrate, the greater penetration depth achieved by this higher-energy beam is used to more deeply probe latchup susceptibility. To further stress latchup susceptibility, the supply voltage of all the DUT supplies is increased to 2.7 V and the temperature raised from room temperature to 131°C. For this testing, a Xenon beam (the highest LET ion available) is angled by 48° to achieve an effective LET of 62 MeV-cm^2/mg. No latchup is observed during testing. In particular, at the highest temperature of 131°C the SVADC-1 is exposed to a total beam fluence of 1.09×10^7 ions/cm^2 while processing an AC input.

To summarize the result of the most thorough latchup testing: the SVADC-1 shows no latchup up to an LET of 62 MeV-cm^2/mg at elevated supply (2.7 V) and temperature (131°C) during 25-MeV/nucleon heavy-ion exposure.

6.3.2.2 Turflinger Analysis Description

Soft errors are assessed by computing the SVADC-1 cross section versus linear energy transfer (LET) (see Section 2.2.1). Measuring a cross section, though, requires a definition of what constitutes an ADC soft error. Technically, this definition is application-dependent. However, a general method for describing ADC soft errors was proposed by *Turflinger and Davey* [1990]. The Turflinger analysis is adopted in this dissertation.

Conceptually, the Turflinger analysis distinguishes between two types of soft errors: noise errors and offset errors. Assume a DC signal is input to the ADC. While a converter ought output a single code in response, in reality circuit noise causes a Gaussian distribution of codes. Multiple soft errors caused by radiation-induced upsets, aggregated over time, can widen this distribution. Errors relating to this widening are classified as noise errors. In contrast, soft errors can also cause large offsets in the output code, well beyond the bounds of the widened noise distribution.

Such errors are classified as offset errors. Determination of the Gaussian distribution bounds that distinguish between noise and offset errors is typically unique for each ADC [*Turflinger et al.*, 1994].

The Turflinger analysis is illustrated in Figure 6.20. First, before any irradiation, a baseline ADC response is measured: a DC signal is input to the ADC and a histogram of the resulting ADC output codes is compiled. The baseline histogram of the SVADC-1 from the 10-MeV/nucleon heavy-ion testing is shown in Figure 6.20(a). The histogram is fit to a Gaussian (indicated by a solid red line) and the output codes "within" this Gaussian dubbed the "signal bins". More quantitatively, assuming a baseline Gaussian with standard deviation σ_B and mean μ_B, the bins within $\mu_B \pm \alpha_B \sigma_B$, for α_B a spread coefficient (discussed later), are signal bins. In Figure 6.20(a), $\alpha_B = 4$ and the consequent signal bins are those bins between the red dashed lines.

To characterize the cross section, the ADC is exposed to a radiation source of a given LET. The same DC signal is input to the ADC and an output code histogram compiled. As the signal bins do not represent soft errors, they are removed: as an example, the histogram for an LET of 58.72 MeV-cm^2/mg, sans signal bins, is shown in Figure 6.20(b). The resulting histogram is then fitted to another Gaussian (indicated by the solid green line). This second Gaussian distinguishes the noise and offset errors. Specifically, assuming the second Gaussian has a standard deviation σ_E and mean μ_E, the bins within $\mu_E \pm \alpha_E \sigma_E$, for α_E a spread coefficient (discussed later), are considered noise errors: the sum of counts over these "inner" bins defines the noise error cross section. Bins without this range, then, are offset errors: the sum of counts over these "outer" bins defines the offset error cross section. In Figure 6.20(b), $\alpha_E = 2$ and the bins between the green dashed lines are noise errors, while those outside are offset errors. If desired, the total errors, defined as the sum of the noise and offset errors, can also be computed.[23]

[23] It is worthwhile to stress that, to conduct a Turflinger analysis, the irradiation flux and exposure duration should be adjusted to guarantee a statistically significant number of both noise and offset errors in the irradiated histograms. However, as it is difficult to predict the proper flux and duration *a priori*, these parameters are often determined by trial and error during testing. Thus single-event experiments must rapidly generate (at least preliminary) analytical results during the experiment; see Section A.3.4 for a more thorough discussion of the consequent impact on experiment design.

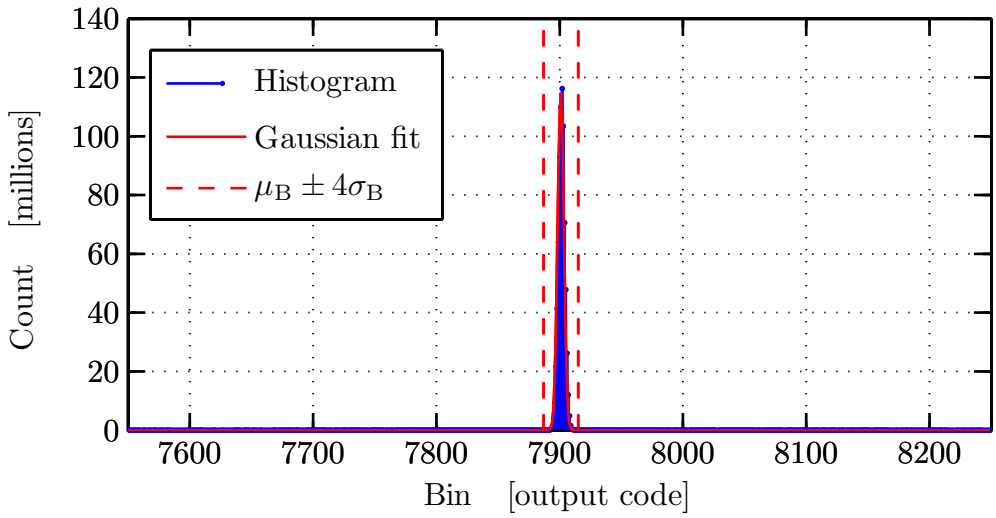

(a) Baseline histogram.

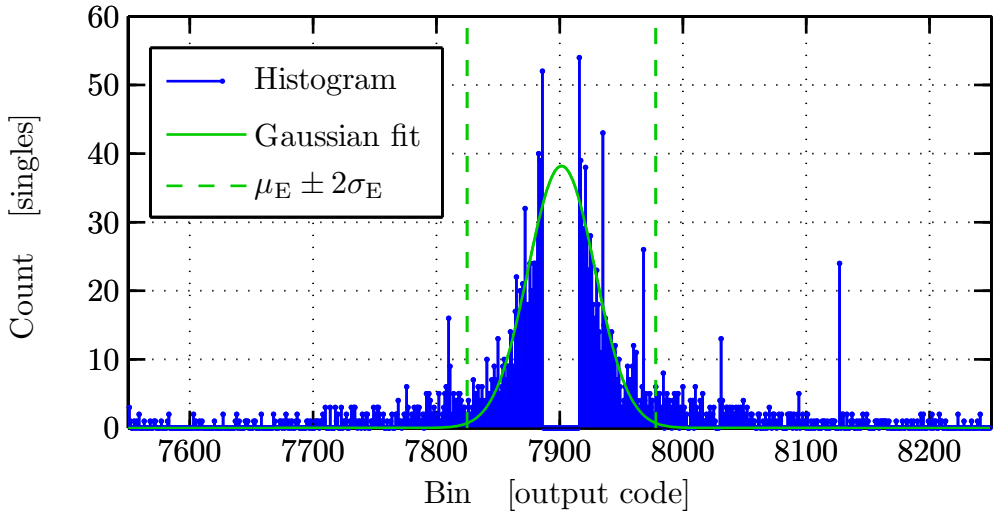

(b) Irradiated histogram (sans signal bins).

Figure 6.20: Illustration of the Turflinger analysis for ADC soft errors. (a) Baseline ADC output histogram. Gaussian fit establishing signal bins shown in red: signal bins lie within red dashed lines. (b) Irradiated ADC output histogram, sans signal bins. Gaussian fit establishing noise error bins shown in green: noise error bins lie within green dashed lines, offset error bins lie without. All data shown captured during 10-MeV/nucleon heavy-ion testing of the SVADC-1; in particular, irradiated data compiled during exposure to a heavy-ion beam of 58.72 MeV-cm^2/mg LET while processing a 1-mV DC input.

Naturally, the Turflinger analysis is sensitive to the choice of α_B and α_E. The choice of α_B, while technically arbitrary, in practice is often limited by concerns in the second Gaussian fit: if α_B is chosen too small, large remaining signal bins can derail convergence of the second fit.[24] The choice of α_E is less stringent, and left to the discretion of the experimenter. α_B and α_E are often empirically determined and should be explicitly stated.

6.3.2.3 Turflinger Analysis Results

For the SVADC-1, a Turflinger analysis is conducted on the 10-MeV/nucleon heavy-ion testing results. It is known that ADCs can show different cross sections given different input conditions [*Buchner et al.*, 2005].[25] Thus, for this testing two different input DC values are tested: a 1-mV DC input wherein the SVADC-1 internal differential circuitry is balanced, and a 0.5-V DC input wherein the SVADC-1 internal differential circuitry is unbalanced.[26] As the Turflinger analysis is more concerned with functional errors, the analysis is performed on the SVADC-1 output before self-calibration. The resulting noise, offset, and total error cross sections are shown as a function of LET in Figure 6.21.[27] For both DC inputs, $\alpha_B = 4$ and $\alpha_E = 2$. Notably, the cross sections for both DC inputs are very similar, with noise errors contributing the majority of the soft errors. The cross sections are also fit (via a nonlinear least

[24] Naturally, if α_B is chosen too large, it can artificially depress the ADC cross section by over-representing the space of non-error signal bins.

[25] An input that stresses a variety of input conditions would be an AC signal such as a sinusoid. Such signals are often avoided, though, due to difficulties in defining the ideal output response from which the DUT deviates during irradiation. However, in theory it should be possible to use least-squares three-parameter or four-parameter fits of the ADC output to a sinusoid (as discussed in footnote 22 of Chapter 3) to define the ideal output sinusoid of the ADC. With this ideal output established, the Turflinger method of defining signal bins and subsequently characterizing noise and offset soft errors can be applied to derive AC cross sections of the ADC. Unfortunately, the single-event testing of the SVADC-1, while including AC input testing, did not capture such data over a sufficient number of LET values to test this hypothesis.

[26] Unfortunately, time constraints precluded the testing of additional input voltages.

[27] For additional information regarding the intermediate products of this analysis, see Section C.4.

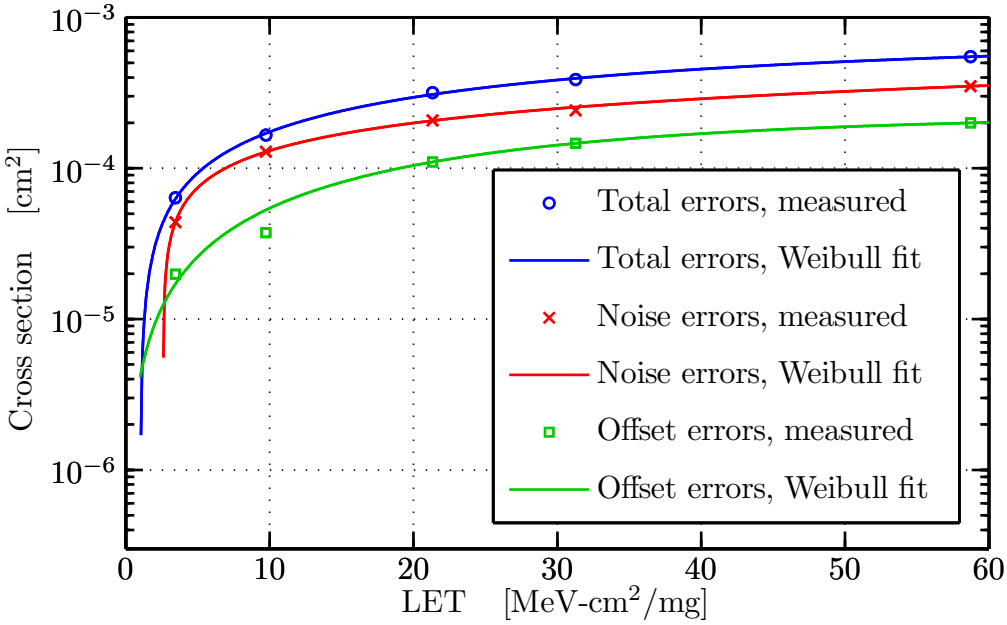

(a) With 1-mV DC input.

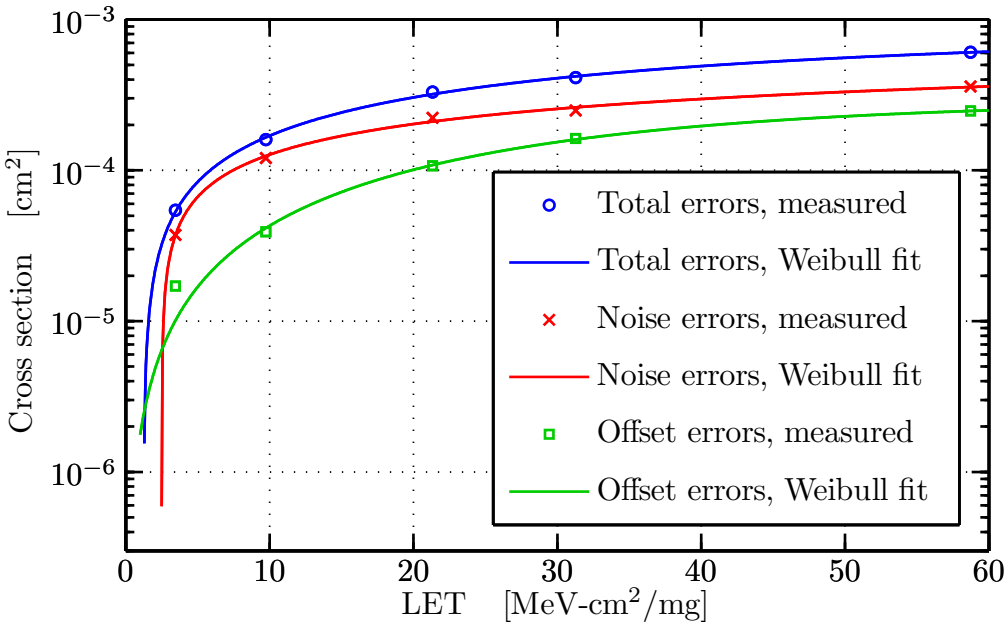

(b) With 0.5-V DC input.

Figure 6.21: Turflinger analysis cross sections of the SVADC-1 for a (a) 1-mV and (b) 0.5-V DC input as measured during 10-MeV/nucleon heavy-ion testing. Markers indicate measured data points, curves indicate Weibull fits (see Table 6.3).

Input	Error type	Weibull parameters			
		A_0	L_0	W	S
1-mV DC	Noise	2.115×10^{-3}	2.600	1561	0.5141
	Offset	2.240×10^{-4}	2.221×10^{-14}	29.77	1.175
	Total	7.798×10^{-4}	1.025	45.75	0.8442
0.5-V DC	Noise	6.905×10^{-4}	2.498	93.25	0.6323
	Offset	2.858×10^{-4}	5.617×10^{-14}	35.96	1.419
	Total	7.798×10^{-4}	1.247	49.28	0.8972

Table 6.3: Weibull parameters for the SVADC-1 cross section according to a Turflinger analysis. Compiled from 10-MeV/nucleon heavy-ion testing results. Refer to Equation (6.4) for parameter interpretation.

squares method) to Weibull distributions of the form:

$$\sigma_{\text{SEE}} = A_0 \begin{cases} 1 - e^{-[(L-L_0)/W]^S} & , \quad L > L_0 \\ 0 & , \quad L < L_0 \end{cases} \tag{6.4}$$

for σ_{SEE} the cross section and L the LET in MeV-cm^2/mg. The fits are plotted as continuous curves in Figure 6.21. The Weibull parameters themselves are summarized in Table 6.3. These parameters can be used to predict SVADC-1 upset rates given a particular radiation environment.

6.4 Application-Specific Metrics

In addition to standalone development, the SVADC-1 has also been integrated into plasma wave receivers on a variety of current scientific instruments, including the SpriteSat instrument (launched February 2009), the Firefly mission (launch scheduled in 2012), and the BBR (BroadBand Receiver) of the WIPER (Wave-Induced Precipitation of Electron Radiation) instrument aboard the DSX (Deployable Structures eXperiment) satellite (currently under development, with instrument

delivery scheduled in 2009 and launch scheduled in 2012).[28] These system integrations required additional design and testing efforts, the salient results of which are presented in this section.

6.4.1 FPGA Implementation

As the SVADC-1 chip contains the analog portion of the converter, for the aforementioned missions the digital portion of the SVADC-1 is written in Verilog and programmed into a radiation-hard FPGA. The digital portion includes the digital reconstruction, self-calibration algorithm, and additional interface logic. The self-calibration algorithm is implemented as per Section 6.2 but with fixed-point rather than floating-point computation. In particular, scaling bitwidths are used for the uncalibrated stages (with reconstruction implemented as per footnote 4 of Chapter 4), while for the calibrated stages all registers and arithmetic operations are performed at 18 bits to reduce roundoff error. The final output word is truncated to 16 bits. Self-calibration can be initiated at any time, and the code includes a preset for returning the lookup tables to before self-calibration defaults.

Table 6.4 presents the resource utilization of the digital portion when programmed into a radiation-hard Actel RTAX250S FPGA [*Actel Corporation*, RTAX-S].[29] The digital portion requires roughly one third of the total gates of the RTAX250S—for an estimated gate count of ∼85,000 gates—and at 5 MS/s consumes only ∼10-mW total dynamic power.[30] It is notable that the digital portion is constructed rather conservatively and not optimized: further reductions are certainly possible.[31]

[28] All dates current as of 12 June 2009.

[29] The RTAX250S is an antifuse-based, 250,000-gate FPGA specially designed to maintain performance up to a total dose of 200 krad(Si), and to maintain functionality up to a total dose of 300 krad(Si). Furthermore, the RTAX250S is latchup immune to an LET of >117 MeV-cm^2/mg and—through the use of special register designs—soft error immune to an LET of >37 MeV-cm^2/mg. The RTAX250S is manufactured in a 0.15-μm technology [*Actel Corporation*, RTAX-S].

[30] The remainder of the power is due to static leakage in the RTAX250S.

[31] For example, the digital reconstruction utilizes extensive pipelining: some of this pipelining can be eliminated for savings in register use. Possible savings are also available in the self-calibration algorithm: while the implementation currently uses 18 bits for all registers and arithmetic, a more complete study of the roundoff error may yield reduced bitwidths (especially for the downstream calibrated stages) and accompanying register savings.

Characteristic	Best	Typical	Worst
Clock frequency		5 MHz	
Cell usage			
Sequential (R cells)		754 (53.6% utilization)	
Combinatorial (C cells)		683 (24.3% utilization)	
Logic (R+C cells)		1437 (34.0% utilization)	
Pin usage			
Input pins		23	
Output pins		24	
Power consumption (all quantities in mW)			
Logic	0.4	0.5	0.5
Output pins	3.9	4.3	4.8
Input pins	0.1	0.1	0.1
Clock	3.9	4.4	4.8
DC	59.0	59.0	270.5
Power consumption by category (all quantities in mW)			
Static	59.0	59.0	270.5
Dynamic	8.4	9.3	10.3
Total	67.4	68.3	280.8

Table 6.4: Resource utilization of the digital portion of the SVADC-1 (including digital reconstruction and self-calibration algorithm) in an Actel RTAX250S. Best, typical, and worst case performance given.

Nonetheless, even as is, these results suggest that, compared to the analog portion, the digital portion can be implemented relatively inexpensively.[32]

[32]Designers interested in implementing the digital portion of the SVADC-1 in a custom integrated circuit are recommended to the dissertation of fellow Ph.D. student Hsiao-Heng K. Lee, which addresses the design of radiation-hard digital circuits in commercial CMOS manufacturing processes [*Lee*, in preparation].

6.4.2 Burn-In Testing

In support of the WIPER instrument, the SVADC-1 underwent burn-in testing. Burn-in testing seeks to screen infant mortality within a population of DUTs by subjecting said DUTs to high temperature stress. For the SVADC-1, all DUTs are stressed at 125°C for at least 160 hours. During this "burn-in", the DUTs are fully operational and under bias.[33] To screen the DUTs, they are evaluated for power consumption, aggregate gain, aggregate SFDR, and SINAD before and after burn-in.[34] Further details of the testing are available in Appendix K.

Two batches of SVADC-1s were burn-in tested. The first batch was composed of 23 SVADC-1s manufactured in the BiCMOS8iED manufacturing process. All proved functional after burn-in testing. In terms of performance, while outliers exist, after burn-in the vast majority showed changes of at most 1% in power consumption, 2% in aggregate gain, 5% in aggregate SFDR, and 3% in SINAD compared to their before burn-in values (see Section K.5 for the complete data tables).

For the second batch of SVADC-1s, the SVADC-1 design was migrated to the BC8BPLUS manufacturing process. BiCMOS8iED and BC8BPLUS are both BiCMOS processes generously provided by National Semiconductor Corporation that couple bipolar NPN devices with a 0.25-μm CMOS technology. Although the two processes differ in the manufacture of their bipolar devices, their CMOS technologies are commercially the same.[35] In addition to being fabricated in a new manufacturing process, the second batch of SVADC-1s were housed in a new package with a smaller chip cavity.[36]

[33]Specifically, the SVADC-1 is clocked at 100 kHz (the sampling frequency of the WIPER instrument) with both inputs set to 1.25 V DC (i.e., 0-V differential DC input). Reference voltages are generated internally (by connecting $V_{\mathrm{REF}+}$ to $V_{\mathrm{SENSE}+}$, and $V_{\mathrm{REF}-}$ to $V_{\mathrm{SENSE}-}$, in Figure 5.26), and the reference current set by an external resistor (see Figure 5.27). See Section K.3 for more information.

[34]The gain and SFDR is measured at multiple input frequencies: the aggregate gain and SFDR collapse these measurements into a single metric. See Section K.5 for the definitions of these aggregate quantities.

[35]The only circuit redesign involved in the migration from BiCMOS8iED to BC8BPLUS is a change in the RC circuit used to trigger the ESD clamp circuitry (see footnote 32 of Chapter 5): the RC time constant is changed from 3.8 μsec in the BiCMOS8iED version of the SVADC-1 to 6.3 μsec in the BC8BPLUS version.

[36]The first batch SVADC-1s were found sensitive to bondwire failures during vibration testing.

The second batch of SVADC-1s was composed of 16 SVADC-1s manufactured in the BC8BPLUS manufacturing process. All proved functional after burn-in testing. In terms of performance, while outliers exist, after burn-in the vast majority showed changes of at most 0.5% in power consumption, 0.5% in aggregate gain, 2% in aggregate SFDR, and 3% in SINAD compared to their before burn-in values (see Section K.5 for the complete data tables).

In all, the burn-in results further demonstrate the robustness of the SVADC-1.

6.5 Conclusion

This chapter has presented measured performance results for the SVADC-1. These results showed the effectiveness of the self-calibration algorithm in measuring, and then eliminating, analog-digital mismatch between the analog stages and digital reconstruction. The result is a more uniform quantization, as evidenced by improvements in the INL from -3.52 LSB and $+2.47$ LSB before self-calibration to -0.76 LSB and $+0.84$ LSB after self-calibration. More uniform quantization in turn translates to improved converter linearity, as proven by improvements in the peak SFDR to at least 90.9 dB over a wide range of input frequencies after self-calibration. In particular, at input test frequencies of 9.99093-kHz and 366.007-kHz, the peak SFDR is just 75.9 dB and 83.4 dB (respectively) before self-calibration, but improves to 91.4 dB and 96.4 dB (respectively) after self-calibration.

In addition, this chapter presented measured results of the radiation performance of the SVADC-1, as assessed by both total-dose testing (by 50-MeV protons up to 2 Mrad(Si)), and single-event testing (by 10-MeV/nucleon heavy ions up to 58.78 MeV-cm^2/mg, by 25-MeV/nucleon heavy ions up to 63 MeV-cm^2/mg, and by pulsed laser).

In total-dose testing, the SVADC-1 showed little change in performance up to the full 2 Mrad(Si) tested, demonstrating the utility of the radiation-hardness-by-design techniques employed in its design. For example, the SVADC-1 exhibited a variation

Hence for the second batch the bondwires were shortened by adopting a package with a smaller chip cavity.

in peak SNDR of only 0.2% dB over the full 2 Mrad(Si) tested, showing that output droop from leaky switches is effectively eliminated by the adoption of the enclosed terminal layout. The efficacy of the nonstandard layout is also apparent in the power consumption, which rises with dose in the supplies connected to logic gates wherein the enclosed terminal layout is not used, but remains steady in the supplies connected to logic gates wherein it is used. Overall, the SVADC-1 maintains its high linearity performance up to very high doses, achieving a peak SFDR of at least 90.1 dB up to 1 Mrad(Si), and of at least 88.2 dB up to 2 Mrad(Si), after self-calibration.

In single-event testing, the SVADC-1 showed no latchup under a variety of radiation sources, including under 25-MeV/nucleon heavy-ion testing up to an LET of 63 MeV-cm^2/mg at elevated supply (2.7 V) and temperature (131°C), affirming the efficacy of the latchup prevention techniques. In addition, the soft error rate of the SVADC-1 was characterized by a Turflinger analysis, and the consequent Weibull parameters presented.

The measured performance of the SVADC-1 after self-calibration is summarized in Table 6.5.

Characteristic	Baseline performance	Radiation performance[a]
Technology	0.25-μm CMOS, 1P5M, non-epitaxial substrate	
Supply voltage	2.5 V	
Sampling rate	5 MS/s	
Input range	2 V differential (1 V_{PP})	
Resolution	$\geq 2^{12}$ output levels (16-bit output word)	
Area (ADC)	9.0 mm^2 (5.2 mm^2 core)	
Peak SFDR (assessed at 100 Hz/bin)		
$f_{in} = 9.99093$ kHz[b]	91.4 dB	$\geq$90.1 dB, $\geq$88.2 dB[d]
$f_{in} = 366.007$ kHz[c]	96.4 dB	
Wideband[e]	$\geq$90.9 dB	
Peak SNDR		
$f_{in} = 9.99093$ kHz[b]	69.9 dB (11.3-bit)	$\geq$70.6 dB (11.4-bit)
$f_{in} = 366.007$ kHz[c]	70.8 dB (11.5-bit)	
Wideband[e]	$\geq$69.1 dB	
Peak SNR[f]	65.6 dB	
INL (assuming 12-bit LSB)		
Minimum	-0.76 LSB	≥ -0.76 LSB
Maximum	$+0.84$ LSB	$\leq +0.89$ LSB
Power consumption		
Analog portion	48.8 mW	$\leq$60.2 mW, $\leq$60.5 mW[d]
Digital portion[g]	10.3 mW	$\leq$10.3 mW
Radiation testing		
Total dose[h]	Up to 2 Mrad(Si)	
Latchup[i]	None up to 63 MeV-cm^2/mg at 131°C, 2.7 V supply	
Soft error rate[j]	By Turflinger analysis, see Table 6.3	

[a] Unless otherwise noted, worst case performance over 2 Mrad(Si) given.
[b] Assessed over 100-Hz to 1-MHz evaluation bandwidth.
[c] Assessed over 100-Hz to 1.2-MHz evaluation bandwidth.
[d] First value is performance up to 1 Mrad(Si), second value up to 2 Mrad(Si).
[e] Wideband quantities evaluated over input frequencies from 103.712 Hz to 366.007 kHz.
[f] Assessed over 0-Hz to 2.5-MHz evaluation bandwidth with 366.007-kHz input.
[g] Dynamic power consumption assuming Actel RTAX250S implementation.
[h] Total-dose testing by 50-MeV protons. Highest tested dose given.
[i] Latchup testing by 25-MeV/nucleon heavy ions. Highest tested LET given.
[j] Soft error rate testing by 10-MeV/nucleon heavy ions.

Table 6.5: Summary of measured performance, after self-calibration, of the SVADC-1.

Chapter 7

Conclusion

Modern advances in satellite technology have enabled the next generation of plasma wave instruments to take advantage of increasing amounts of available on-board digital computation and storage. In these new instruments, the objective of the analog electronics is to simultaneously capture as much of the frequency and power ranges of the signal as possible, relying on digital signal processing to balance issues of data storage and telemetry. Naturally, since these instruments are flown aboard satellites, they must also be low power and radiation tolerant.

This dissertation presented the development of the SVADC-1, a high-fidelity, radiation-hard, pipeline ADC developed for use in these new instruments. To achieve high fidelity, the SVADC-1 uses a novel self-calibration technique that combines DAC differencing with alterations to the stage transfer function. This technique measures the discontinuity heights of the transfer function implemented by the analog stage and uses these measurements to adjust the equivalent values in the digital reconstruction, ultimately correcting for analog-digital mismatch and mitigating the dominate contributor to pipeline converter nonlinearity.[1] To achieve radiation hardness, the SVADC-1 uses a commercial process coupled with a philosophy of radiation-hardness by design. Examples of specific radiation-hardness-by-design techniques include the use of nonstandard enclosed-terminal layouts for select NMOS

[1] Another contributor to converter nonlinearity—operational amplifier nonlinearity—is assessed by a novel two-gain technique.

devices, the use of self-resetting architectures for amplifiers, the extensive use of guard rings to protect against latchup (important given the non-epitaxial substrate of the SVADC-1), and the overdesign of select parameters of particular circuits to compensate for radiation degradation. However, it should be noted that the radiation-hardness-by-design philosophy adopted by the SVADC-1 is broader than just these techniques, and influenced many other design decisions as well.

The power consumption of the SVADC-1 is kept low by a careful consideration of the converter specifications. Given the spectrographic nature of plasma wave analysis, as well as the expected spectra of the waves themselves, SFDR is the key specification for the SVADC-1: to capture signals over the 100-Hz to 1-MHz bandwidth of interest, an SFDR of at least 90 dB (assuming 100-Hz/bin spectral resolution) is needed. A combination of theoretical and numerical studies show that this SFDR can be achieved by a 12-bit converter with a circuit-noise-limited (full bandwidth) SNR of only 58–66 dB. Both the small number of bits and the relatively low SNR (which corresponds to only 9.3-bit ENOB) of this converter aid in decreasing power consumption. These power savings are especially important given the overdesign incorporated through radiation-hardness by design. The result is that the power consumption of the SVADC-1 is such that a complete plasma wave instrument of 6 wideband channels incorporating the SVADC-1 has power consumption competitive to the current state-of-the-art instruments.

To demonstrate the efficacy of all these techniques and analyses, an SVADC-1 chip was fabricated in a commercial, 0.25-μm, 1P5M, single-well CMOS process that uses STI interdevice isolation and a non-epitaxial substrate. Baseline measurements confirm the utility of the self-calibration technique: the SVADC-1 chip achieves a wideband peak SFDR of at least 90.9 dB while sampling at 5 MS/s and consuming just 48.8 mW. And radiation measurements—including both total-dose (by 50-MeV protons) and single-event (by 10-MeV/nucleon heavy ions, 25-MeV/nucleon heavy ions, and pulsed laser) testing—confirm the utility of the radiation-hardness-by-design philosophy. In particular, the SVADC-1 chip maintains at least 90.1-dB SFDR while sampling at 5 MS/s and consuming at most 60.2 mW up to a total dose of 1 Mrad(Si), experiencing a slight performance loss to 88.2-dB SFDR and 60.5 mW at 2 Mrad(Si)

(the highest tested dose). In addition, the SVADC-1 displays no latchup up to an LET of 63 MeV-cm^2/mg (the highest tested LET) at elevated temperature (131°C) and supply voltage (2.7 V).

To place the performance of the SVADC-1 in perspective, Table 7.1 lists the currently available, high-fidelity, radiation-tolerant ADCs as originally presented in Table 3.2, but now including the SVADC-1. In comparison to the other converters, not only does the SVADC-1 meet all the specifications, but it does so to a much higher tested radiation dose.[2] It also does so at a much lower power although, to be fair, it should be noted that the listed power consumption includes only the power consumption of the analog portion of the SVADC-1. Nonetheless, it should be clear that the SVADC-1, and the techniques employed in its design, are applicable wherever high-fidelity, wideband, radiation-hard analog-to-digital conversion is required.

7.1 Suggestions for Future Research

The SVADC-1 chip manufactured for this dissertation implemented the analog portion of the SVADC-1 converter. Future research, then, could integrate the digital portion of the SVADC-1 converter, either by including the digital design on the same die or by fabricating a separate digital die and bonding the two chips in the same package. Such a design could use the Verilog code described in Section 6.4.1 as a starting point, although the design would conceivably be optimized for implementation in an integrated circuit environment. In addition, the digital design may have to be reconsidered in light of the requirement that it be radiation-hard. Radiation-hardness by design for digital designs is currently an area of active research, fostered by observed digital upsets in chips manufactured in deep-submicron processes in even terrestrial environments [*Dodd and Massengill*, 2003; *Seifert et al.*, 2002]. Readers interested in radiation-hard digital design—in particular, upset-tolerant digital design—are recommended to the dissertation of fellow Ph.D. student

[2] It is worth noting that most commercial radiation-tolerant parts are only tested to 300 krad(Si). However, the performance of the SVADC-1 up to 1 Mrad(Si) is still quite remarkable, and represents a significant increase in the lifetime of the SVADC-1 in actual radiation environments over the original 100 krad(Si) specification.

Bits	Manufacturer	Part	Power	Rate	SFDR	SEL	TID
			[mW]	[MS/s]	[dB]	[MeV-cm^2/mg]	[krad(Si)]
	Required specifications		60	5	90	No latchup	100
≥12	SVADC-1		60	5	90	No latchup	1000
			61	5	88	No latchup	2000
12	Honeywell	RH9225	240	20	85	No latchup	300
14	Analog Devices	AD6644	1300	65	92	No latchup	>100
14	Analog Devices	AD6645	1500	105	93	No latchup	>100
14	Analog Devices/Maxwell	7871RP	50	0.083	88	No latchup	30–100
16	Linear Technology	LTC1604	220	0.333	93	55–70	100
16	Texas Instruments/Maxwell	7809LPTRP	150	0.1	100	No latchup	100

Table 7.1: Performance summary of currently available, radiation-tolerant ADCs, including the SVADC-1 (with performance listed up to 1 Mrad(Si) and 2 Mrad(Si)). Required converter key specifications included: value highlighted in red if it does not satisfy specification. Culled from the most recent JPL A/D Selection Guide [*JPL*, 2005].

Hsiao-Heng K. Lee, which describes application of radiation-hardness by design to an FFT processor [*Lee*, in preparation].

In addition to integration of the analog and digital portions of the SVADC-1 converter, the SVADC-1 converter can also be integrated with the SVLNAE-3—the LNA/AAF chip designed by fellow Ph.D. student Benjamin J. Mossawir [*Mossawir et al.*, 2006; *Mossawir*, in preparation]—to create a plasma wave receiver on a chip. Such research would have to account for not only concerns of drive between the AAF and ADC,[3] but also system-on-a-chip integration issues (such as coupling of digital switching and clock noise into sensitive analog circuits).

As for the SVADC-1 converter architecture itself, further research into decreasing the power consumption, and/or increasing the SFDR, should be considered. Towards the former, research into better scaling and optimization of the pipeline stages—especially in light of the digital self-calibration technique's ability to mitigate capacitor-mismatch concerns, allowing smaller signal capacitors—remains intriguing. Towards the latter, the results of Table 4.1 suggest that higher SFDR can be accomplished by increasing the number of bits in the uncalibrated ADC, for example by increasing the resolution of the terminating ADC or adding backend pipeline stages.[4] In addition, exploration of other high-fidelity ADC architectures (such as the $\Sigma\Delta$-modulator) is of interest, especially given the radiation-hardness results of the SVADC-1. Indeed, in general the results of the SVADC-1 show that, with proper radiation-hardness by design, it is possible to manufacture a wide array of radiation-hard circuits in commercial manufacturing processes. Naturally, the extension of these techniques to other circuits—such as phase-locked-loops (PLLs),[5] mixers, and

[3] In current board-level designs involving the SVLNAE-3 and SVADC-1, an external buffer is used to provide adequate drive strength of the switched-capacitor ADC input, as well as limit backdrive of the ADC sampling transient into sensitive front-end circuitry; see Appendix G for an example.

[4] Although it is notable that at SFDRs beyond ~100 dB, the linearity of the sampling circuit becomes limiting (see Section 5.4.1.1).

[5] As discussed in footnote 4 of Chapter 6, inclusion of an on-chip PLL for the SVADC-1 is desirable. Of course, the PLL would have to be radiation-hard. Such a design would thus most probably incorporate many of the techniques presented in this dissertation, as well as develop many additional techniques specialized for the particulars of the PLL architecture. Also beneficial would be inclusion of more robust on-chip generation of the reference voltages and master current source, and that in a radiation-hard manner.

references—offers an open avenue of research.

Appendix A

Introduction to Radiation Testing

This appendix presents a general introduction to the radiation testing of integrated circuits. Radiation testing is typically divided into two classes: single-event testing and total-dose testing. In the former, the device-under-test (DUT) is assessed for short-term effects, often stemming from a single radiation strike. In the latter, the DUT is assessed for long-term effects due to accumulated radiation damage. The physics underlying both effects are discussed in more detail in Chapter 2, particularly Sections 2.2 and 2.3, respectively. Instead, this appendix focuses on testing methods for measuring these effects.

This appendix begins with terminology, presenting definitions for linear energy transfer (LET) and dose. With the nomenclature established, it then turns to total-dose and single-event testing themselves. The radiation sources and procedures for each are described. In addition, some practical concerns are discussed. Examples of these principles in practice can be found in Appendices B and F (for total-dose testing) and Appendices C, D, and E (for single-event testing), which describe the particular radiation testing experiments performed in support of this dissertation.

All the test procedures described here (and all the radiation testing experiments undertaken for this dissertation) substantially conform to Military Standard 883G, Test Standard Method, Microcircuits [*MIL-STD-883G*, 2006], referred to simply as the Military Standard throughout this dissertation. Notably, the Military Standard is intended for qualification of commercial parts. This dissertation, though, assumes

a research—and not product certification—setting: the objective is more to measure and understand performance changes than to certify chip yields. Thus, on the one hand, certain procedures and reasonings of the Military Standard are not strictly applicable. On the other, the reader is cautioned that research results should not be substituted for certification qualifications *prima facie*.

A.1 Terminology

In both total-dose and single-event testing, the DUT is irradiated by a radiation source and its response measured. Ideally, the radiation source mimics an anticipated radiation environment. Exactly recreating the anticipated environment, though, is usually impractical; for example, the near-Earth space environment of the SVADC-1 is composed of dozens of particle species of differing energies and fluxes [*Rasmussen*, 1988]. Instead, radiation testing mimics the environment with a series of proxy radiation sources that recreate the dominant radiation damage mechanisms, hence recreating the essential physics. This section briefly introduces the terminology used to describe these sources.

A.1.1 Linear Energy Transfer

In single-event testing for aerospace applications, the metric of linear energy transfer, or LET, is often used to characterize incident radiation. Broadly, LET is the amount of energy lost by an incident particle per unit distance traveled in a target material.[1] LET is target material dependent: this dissertation adopts the standard aerospace practice of assuming silicon as the target material. This dissertation also adopts the standard aerospace practice of normalizing LET by the target material density: LET values are thus given in units of MeV-cm^2/mg.

LET is a simplified classification that characterizes a wide range of incident

[1]Note that LET focuses on the energy loss of the incident, but does not necessarily describe the eventual absorption of that energy by the target material [*Koga*, 1996]. Instead, the absorbed energy is described by the metric of total absorbed dose, as described in Section A.1.2.

ionizing particles of different species and energies under a single metric.[2] As such ionizing radiation passes through a target material, it imparts energy to the material through a complex series of radiation-matter interactions. The vast majority of these processes lead to electron excitation and atomic ionization, with the result a track of electron-hole pairs in the wake of the incident ion the shape and concentration of which depends on the properties of the incident particle (such as mass, energy, and angle of incidence) and the target material [*Koga*, 1996]. As these pairs migrate (subject to local structures, materials, and electric fields), both soft and hard errors can occur. LET encapsulates these processes, essentially assuming that different ionizing particles, so long as they have the same LET, have the same ultimate effect on the single-event performance of the DUT [*Petersen et al.*, 1992]. This broad claim is supported by early experiments that showed that, at least for charged particles, the "spectrum of energy losses in primary collisions is to a good approximation the same for various types and speeds of charged particles" [*Neary*, 1969], and thus "a useful simple index of radiation quality is the amount of energy lost per unit length of track, the linear energy transfer" [*Neary*, 1969]. The definition has since often been refined to restrict the LET to consider only energy loss of the incident particle due to atomic ionization and electron excitation in the target material.[3]

Of course, the LET concept does have limitations. First, there are caveats to the simple expansion of LET to different types of radiation. For example, heavy-ion and proton sources of the same LET may nonetheless generate different error rates. This discrepancy arises since heavy ions cause ionization and excitation directly, while protons cause ionization and excitation more through secondary processes: as LET concentrates on the energy loss of the incident particle, it does not capture this secondary radiation damage [*Rasmussen*, 1988].[4] Second, as perhaps anticipated

[2]Many single-event effects such as latchup and nodal upsets are primarily due to radiation-induced charge in the silicon substrate. This charge is caused by energy imparted by the incident radiation to the target material through the excitation of electrons and ionization of atoms. Such incident particles are thus called ionizing particles, and such incident radiation is called ionizing radiation.

[3]A more general metric that ignores this restriction of energy loss mechanisms is stopping power, which includes all energy loss mechanisms by the incident particle. However, traditionally the aerospace community has preferred to concentrate on radiation-induced ionization processes, and hence has favored LET.

[4]As a result, "devices with LET thresholds well above the LET of protons may nevertheless

by its concentration on ionizing radiation effects, LET does not well-characterize displacement damage: particles with the same LET do not necessarily cause the same displacement damage [*Koga*, 1996].

Another important subtlety is that LET varies with the penetration depth of the incident ionizing particle into the target material. This variation is characterized by a Bragg curve: examples of Bragg curves for various heavy ions incident on silicon are shown in Figure A.1. Bragg curves typically show a peak in the LET—called the Bragg peak—before a precipitous fall as the incident particle comes to rest. The curves are strongly influenced by the energy of the incident particle, with higher energy particles usually yielding deeper Bragg peaks. Typically, for each ion the Bragg curve is characterized by a single LET value: for example, the surface LET (i.e., the LET at a penetration depth of 0) may be used. Note that, to ensure this selection is a valid characterization, the single-event sensitive regions of the DUT should lie at depths less than the Bragg peak. For most integrated circuits, these sensitive regions are indeed at low penetration depths near the surface of the silicon. However, in some cases the sensitive regions may lie deeper. Hence LET values should always be accompanied by a statement of the incident particle energy and type; for example, the particle energy in MeV/nucleon is typically given when presenting the results of testing by heavy ions.

Despite its limitations, LET has nonetheless retained wide acceptance and use in the aerospace community [*Koga*, 1996]. Indeed, as discussed in Section A.3.1, when combined with the concept of cross section, LET provides a useful way to estimate the error rates of a DUT in a given radiation environment. This dissertation thus adopts the concept of LET when characterizing both hard and soft single-event errors.

A.1.2 Dose

Whereas single-event testing assesses the impact of a single incident particle, total-dose testing concentrates on the cumulative effect of radiation over prolonged exposure. In these contexts the alternate metric of total absorbed dose is used to

experience high upset rates from proton environments" [*Rasmussen*, 1988].

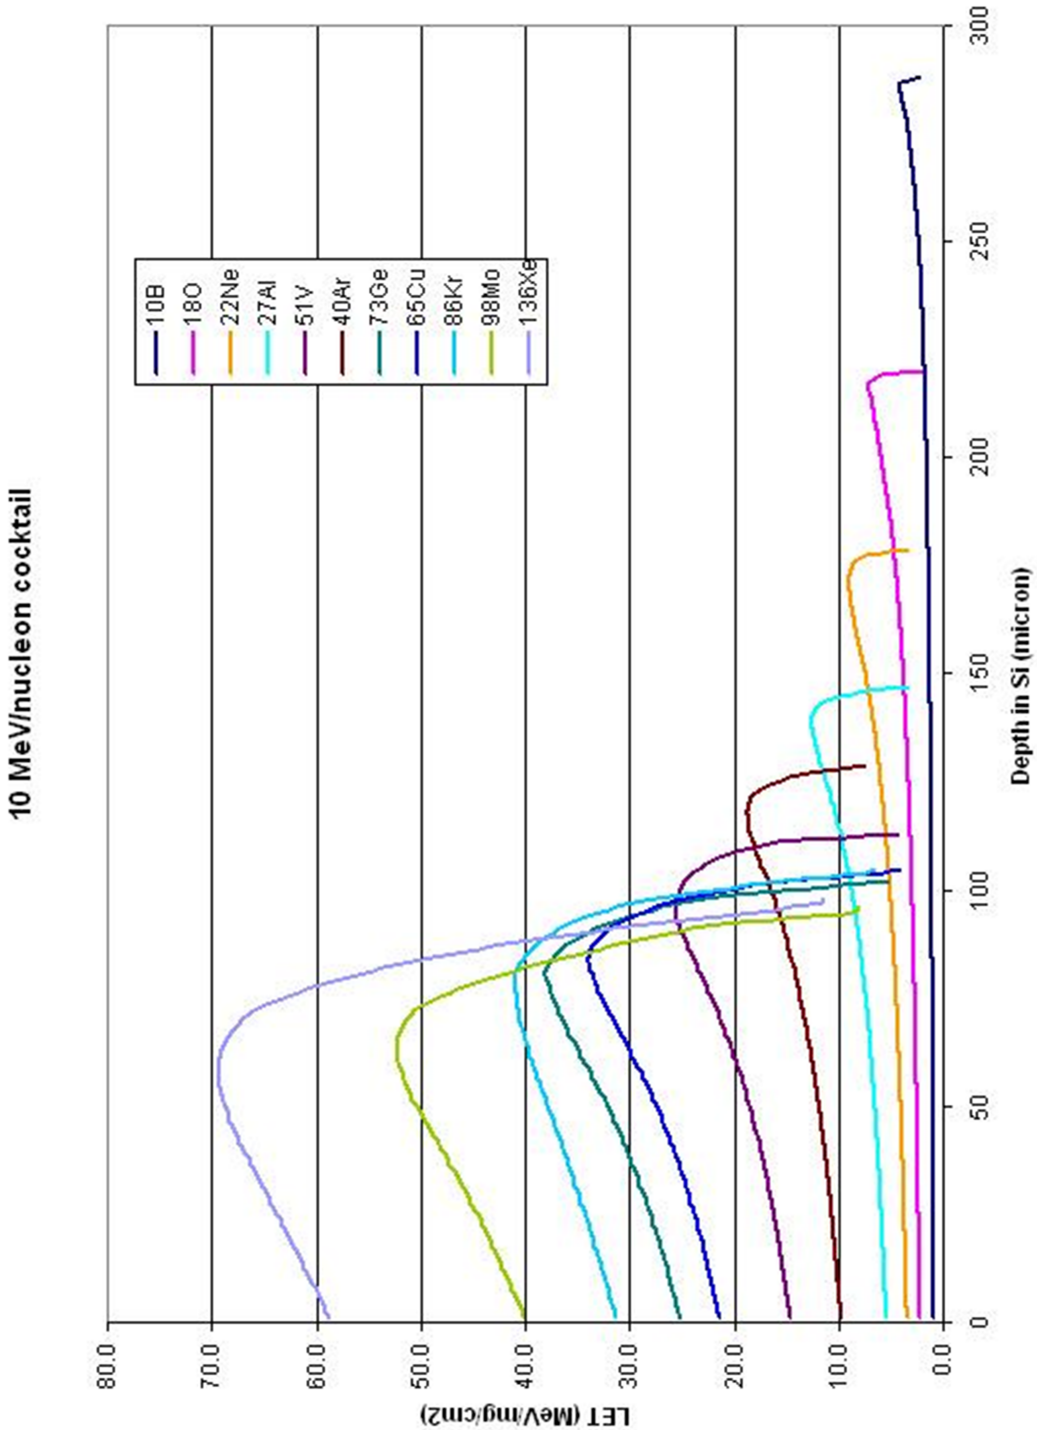

Figure A.1: Examples of Bragg curves for select heavy ions in silicon assuming normal incidence (i.e., the ion is incident at an angle perpendicular to the surface of the silicon). For further details about these heavy ions refer to Table A.1. Reproduced in total from *LBNL* [2009a].

characterize radiation.

The dose is defined as the amount of energy absorbed by the target material from a radiation source, normalized by the mass of the target material. That is, if an energy $\mathrm{d}E$ is imparted to a target of mass $\mathrm{d}m$, the resulting dose D is [*Kerris*, 1989, p. 461]:

$$D = \frac{\mathrm{d}E}{\mathrm{d}m} \tag{A.1}$$

Dose thus has units of energy per mass. It is typically given in units of rad, where [*Kerris*, 1989, p. 461]:

$$1 \text{ rad} = 100 \frac{\text{ergs}}{\text{g}} \tag{A.2}$$

The unit rad is often prefixed as if an SI unit.[5]

Notably, dose is material dependent: the same radiation source can be absorbed in different proportions by different materials. By convention, the absorbing material is included in parentheses after rad: for example, rad(Si) indicates the absorbed dose in silicon, while rad(SiO_2) indicates the absorbed dose in silicon dioxide. By convention, all total-dose radiation testing data measured in support of this dissertation is given in rad(Si).[6]

[5]Technically, the SI unit of absorbed dose is the Gray (Gy) [*NIST 811*, 2008, p. 5]:

$$1 \text{ Gy} = 100 \frac{\text{J}}{\text{kg}}$$

Conversion between rad and Gray is simple [*NIST 811*, 2008, p. 10]:

$$1 \text{ Gy} = 100 \text{ rad}$$

While the U. S. National Institute of Standards and Technology "strongly discourages" the use of the more senior unit rad [*NIST 811*, 2008, p. 10], rad is nonetheless still widely accepted and actively used throughout the aerospace community. It is thus adopted in this dissertation.

[6]The use of rad(Si) for characterizing the total-dose performance of silicon circuits is fairly standard throughout the aerospace community. The alternate unit rad(SiO_2) finds more use in studies of specific devices and structures: see footnote 35 of Chapter 2.

A.2 Total-Dose Testing

With the terminology established, this appendix now turns to radiation testing, beginning with total-dose testing. In total-dose testing the DUT is irradiated by a radiation source until it has accumulated a desired dose. The DUT performance at this dose is then characterized. If additional dose steps are desired, the DUT is again irradiated by the source and the process repeated. In this way the DUT performance is charted as a function of the total accumulated dose. Historically, the dose steps are logarithmically spaced.

The following subsections address this process in more detail, including considerations of both the radiation sources used and test procedures executed, sprinkled with some general advice gleaned from the author's own experience.

A.2.1 Total-Dose Sources

Total-dose testing is typically accelerated testing: whereas the DUT often accumulates a total dose over the course of months or years in its anticipated environment, in radiation testing the same dose is delivered in just hours or days. Hence the radiation source used must be capable of delivering high dose rates.[7]

A popular radiation source for total-dose testing is Cobalt-60 (^{60}Co), which produces γ-ray radiation [*Kerris*, 1989, pp. 454–455]. Produced artificially by neutron irradiation of ^{59}Co in nuclear reactors, ^{60}Co decays via β^- emission to a metastable state of ^{60}Ni. The decay half-life is 5.26 years. The metastable ^{60}Ni then instantaneously decays to its ground state by emission of a 1.17-MeV γ-ray and

[7] Notably, while high dose rates are generally desirable in total-dose testing of CMOS technologies, they are not necessarily desirable for other technologies. For example, bipolar devices are known to exhibit enhanced radiation degradation at low dose rates, dubbed enhanced low dose rate sensitivity, or ELDRS [*Pease*, 2003]. Hence circuits from bipolar technologies may be radiation tested at low dose rates, for example, as prescribed by Condition D of Method 1019.7 of the Military Standard [*MIL-STD-883G*, 2006]. (Although recent research has shown that ELDRS can be eliminated by altering the top passivation layer [*Shaneyfelt et al.*, 2006], which may have implications on the necessity of low-dose-rate testing of such devices in the future.) Recently, there has been some suggestion that CMOS technologies may also display such behavior, however, not all CMOS technologies tested display such effects [*Witczak et al.*, 2005]. For this dissertation, then, only accelerated, high-dose-rate testing is undertaken.

a 1.13-MeV γ-ray.[8] These γ-rays are the desired radiation: as they pass through integrated circuits, γ-rays create ionization damage through a variety of physical processes.[9] From a practical perspective, ^{60}Co sources can deliver dose rates as high as Mrad(Si) per hour [*Kerris*, 1989, p. 455] and achieve very good penetration.[10] In addition, compared to other sources that require linear accelerators or cyclotrons (such as the proton beams discussed below), ^{60}Co sources suitable for radiation testing can be had with relatively small commitments in infrastructure, making them more easily accessible.

A ^{60}Co source creates ionization damage but very little displacement damage. An alternative that creates both is a proton source: as protons pass through integrated circuits, they not only create ionization damage, but also create displacement damage through collisions with atomic nuclei. Proton beams are commonly generated by cyclotrons, and hence are often less readily available. Depending on the particular test facility, proton beams in energies ranging from single MeV through hundreds of

[8] In practice, the actual γ-ray spectrum is somewhat more complicated owing to scattering within the ^{60}Co source itself, to the source encapsulation containers, to the chamber housing the entire test setup. An introduction to the effects of these source nonidealities is given in *Kerris* [1989, pp. 454–455]. In the end, it is always advisable to measure the actual dose delivered (instead of relying on computations from first principles) before beginning total-dose testing.

[9] For photons in general, the primary radiation-matter interactions are the photoelectric effect, Compton scattering, and pair production [*Srour and McGarrity*, 1988]. All result in the generation of liberated electrons (and, in the case of pair production, positrons), while some also result in the emission of secondary radiation. Briefly:

Photoelectric effect. The incident photon is completely absorbed by an electron, which is then liberated. If the liberated electron is from an inner shell, an outer shell electron can drop into the vacated energy state, emitting either a characteristic X-ray or low-energy Auger electron (depending on the atomic number of the target atom) as it does so.

Compton scattering. Similar to the photoelectric effect, the incident photon is completely absorbed and liberates an electron. However, the incident photon also causes the emission of a lower energy photon which then continues on through the material.

Pair production. The incident photon is completely absorbed and causes formation of a positron-electron pair (the positron has the same rest mass and charge as an electron, but is positively charged).

Which process occurs depends on the energy of the incident photon, the target material, and the attachment of the electron. However, in silicon, the photoelectric effect dominates for photon energies less than ∼50 keV, whereas pair production dominates for photon energies greater than ∼20 MeV. For energies in between (as is the case for ^{60}Co-generated γ-rays), Compton scattering dominates [*Evans*, 1955, p. 712].

[10] Indeed, so much so that boards populated by DUTs can be stacked during irradiation so that the γ-rays pass through multiple DUTs [*Holmes-Siedle and Adams*, 2002, p. 439].

MeV can be had [*Holmes-Siedle and Adams*, 2002, p. 447]: the particular energy desired depends on the application and the desired dose rate. Dose rates can again be very high, sufficient for achieving Mrad(Si) on hour time scales.

Other radiation sources for total-dose testing include pulsed X-rays and electron beams. Readers interested in a more comprehensive overview are recommended to Section 13.2 of the *Handbook of Radiation Effects* [*Holmes-Siedle and Adams*, 2002].

A.2.2 Total-Dose Procedure

Guidelines for total-dose testing procedures are given in Method 1019.7 of the Military Standard [*MIL-STD-883G*, 2006]. Total-dose testing involves irradiating the DUT by a radiation source to accumulate a desired dose and then assessing its performance at that dose. The process is thus readily decomposed into two tasks: irradiation (wherein the DUT is irradiated) and characterization (wherein the DUT performance is measured). Typically, the DUT is not irradiated during characterization, but rather removed from the radiation source altogether.[11]

In irradiation, the DUT is exposed to the radiation source. During this time, the DUT should be biased and functioning, though it need not be performing to its full specifications. This distinction is important as it means the test board used during irradiation need not be exceptionally sophisticated: not only can it use simpler, more radiation-hard components (such as passives), but also it can use less overall hardware (desirable since in many irradiation systems—such as ^{60}Co systems—the test board is placed in a small, cramped space). Technically, DUT monitoring during irradiation is not required. In practice, though, simple checks of the DUT bias are usually performed immediately before irradiation to confirm the bias.

In characterization, the DUT performance is assessed by a batch of custom-designed tests. As the DUT need not be irradiated during characterization, the experimental setup need not be radiation tolerant and can be as sophisticated as

[11] Indeed, often irradiation and characterization occur in separate physical rooms. In some special cases, DUT performance may be assessed while the DUT is still under irradiation. However, since performing such testing can be quite difficult from a practical perspective, it is usually only done if strictly required.

needed. However, characterization is strongly time-constrained. To reduce annealing effects, Method 1019.7 of the Military Standard dictates that characterization begin at most one hour after the end of irradiation [*MIL-STD-883G*, 2006, Method 1019.7, Sect. 3.10]. Furthermore, if another dose step is desired, the next irradiation should begin at most two hours following the end of the previous irradiation [*MIL-STD-883G*, 2006, Method 1019.7, Sect. 3.10]. Thus characterization can run at most two hours. Practically, though, once DUT recovery,[12] DUT transport, measurement setup, and retest margin[13] are considered, characterizations are often designed to run much shorter. Under such conditions, it is often valuable to automate characterization as much as possible to both speed data acquisition and reduce operator error.

Typically, a baseline characterization is made before the onset of any irradiation and an anneal characterization is made well after the completion of all irradiation. For the latter, following the completion of all doses the DUT is left unirradiated to allow for annealing of the accumulated radiation damage. After this anneal, the DUT is characterized. In this dissertation, anneal testing is used as a diagnostic tool: since different radiation damage mechanisms tend to anneal at different rates, the anneal step can help assess their relative impacts. In such contexts, the design of the anneal—such as its duration and the bias of the DUT—is left to the discretion of the experimenter.[14]

Finally, a couple of practical points. First, given the short characterization time, often large amounts of data are rapidly acquired during characterization and then interpreted at a more leisurely pace after the completion of all irradiation. The only

[12] That is, recovering the DUT from the radiation source. This recovery can take non-trivial time: for example, in proton testing at the 88-inch cyclotron at Lawrence Berkeley National Laboratory, the entire test chamber is also irradiated during DUT irradiation. If the resulting test chamber radiation level is high enough, experimenters may need to wait up to an hour before they can safely enter the chamber and retrieve the DUT (see Section B.3.2 for an example).

[13] That is, time allotted for repeating portions of the characterization in case of errors.

[14] Method 1019.7 of the Military Standard does provide explicit procedures for two anneal tests. In the first, the anneal procedure is to improve yield by recovering devices that have experienced performance failure during radiation testing by annealing radiation-induced damage [*MIL-STD-883G*, 2006, Method 1019.7, Sect. 3.11]. Hence it is more appropriate for product certification than investigative research. In the second, the anneal procedure is to test for low-dose-rate effects [*MIL-STD-883G*, 2006, Method 1019.7, Sect. 3.12]. For this dissertation, though, low-dose-rate effects are not investigated (see footnote 7 of this appendix).

caveat here is to make sure that certain preliminary analytical results are available during characterization as these results can influence future doses: for example, if a parameter has begun to display large shifts the experimenter may wish to switch to finer dose steps. Second, typically different test boards house the DUT during irradiation versus during characterization: the DUT is thus often being moved back and forth between an irradiation bias board and a characterization board. To prevent physical damage to the DUT during these several populations and depopulations, use of a zero-insertion force (ZIF) socket on both boards should be given strong consideration.

A.3 Single-Event Testing

In contrast to total-dose testing where irradiation and characterization are largely separate, in single-event testing the two are more intertwined. Broadly, single-event testing involves monitoring and measuring the DUT during irradiation to catch instances of single-event effects. This collusion, conflated with the fact that often experimenters are often not allowed near the DUT during irradiation for safety reasons, introduces its own challenges.

This section begins by introducing the concept of cross section, which is used to characterize DUT error rates. It then turns to single-event testing itself, reviewing the procedures for measuring and analyzing the cross section and introducing common radiation sources used in such testing. To highlight many of the practical concerns involved, this section concludes by considering the design of heavy-ion testing experiments at the 88-inch cyclotron at Lawrence Berkeley National Laboratory.

A.3.1 Cross Section

Single-event testing seeks to predict the occurrence of hard and soft errors in the DUT for an anticipated radiation environment. For CMOS technologies in many environments (including near-Earth space environments) this prediction is usually predicated on measurements of the cross section of the DUT. In this method, it is

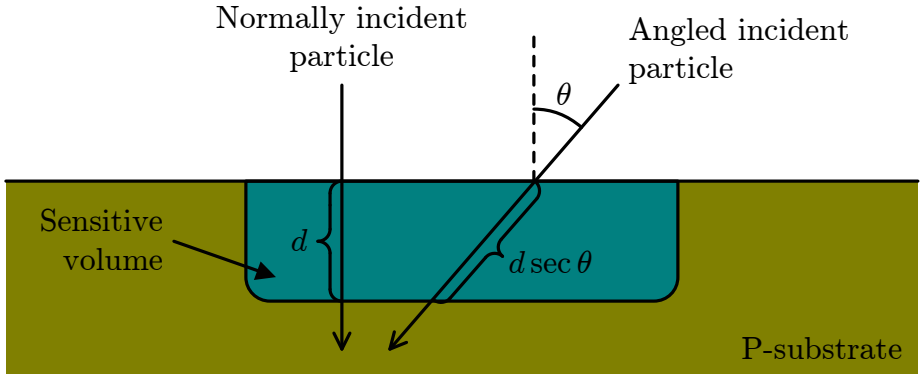

Figure A.2: Geometry of effective LET. If a particle traverses a total distance d through the sensitive volume (shown in teal) when normally incident, then it traverses a total distance $d \sec\theta$ when incident at an angle θ. Assuming the particle LET is constant over both distances, the angled incident particle thus deposits $\sec\theta$ times greater energy than the normally incident particle.

assumed that the DUT is composed of a collection of upset-sensitive volumes that trigger at some LET. The cross section then characterizes the total sensitive volume at a given LET [*Petersen et al.*, 1992].

To measure the cross section, the DUT is exposed to a radiation source of known LET and the number of errors counted. The definition of error is an open question and depends on the target application of the DUT. Naturally, the number of errors depends on the number of particles incident on the DUT during exposure. Hence, the raw number of errors is normalized by the beam fluence (the number of particles per unit area), which is the integral of the beam flux (the number of particles per unit area over a unit time) over the exposure time. As the result has units of area it is dubbed the cross section σ_{SEE}:[15] intuitively, σ_{SEE} characterizes the sensitive area of the DUT at a particular LET.

It is known that the angle of incidence of the incident particle often affects the measured cross section. This phenomenon is typically explained by noting that sensitive volumes are often shallow: assuming the particle penetration depth is much larger than the sensitive volume depth, then a particle passing through at an askew

[15] In most aerospace applications, the cross section is just denoted σ, but as σ is rather overloaded in this dissertation, the more descriptive σ_{SEE} is used here.

angle traces a longer trajectory through the sensitive volume as shown in Figure A.2. The longer trajectory translates to greater energy deposition in the sensitive volume and hence a greater effective LET [*Petersen et al.*, 1992]. Indeed, from the geometry of Figure A.2, the effective LET is:

$$\mathrm{LET}_{\mathrm{eff}} = (\sec\theta)\,\mathrm{LET}_{\mathrm{normal}} \tag{A.3}$$

for $\mathrm{LET}_{\mathrm{normal}}$ the LET under normal incidence (that is, the LET assuming the incident radiation is at an angle perpendicular to the surface of the target material). The relationship of Equation (A.3) is often exploited in single-event testing by intentionally angling the DUT with respect to the beam to boost the effective LET.[16]

A.3.2 Single-Event Procedure

To measure the cross section, the DUT is exposed to a radiation source of known LET and fluence, and the number of errors counted. This process is repeated over a range of LETs to measure the cross section at each LET. Typically, the data is then fit to a Weibull distribution. The Weibull distribution is often used in reliability analysis, as it well describes systems wherein a system failure can be modeled as a failure in the weakest link of many competing failure processes [*Petersen et al.*, 1992; *Tobias and Trindade*, 1986, pp. 70–72]. When used to characterize the cross section (σ_{SEE}) to LET (L) curve, the Weibull distribution is often parameterized as [*Petersen et al.*, 1992]:

$$\sigma_{\mathrm{SEE}} = A_0 \begin{cases} 1 - e^{-[(L-L_0)/W]^S} & , \quad L > L_0 \\ 0 & , \quad L < L_0 \end{cases} \tag{A.4}$$

[16] There are some limitations to the concept of effective LET, especially when the width of the sensitive volume is comparable to the width of the track left by the incident particle. In this case, grazing particles that only illuminate a small corner, say, of the sensitive volume may not trigger errors. Hence discontinuities sometimes occur when cross sections measured at LETs with normally incident particles, and cross sections measured at LETs with angled incident particles, are plotted together. This effect can be accounted for and corrected: see *Petersen et al.* [1992].

where L_0, W, S, and A_0 are the open parameters.[17] The Weibull fit gives the complete cross section as a function of LET and is used to predict the DUT error rate for an anticipated environment.[18] Separate Weibull fits can then be compiled for different types of soft errors to more precisely describe the overall DUT soft error performance.

Hard errors are typically measured in the same fashion, although for errors such as latchup, it is the occurrence, not the number, of the error that is assessed. In such cases, the result is not fit to a Weibull distribution, but instead the threshold LET is reported.

A.3.3 Single-Event Sources

A popular radiation source for single-event testing is a heavy-ion beam, that is, a beam of ionized elements such as boron or xenon. Heavy-ion beams are typically generated by cyclotrons such as the 88-inch cyclotron at Lawrence Berkeley National Laboratory, or the K500 superconducting cyclotron at Texas A&M University. These cyclotrons are fed by electron cyclotron resonance ion sources and provide a "cocktail" of heavy ions of different energies and LETs.[19] As an example, Table A.1 summarizes the

[17] Technically, Equation (A.4) represents the cumulative distribution function (CDF) of a Weibull distribution scaled by A_0 (that is, technically $\sigma_{\mathrm{SEE}}/A_0$ is a Weibull distribution). In Equation (A.4), though, A_0 is included as the equation describes a cross section, not a strict CDF. Regarding the parameters, L_0 is often called the threshold, W the width, and S the shape of the Weibull distribution. Furthermore, regarding S: it can be shown that if $S = 1$, then the Weibull distribution reduces to the exponential distribution; if $S = 2$, to the Rayleigh; and if $S = 4$, to the lognormal [*Petersen et al.*, 1992]. Finally, it is notable that this parameterization of the Weibull is neither unique nor universally adopted: different parameterizations are often used [*Tobias and Trindade*, 1986, p. 64].

[18] Conceptually, given a radiation environment with a spectrum of particles of known fluence and LET, the number of errors is the product of the fluence with the measured cross section at that LET. In reality, this calculation is more complicated, involving incident particle types (for example, in the case of heavy-ion or proton test results as previously mentioned), incident particle angles, and directional shielding around the DUT (and any possible secondary radiation created by the shielding materials). Hence accurate DUT error rate computations must generally take into account the three dimensionality of the sensitive volumes, the incident particle track, and the radiation environment [*Petersen et al.*, 1992].

[19] Broadly speaking, the electron cyclotron resonance ion source injects ions of various masses into the cyclotron. The cyclotron then acts as a mass analyzer, accelerating a selective mass depending on the frequency of the alternating electric potential [*Johnson et al.*, 2007]. The basic operation of a cyclotron is described in many standard college freshman physics texts (e.g., *Halliday et al.* [1992, pp. 741–742]). Further details on the ion source injection as implemented at the 88-inch cyclotron

Ion	Cocktail [MeV/nucleon]	Z	A	Charge State	Energy [MeV]	LET [MeV-cm^2/mg]	Range [μm]
B	4.5	5	10	+2	44.90	1.65	78.5
N	4.5	7	15	+3	67.44	3.08	67.8
Ne	4.5	10	20	+4	89.95	5.77	53.1
Ar	4.5	18	40	+8	180.00	14.32	48.3
Cu	4.5	29	63	+13	301.79	29.33	45.6
Kr	4.5	36	84	+17	387.08	38.96	48.0
Ag	4.5	47	109	+22	499.50	58.18	46.3
Xe	4.5	54	136	+27	602.90	68.84	48.3
Tb	4.5	65	159	+32	724.17	77.52	52.4
Ta	4.5	73	181	+36	805.02	87.15	53.0
Bi	4.5	83	209	+41	904.16	99.74	52.9
B	10	5	11	+3	108.01	0.89	305.7
O	10	8	18	+5	183.47	2.19	226.4
Ne	10	10	22	+6	216.28	3.49	174.6
Ar	10	18	40	+11	400.00	9.74	130.1
Cu	10	29	65	+18	659.19	21.17	108.0
Kr	10	36	84	+24	906.45	30.23	113.1
Ag	10	47	107	+29	1039.42	48.15	90.0
Xe	10	54	124	+34	1232.55	58.78	90.0
N	16	7	14	+5	233.75	1.16	505.9
O	16	8	17	+6	277.33	1.54	462.4
Ne	16	10	20	+7	321.00	2.39	347.9
Cl	16	17	35	+12	539.51	6.61	233.6
Ar	16	18	40	+14	642.36	7.27	255.6
Cu	16	29	63	+22	1007.34	16.53	190.3
Kr	16	36	78	+27	1225.54	24.98	165.4
Xe	16	54	124	+43	1954.71	49.29	147.9

Table A.1: Select heavy ions available at the 88-inch cyclotron at Lawrence Berkeley National Laboratory. Ions described by atomic number (Z), mass number (A), charge state, and energy (both per nucleon and total). LET and range[20] given assuming normal incidence on silicon. Reproduced from *LBNL* [2009b].

cocktails available at Lawrence Berkeley National Laboratory: a wide range of LETs is available. The cocktails are organized by the nucleon energy (given in MeV/nucleon): it is worth repeating that different ion energies yield different penetration depths and hence energies should be chosen appropriate for the particular DUT.[20] Finally, note that heavy-ion sources also accumulate total-dose damage over exposure: single-event testing should be careful to sufficiently limit DUT exposure to prevent conflating the two effects.

Other radiation sources for single-event testing include neutron sources[21] and laser sources. The latter has gained popularity over the years as it does not require the costly infrastructure of a cyclotron, but instead can be built in a standard optical laboratory setting. Laser testing uses energy deposition by picosecond laser pulses: the energy conveyed to the silicon causes atomic ionization and electron excitation. It has many advantages. First, it does not incur total-dose damage [*Moss et al.*, 1995]. Second, laser testing allows the experimenter to conveniently interact with the DUT during exposure—which is in contrast with particle-based single-event testing (see Section A.3.4)—alleviating many practical concerns. Third, given the small laser

at Lawrence Berkeley National Laboratory can be found in *McMahan et al.* [1986]. The injection technique has been subsequently expanded to increase the range of available LETs and ion energies: the curious are recommended to follow-up articles *McMahan et al.* [2004] and *Johnson et al.* [2007].

[20]To characterize the penetration depth, Table A.1 gives the "range" of the heavy ions. As a heavy ion travels through the silicon, it traverses a path depending on the individual interactions encountered along its traverse before coming to rest. This traverse can be characterized by the pathlength of the traverse, or by the distance between the end-points of the traverse (that is, the "crow's-flight" length of the traverse). For heavy particles which experience few elastic collisions, these two quantities approach each other. It is assumed, then, that Table A.1 follows the convention wherein "range" refers to the expected value of the distance between the end-points of the traverse projected along the extrapolated straight-line-trajectory of the incident particle.

[21]Neutron beams cause both direct and indirect ionization (the latter being secondary interactions following the principle radiation-matter interaction) [*Koga*, 1996]. In addition, they can cause displacement damage [*Holmes-Siedle and Adams*, 2002, p. 448]. Traditionally, neutron testing has been motivated by the fact that high-energy atmospheric neutrons (created by interactions of cosmic rays with ambient oxygen and nitrogen) are the main cause of single-event effects at avionics altitudes [*Normand*, 1996]. The neutron flux, though, rapidly diminishes at lower altitudes. However, given the ever smaller feature sizes of modern CMOS technologies, there have recently been concerns that the neutron flux at terrestrial altitudes, while small, is nonetheless sufficient to cause upsets [*Dodd and Massengill*, 2003; *Seifert et al.*, 2002]. As such, neutron testing is becoming more popular: for instance, the 88-inch cyclotron at Lawrence Berkeley National Laboratory has recently added neutron testing capabilities [*George et al.*, 2008].

spot size and the ability to synchronize the laser pulses with other test equipment in time, laser testing can pinpoint single-event sensitive volumes to sub-micron precision [*Moss et al.*, 1995] and nanosecond temporal accuracy [*Pouget et al.*, 2004]. However, laser testing also has its disadvantages. First, it cannot penetrate metallization.[22] Second, it cannot expose the entire DUT at a time.[23] Third, there is doubt as to how laser results correlate with the more widely accepted particle-based results.[24] If nothing else, then, laser testing remains a useful diagnostic with promising wider applicability.

A.3.4 A Practical Example

Particle-based single-event testing presents a host of practical challenges. To better describe these concerns, this section uses the example of heavy-ion testing at the 88-inch cyclotron at Lawrence Berkeley National Laboratory (LBNL). Many of the concerns addressed here, though, are more widely applicable to other single-event testing venues.

The floorplan of the 88-inch cyclotron is shown in Figure A.3. The ion beam is generated in the 88-inch cyclotron vault and then steered to one of several caves. Each cave supports different types of experiments: heavy-ion testing takes place in Cave 4B (in the lower left). The cyclotron is thus a shared facility, with valuable "beam

[22] A solution to this problem is backside illumination, where the laser energy is injected from the back of the DUT, that is, up through the substrate instead of down through the metal layers [*Lewis et al.*, 2001]. This technique circumvents the erratic shadowing of the metal layers. However, since the substrate absorbs laser energy, the energy delivered to the single-event sensitive regions (which favor the substrate surface where the transistors are fabricated) is diminished. An interesting solution to this caveat couples backside illumination with two-photon absorption. In two-photon absorption, the laser wavelength is trimmed to less than the semiconductor bandgap. At such wavelengths, the laser experiences little optical absorption at low intensities, but high optical absorption (owing to simultaneous absorption of two photons), and consequently strong electron-hole generation, at high intensities. Focusing the laser focal length to depths within the substrate, then, allows energy injection at the sensitive regions [*McMorrow et al.*, 2004].

[23] There are solutions to this problem as well, including scanning the laser across the DUT. Universality of DUT condition at different locations during such a scan can be guaranteed by proper temporal synchronization of the laser pulses with the DUT inputs and clocks [*Pouget et al.*, 2004]. However, given the micron size of laser spots, scanning through a DUT a few millimeters on a side can be rather slow.

[24] Although recent work suggests that the upset waveforms from both sources are the same, lending credence to the concept of a laser-source effective LET [*Buchner et al.*, 2004].

time" divided between a wide variety of different experiments and institutions. To maximize the use of the available beam time, experiments are scheduled back-to-back with no intervening downtime:[25] extra time is available only at the good graces of the next experimenter. Furthermore, note that if succeeding experiments use the same cave, then experiment set-up and tear-down time also consume precious beam time.

Given the time pressure, it is worthwhile to automate as much of the experiment as possible. In addition to data acquisition, this automation should also include data analysis of at least a preliminary sort. If nothing else, such analysis is usually needed to set the beam fluence: several preliminary exposures are done to determine a fluence that provides a statistically significant number of errors in a reasonable time. Longer exposures then follow for more formal measurements. Preliminary data products are also important in testing as they can provide insights into effects worth exploring: while experimenters enter single-event testing with notions of what to expect, unexpected effects are often unearthed. Naturally, studying these effects requires an experimental setup flexible both in how it conducts testing and in how it gathers and analyzes the consequent data. Single-event testing is thus often more dynamic than total-dose testing, with not only experiment parameters, but sometimes even experiment hypotheses, generated on-the-fly.

But perhaps the most obvious difference between single-event testing and total-dose testing is that in single-event testing the DUT is actively monitored and measured during exposure. These tasks are complicated by the fact that, for safety reasons, experimenters are not allowed near the DUT during irradiation. At LBNL, the DUT and accompanying test board are sealed inside a vacuum chamber at the end of the beamline.[26] During exposure, the vacuum chamber is evacuated and the cave sealed.[27] No experimenters are allowed inside the cave during irradiation. Instead,

[25] Quenching then restarting and retuning the cyclotron involves more than a little time and effort. Indeed, even steering the beam from one cave to another can take a few hours. Hence, aside from maintenance and emergencies, the cyclotron is run well nigh continuously.

[26] The vacuum chamber is indicated in Figure A.3 by the small black box in Cave 4B. The vacuum chamber is a cube about a yard in each dimension. The test board is typically mounted to a motorized arm set inside the cube: a camera system in the cave roof command center can then be used to center the DUT in the beam. Notably, these motors need to be actively run to maintain their position, and can thus emit unwanted electromagnetic interference.

[27] Note that physically changing configurations on the DUT test board thus requires depressurizing

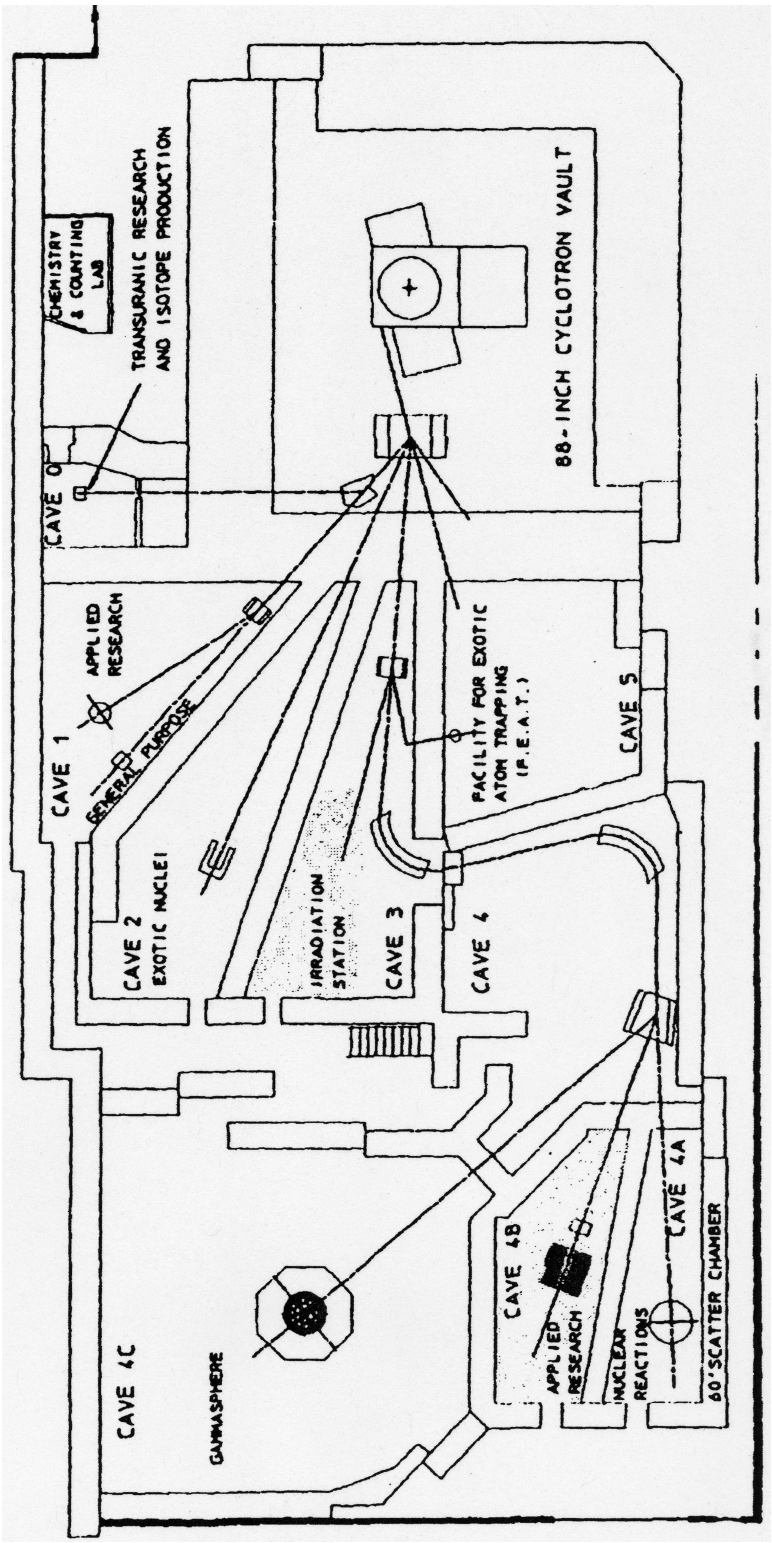

Figure A.3: Floorplan of the 88-inch cyclotron facility at Lawrence Berkeley National Laboratory (LBNL). Reproduced in total from *McMahan* [1999].

experimenters are housed on the cave roof. A command center on the cave roof enables experimenters to control the beam. Furthermore, a small hole in the cave ceiling allows signaling down into the cave, and headers built into the side of the vacuum chamber allow signaling to the DUT and test board. Experimenters thus set up control equipment on the cave roof and then command and measure the DUT remotely through a cabling distance of ∼30 feet all told, which can introduce noise problems for both analog and digital signals.[28] In addition, the headers on the side of vacuum chamber limit the number of analog and digital signals can be passed through to the DUT test board. Thus, when designing experiments for LBNL, control and measurement circuitry must be carefully partitioned between the test board, the cave, and the cave roof, all the while weighing the limited number of signals at each interface and concerns of noise and signal integrity.[29]

A complete single-event testing experiment must therefore accomplish the demands of automation, flexibility, and remote control, and do so given limited beam time. For an example of a particular solution to these challenges, the reader is recommended to Appendix C, which details the heavy-ion single-event testing of the SVADC-1 at LBNL. Finally, as a practical note it is always recommended that experimenters arrive early to complete as much experiment set up and baseline testing as possible, maximizing use of beam time for making actual measurements.

the vacuum chamber, and then repressurizing the chamber before recommencing irradiation. This cycles consumes valuable beam time, and thus such changes should be minimized.

[28]Noise abounds at LBNL: if nothing else, the cyclotron produces a large RF interferer. Perhaps more disconcerting, the noise environment is dynamic. LBNL houses a large amount of support machinery (such as an overhead gantry crane) as well as numerous other experiments that may or may not be operational at any given time. All of these sources can emit varying amounts of electromagnetic interference that can readily couple into long signal cables. Especially for measurement of low noise DUTs, experimenters must thus be prepared to adapt to a noise environment that can change day to day, if not shift to shift.

[29]Luckily, the heavy-ion beam is only about 3 inches in diameter and hence the test board can house active, radiation-soft command and control circuitry.

Appendix B

SVADC-1 Total-Dose Testing, 50-MeV Proton

This appendix describes total-dose testing[1] of the SVADC-1 converter by 50-MeV protons at Lawrence Berkeley National Laboratory (LBNL). This experiment seeks to determine the total-dose response of the SVADC-1 over a series of performance metrics, including power consumption and a variety of sinusoid-based metrics (e.g., SFDR). To this end, the SVADC-1 performance is assessed at logarithmically-spaced dose steps up to 2 Mrad(Si). Following the completion of irradiation, an anneal assessment is made after an unbiased, 65-hour anneal at room temperature.

The experiment was conducted on 8 and 11 February 2008 in Cave 4A of the 88-inch cyclotron facility at Lawrence Berkeley National Laboratory.[2] The experimenters were Charles C. Wang with Stanford University, Jeffery S. George and Rocky Koga with The Aerospace Corporation, and Michael B. Johnson with Lawrence Berkeley National Laboratory. The experiment was conducted under the aegis of Bernie Blake of The Aerospace Corporation.

[1] An introduction to radiation testing, including total-dose testing, is found in Appendix A.

[2] A floorplan of the 88-inch cyclotron, with Cave 4A labeled, is given in Figure A.3. *McMahan* [1999] gives a good introduction to radiation testing at Cave 4A of the 88-inch cyclotron.

B.1 Device-Under-Test

The device-under-test (DUT) is the converter portion of the SVADC-1. A single DUT is tested. The DUT is de-lidded throughout the experiment, including during both irradiation and characterization. All testing (including both irradiation and characterization) is performed at room temperature.

B.2 Irradiation

The DUT is irradiated by 50-MeV protons in Cave 4A of the 88-inch cyclotron at Lawrence Berkeley National Laboratory.

B.2.1 Irradiation Setup

During irradiation, the DUT is biased by a custom irradiation bias board designed by Mark Turpin at The Aerospace Corporation. The schematic of the irradiation bias board is shown in Figure B.1. It should be noted that this board was originally intended for use with the ^{60}Co γ-ray chamber available at The Aerospace Corporation's facility in El Segundo, California, and hence constructed to conform to the strict area and height restrictions imposed by that chamber.[3] While there are no such space restrictions in proton testing at Lawrence Berkeley National Laboratory, nonetheless, this origin explains many of the decisions made in the design of the irradiation bias board.

On the bias board, the DUT is housed in the zero-insertion force (ZIF) clam-shell socket described in Section 6.1.1. The board is powered by a 2.5 V supply provided by a Tenma 72-2010 DC power supply that both powers the DUT supplies and provides the DUT input voltage through two potentiometers trimmed to 1.75 V and 0.75 V for $V_{\text{in}+}$ and $V_{\text{in}-}$, respectively. The DUT is clocked by an Agilent 33250A outputting a 5 MHz square wave of 2.5 V_{PP} amplitude and 1.25 V offset. During irradiation, both pieces of equipment are placed inside the cave far from the proton beam, and

[3] Adding to the space constraints, the bias board was designed to accommodate up to nine independent DUTs simultaneously.

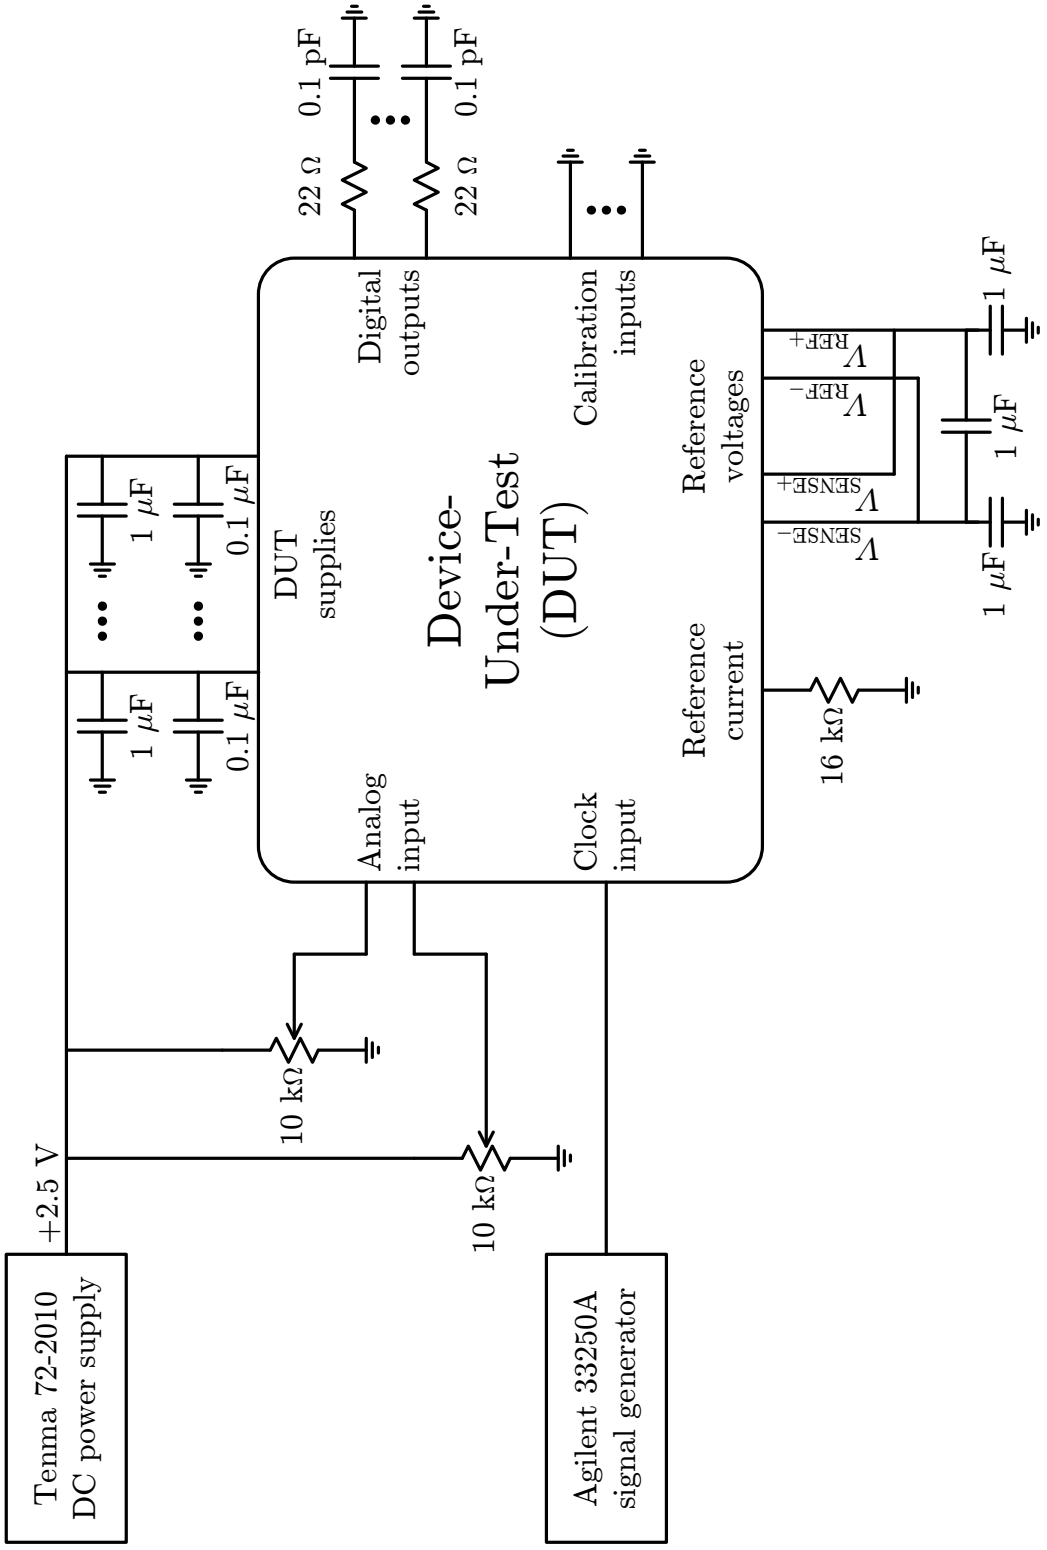

Figure B.1: Biasing of the DUT during irradiation for total-dose testing by 50-MeV protons.

further sheltered behind radiation-absorbing "yellow bricks", to reduce their radiation exposure.

The reference voltages are provided by internal generation to conserve board space, and the reference current provided by an off-chip resistor. Finally, the DUT is configured to operate normally (i.e., not in calibration) by tying the appropriate digital inputs to either ground or the supply, and the digital outputs loaded by simple series RC networks.

B.2.2 Irradiation Procedure

The DUT is irradiated by 50-MeV protons in logarithmic dose steps to a total dose of 2 Mrad(Si).[4,5] The complete irradiation schedule is given in Table B.1.

[4]Future experiments should consider using 30-MeV protons instead. Compared to 50-MeV protons, 30-MeV protons yield higher dose rates, reducing irradiation times. More significantly, 30-MeV protons result in less irradiation of the cave itself, reducing the delay between the end of irradiation and onset of characterization (see Section B.3.2). Lower-energy protons do have lesser penetration depth, but 30-MeV should still be high enough for integrated circuits such as the SVADC-1.

[5]The internal sensors within Cave 4A record the total fluence of the proton beam. For charged particles, the fluence ϕ can be related to a dose D as per [*Attix*, 1986, p. 188]:

$$D = \left(\frac{1}{\rho}\,\frac{\mathrm{d}T}{\mathrm{d}x}\right)\phi$$

where $(1/\rho)(\mathrm{d}T/\mathrm{d}x)$ is the mass stopping power of the incident radiation in the target material (ρ is the density of the target material, T the kinetic energy of the incident particle, and x the path length traversed by the incident particle in the target material). Note that this relation assumes that the mass stopping power is constant through the material, that scattering is negligible, and that the energy lost by the incident particles is fully absorbed by the target material. All of these assumptions hold to a good approximation for the case of high-energy protons incident on silicon. Furthermore, in this case the mass stopping power can be approximated by the LET L, yielding the relation:

$$D = L\phi$$

Given the proton energy (and accommodating for any energy losses through sensors before the beam is incident on the DUT), the LET of the beam can be determined by simulations, for example, by using the GEANT (GEometry ANd Tracking) program [*Holmes-Siedle and Adams*, 2002, p. 115], or the SRIM (Stopping Range of Ions in Matter) software package [*Ziegler et al.*, 2008]. Such simulations reveal that, for the 50-MeV proton beam at the 88-inch cyclotron:

$$D = (1.666\times10^{-7})\phi$$

where ϕ is in units of particles/cm^2, and D in units of rad(Si).

Date	Start time	End time	Exposure	Total fluence	Irradiated dose	Accumulated dose	Assigned dose
			[sec]	[protons/cm^2]	[krad(Si)]	[krad(Si)]	[krad(Si)]
08 Feb 2008	10:48	10:49	31	6.30×10^{10}	10.49	10.49	10
08 Feb 2008	11:34	11:46	133	2.50×10^{11}	41.71	52.20	50
08 Feb 2008	12:20	12:22	158	3.13×10^{11}	52.10	104.30	100
08 Feb 2008	13:14	13:19	313	6.25×10^{11}	104.18	208.48	200
08 Feb 2008	14:13	14:30	1013	1.88×10^{12}	313.16	521.74	500
08 Feb 2008	15:45	15:58	770	1.25×10^{12}	208.29	730.03	700
08 Feb 2008	17:03	17:23	1212	1.88×10^{12}	313.26	1043.29	1000
08 Feb 2008	18:33	18:47	848	1.25×10^{12}	208.29	1251.57	1200
08 Feb 2008	19:57	20:57	3617	5.00×10^{12}	833.14	2084.72	2000

Table B.1: Irradiation schedule for total-dose testing of the DUT by 50-MeV protons.

Dose step	2.5 V at supply	2.5 V on board	V_{CM}
[krad(Si)]	[V]	[V]	[V]
10	2.500	[n/a]	1.242
50	[n/a]	2.475	1.240
100	2.501	[n/a]	1.245
200	2.501	2.480	1.244
500	2.503	2.479	1.215
700	2.500	2.484	1.245
1000	2.500	2.455	1.231
1200	2.500	2.468	1.237
2000	2.504	2.441	1.224

Table B.2: Measurements taken before irradiation to confirm DUT functionality. DUT is unpopulated for check of 2.5 V at supply, populated for check of 2.5 V on board. Note that procedure only finalized at 200 krad(Si): beforehand, 2.5 V only measured either at supply or on board.

Before irradiation, the functionality of the DUT and bias board is confirmed by measuring the 2.5 V supply both at the supply and on the board, and by measuring the output voltage of the extra V_{CM} generator (see footnote 36 of Chapter 5). The measured values are given in Table B.2. Following functional confirmation, the DUT-and-bias-board combination is oriented perpendicular to the proton beam. A reference laser is used to center the DUT within the beam, and a stop guard used to ensure the distance from the DUT to the proton beam source is constant throughout all irradiations.[6] Following these alignments, the cave is sealed and the DUT exposed to the 50-MeV proton beam as per Table B.1.

[6]Specifically, the distance between the DUT and a 2-inch collimator (which restricts the beam to a 2-inch diameter) mounted after the beam source output is maintained at ∼10 inches.

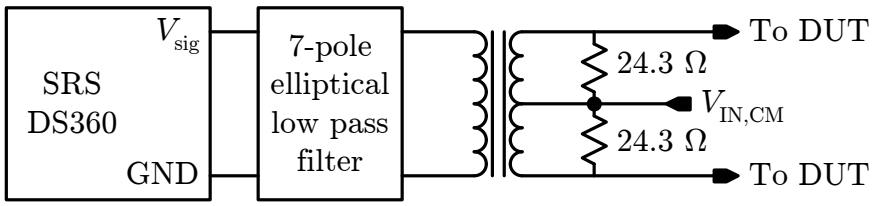

Figure B.2: Input signal path used in nonharmonic performance assessment. Low-pass filter implemented by Coilcraft P7LP-604L. Transformer implemented by Coilcraft AS8456-A.

B.3 Characterization

Ideally, the radiation characterization of the DUT should be the same as the baseline characterization described in Sections 6.1 and 6.2. However, in this experiment the radiation characterization is amended to reflect the limited time available for DUT characterization under the Military Standard [*MIL-STD-883G*, 2006, Method 1019.7]. In particular, only low-frequency 9.9909-kHz sinusoidal testing is conducted: high-frequency 366.007-kHz sinusoidal testing is omitted due to time constraints.

B.3.1 Characterization Setup

The custom test board of Section 6.1 is used for DUT characterization. Indeed, the entire experimental setup is the same as in Section 6.1 with the following amendment:

- For the nonharmonic performance assessment, the input signal path of Figure B.2 is used instead of that of Figure 6.2(b). (Note that for harmonic performance assessment, the input signal path of Figure 6.2(a) is maintained.) In baseline testing, switching between the harmonic and nonharmonic signal paths is done by depopulating and populating the RC filter. This switching consumes time and provides opportunity for operator error. The alternate input signal path of Figure B.2 circumvents this issue by employing a separate signal path, complete with its own Stanford Research Systems (SRS) DS360, for nonharmonic measurements.[7]

[7]Hence in all there are two SRS DS360s used during radiation characterization: one used for harmonic performance assessment and coupled with the input signal path of Figure 6.2(a), the other

B.3.2 Characterization Procedure

The DUT characterization schedule is given in Table B.3. Radiation characterization took place on 8 February 2008, with a final anneal dataset taken on 11 February 2008. Note that, especially for higher dose steps, there is a significant delay between the end of irradiation and onset of characterization. This delay is due to irradiation of Cave 4A itself during exposure: especially following longer exposures, radiation levels in the Cave were sufficiently high that operators needed to wait for the Cave to "cool" before entering and retrieving the DUT. Nonetheless, even with the delay, all characterizations nonetheless began within the maximum 1-hour onset delay prescribed by the Military Standard [*MIL-STD-883G*, 2006, Method 1019.7].

For radiation characterization, the following datasets are collected:

- *Functional check*

  To confirm that the DUT is still operational, three voltages are measured: the output voltage of the extra V_{CM} generator, the voltage at the current source pad V_{master} (see Figure 5.27), and the current through the external current reference resistor I_{master}. These measurements are summarized in Table B.4. In addition, measurements of the characterization board are performed to confirm functionality of the experimental setup itself;[8] these measurements are also included in Table B.4.

- *Power measurement*

  A 200-kHz sinusoid of varying amplitude[9] is input to the DUT through the harmonic input signal path of Figure 6.2(a). For each input signal amplitude, the current of each DUT supply (as reported by the HP6627A and HP6629A DC power supplies) is recorded.

used for nonharmonic performance assessment and coupled with the (separate) input signal path of Figure B.2.

[8] In particular, measurements of the DUT supplies, DUT reference voltages, input common mode voltage, and the gain along the nonharmonic signal path, are taken.

[9] Specifically, the input signal is swept through peak-to-peak amplitudes of 20 $\mu\mathrm{V_{PP}}$, 200 $\mu\mathrm{V_{PP}}$, 2 $\mathrm{mV_{PP}}$, 20 $\mathrm{mV_{PP}}$, 100 $\mathrm{mV_{PP}}$, 200 $\mathrm{mV_{PP}}$, 500 $\mathrm{mV_{PP}}$, 700 $\mathrm{mV_{PP}}$, 1.0 $\mathrm{V_{PP}}$, 1.2 $\mathrm{V_{PP}}$, 1.4 $\mathrm{V_{PP}}$, 1.6 $\mathrm{V_{PP}}$, 1.7 $\mathrm{V_{PP}}$, 1.8 $\mathrm{V_{PP}}$, 1.9 $\mathrm{V_{PP}}$, 2.0 $\mathrm{V_{PP}}$, 2.1 $\mathrm{V_{PP}}$, and 2.2 $\mathrm{V_{PP}}$. Note that 2.0 $\mathrm{V_{PP}}$ corresponds to a full scale peak-to-peak amplitude.

Dose step	Date	Irradiation stop time	Characterization		Delay	
[krad(Si)]			Start time	End time	Onset[a]	Completion[b]
10	08 Feb 2008	10:49	10:55	11:29	0:05	0:39
50	08 Feb 2008	11:46	11:40	12:16	0:03	0:39
100	08 Feb 2008	12:22	12:31	13:08	0:08	0:45
200	08 Feb 2008	13:19	13:32	14:06	0:12	0:46
500	08 Feb 2008	14:30	14:57	15:31	0:26	1:00
700	08 Feb 2008	15:58	16:24	16:58	0:25	0:59
1000	08 Feb 2008	17:23	17:48	18:24	0:24	1:00
1200	08 Feb 2008	18:47	19:17	19:52	0:29	1:04
2000	08 Feb 2008	20:57	21:58	22:58	1:00	2:00
Anneal	11 Feb 2008		13:45	14:21	64:47	65:23

[a] Onset delay is the time between the end of irradiation and the onset of characterization. Given in format of [hours]:[minutes].
[b] Completion delay is the time between the end of irradiation and the end of characterization. Given in format of [hours]:[minutes].

Table B.3: Characterization schedule for total-dose testing of the DUT by 50-MeV protons.

Dose step	V_{CM}	V_{master}	I_{master}	$V_{IN,CM}$	V_{REF+}	V_{REF-}	$V_{DD,A}$	$V_{DD,D}$	$V_{DD,CLK}$	$V_{DD,IO}$	$V_{DD,REF}$	Gain[a]
[krad(Si)]	[V]	[mV]	[μA]	[V]	[V]	[V]	[V]	[V]	[V]	[V]	[V]	
Prerad	1.251	924	64.9	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.830
10	1.251	923	64.8	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.830
50	1.251	923	64.8	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.830
100	1.251	923	64.8	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.830
200	1.251	923	64.8	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.831
500	1.251	922	64.8	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.830
700	1.251	923	64.8	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.832
1000	1.251	921	64.7	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.831
1200	1.252	920	64.6	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.828
2000	1.252	919	64.6	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.832
Anneal	1.251	921	64.7	1.250	1.750	0.750	2.505	2.502	2.507	2.500	2.506	0.829

[a] Gain defined as gain from SRS DS360 amplitude setting to DUT input.

Table B.4: Measurements taken during characterization to confirm the functionality of the DUT and the characterization test setup. All measurements made directly with an HP973A multimeter, except I_{master} (measured with an HP973A multimeter probing the voltage across a 10 kΩ test resistor) and gain (measured with an Agilent 54642A oscilloscope). All measurements made with DUT populated and active, except gain (measured with DUT unpopulated to reduce interference from sampling transients).

- *Clock sweep*

  As discussed in Section 6.1.3, a sweep is made of the clock signal offset and amplitude and a preliminary analysis of the hSFDR at each setting performed. From that analysis, three clock settings are selected for the more extensive harmonic and nonharmonic testing datasets described below.

- *Harmonic testing*

  A 9.99093-kHz sinusoid of varying amplitude[10] is input to the DUT through the harmonic input signal path of Figure 6.2(a). For each input signal amplitude, two records of 5,010,000 samples are acquired. In addition, one record of 1,050,000 samples is acquired for each calibration code. Harmonic testing is repeated at each of the three clock settings.

- *Nonharmonic testing*

  A 9.99093-kHz sinusoid of varying amplitude[11] is input to the DUT through the nonharmonic input signal path of Figure B.2. For each input signal amplitude, two records of 5,010,000 samples are acquired. In addition, one record of 1,050,000 samples is acquired for each calibration code. Nonharmonic testing is repeated at each of the three clock settings.

B.4 Results and Analysis

Although datasets at each of three clock settings are collected at each dose step, for the results presented in this dissertation, only datasets corresponding to the clock settings of Table B.5 are used.

The results of this experiment are presented in detail in Section 6.3.

[10] Specifically, the input signal amplitude is swept through the values given in footnote 9 of this appendix.

[11] Specifically, the input signal amplitude is swept through peak-to-peak amplitudes of 24 μV_{PP}, 240 μV_{PP}, 2.4 mV_{PP}, 24 mV_{PP}, 120 mV_{PP}, 240 mV_{PP}, 600 mV_{PP}, 840 mV_{PP}, 1.2 V_{PP}, 1.4 V_{PP}, 1.7 V_{PP}, 1.9 V_{PP}, 2.0 V_{PP}, 2.2 V_{PP}, 2.3 V_{PP}, 2.4 V_{PP}, 2.5 V_{PP}, and 2.7 V_{PP}. Note that these values are essentially the values used in the harmonic testing divided by 0.830, the nominal gain of the nonharmonic input signal path as reported in Table B.4.

Dose step	Clock offset	Clock amplitude
[krad(Si)]	[V]	[V_{PP}]
Prerad	1.3	1.8
10	1.3	1.8
50	1.3	1.7
100	1.3	1.9
200	1.3	1.9
500	1.3	1.9
700	1.3	1.9
1000	1.3	1.7
1200	1.3	1.9
2000	1.3	1.9
Anneal	1.3	1.7

Table B.5: Clock settings used in assessing DUT radiation performance.

Appendix C

SVADC-1 Single-Event Testing, 10-MeV Heavy Ion

This appendix describes the single-event testing[1] of the SVADC-1 converter by 10-MeV/nucleon heavy ions at Lawrence Berkeley National Laboratory (LBNL). This experiment seeks to ascertain the susceptibility of the SVADC-1 to latchup, and to determine the soft error rates of the SVADC-1. The experiment includes exposures to LETs as high as 58.72 MeV-cm^2/mg; this LET was the highest LET available in the 10-MeV/nucleon cocktail on the experiment dates.

The experiment was conducted on 26 and 27 September 2006 in Cave 4B of the 88-inch cyclotron facility at Lawrence Berkeley National Laboratory.[2] The experimenters were Charles C. Wang and Benjamin J. Mossawir with Stanford University, and Jeffery S. George, Rocky Koga, and Van T. Tran with The Aerospace Corporation. The experiment was coordinated by James L. Roeder of The Aerospace Corporation, principal investigator of the PARX project.

[1] An introduction to radiation testing, including single-event testing, is found in Appendix A.
[2] A floorplan of the 88-inch cyclotron, with Cave 4B labeled, is given in Figure A.3.

C.1 Device-Under-Test

The device-under-test (DUT) is the converter portion of the SVADC-1. A single DUT is tested. The DUT is de-lidded throughout the experiment. All testing is performed at room temperature and at standard supply voltages (i.e., with DUT supplies of 2.5 V).

C.2 Setup

This experiment requires a custom-designed setup particular to the constraints of Cave 4B. A general introduction to single-event testing at this venue is presented in Section A.3.4.[3]

A brief review of the physical layout of Cave 4B is helpful for understanding the setup of this experiment. The experiment takes place inside Cave 4B, a large room formed of thick concrete walls. Inside the cave, the heavy-ion beam is brought to a metal vacuum chamber that also houses the DUT: the DUT is exposed to the beam inside this chamber. During irradiation, the vacuum chamber is evacuated and the cave is sealed. As no experimenters are allowed inside the cave during irradiation, experimenters are instead stationed atop the cave roof. A permanent control trailer installed on the cave roof contains operator stations for controlling and monitoring the heavy-ion beam. Additional, experiment-specific operator stations can be established either inside the control trailer or elsewhere on the cave roof.

Naturally, the DUT must thus be remotely commanded and controlled. To signal down into the cave, a small hole in the cave ceiling near the control trailer allows cabling to pass between the cave roof and the cave interior. Typically, experiment-specific stations are constructed near this passage to reduce cable lengths. To signal into the vacuum chamber, specially-designed bulkheads mounted in the chamber wall are used to pass signals in and out of the chamber without compromising the vacuum.

In light of these constraints, the experimental setup constructed for this experiment is spread over all three physical regions, that is, throughout the vacuum

[3]Further information about radiation testing in Cave 4B can be found in *McMahan* [1999].

chamber, cave, and cave roof. The complete setup is shown in Figures C.1 and C.2. While the following subsections describe this setup in more detail, three things are worth mentioning beforehand. First, the vacuum chamber bulkheads limit both the type and number of signals that can pass to and from the chamber interior. In particular, the vacuum chamber allows for a maximum of four bulkheads, with two available on opposite sides of the chamber. For this experiment, two bulkheads are used: a BNC bulkhead and a 40-pin ribbon cable bulkhead. The former practically accommodates about 10 BNC signals[4] and the latter accommodates up to four 40-pin ribbon cable connections.[5] Second, ideally all signals input to and output from the vacuum chamber are cabled up to the cave roof for operator control. However, such a scheme proved unviable for this experiment due to signal degradation and noise pickup over such long cable traverses (>30 feet all told from the chamber bulkhead to the cave roof). Hence, equipment was placed both on the cave roof and inside the cave itself: equipment at the former can be directly controlled but requires long cable traverses, while equipment at the latter allows shorter cable traverses but requires remote control. Third, to prevent charge buildup during irradiation, this experiment eschews floating nodes in the DUT.

C.2.1 Test Board

The DUT is populated in the clam-shell socket on the custom test board described in Section 6.1. The test board is then attached to a mounting bracket inside the vacuum chamber. The bracket is motorized and coupled with a camera system: this system enables operators to align the DUT with the heavy-ion beam from the control trailer after the vacuum chamber has been sealed.[6] The heavy-ion beam itself is ∼3 inches in diameter, narrow enough that only the DUT and some passive elements near the

[4] In reality, this bulkhead provides additional BNC connectors, but the bulkhead geometry makes these connectors very difficult to connect to.

[5] Although the pins on some of these connectors are bent, discouraging their use.

[6] Notably, while the metal walls of the chamber provide a degree of shielding from external electromagnetic interference, since the motors lie inside the chamber, their emitted low frequency noise can readily couple into the DUT signal lines.

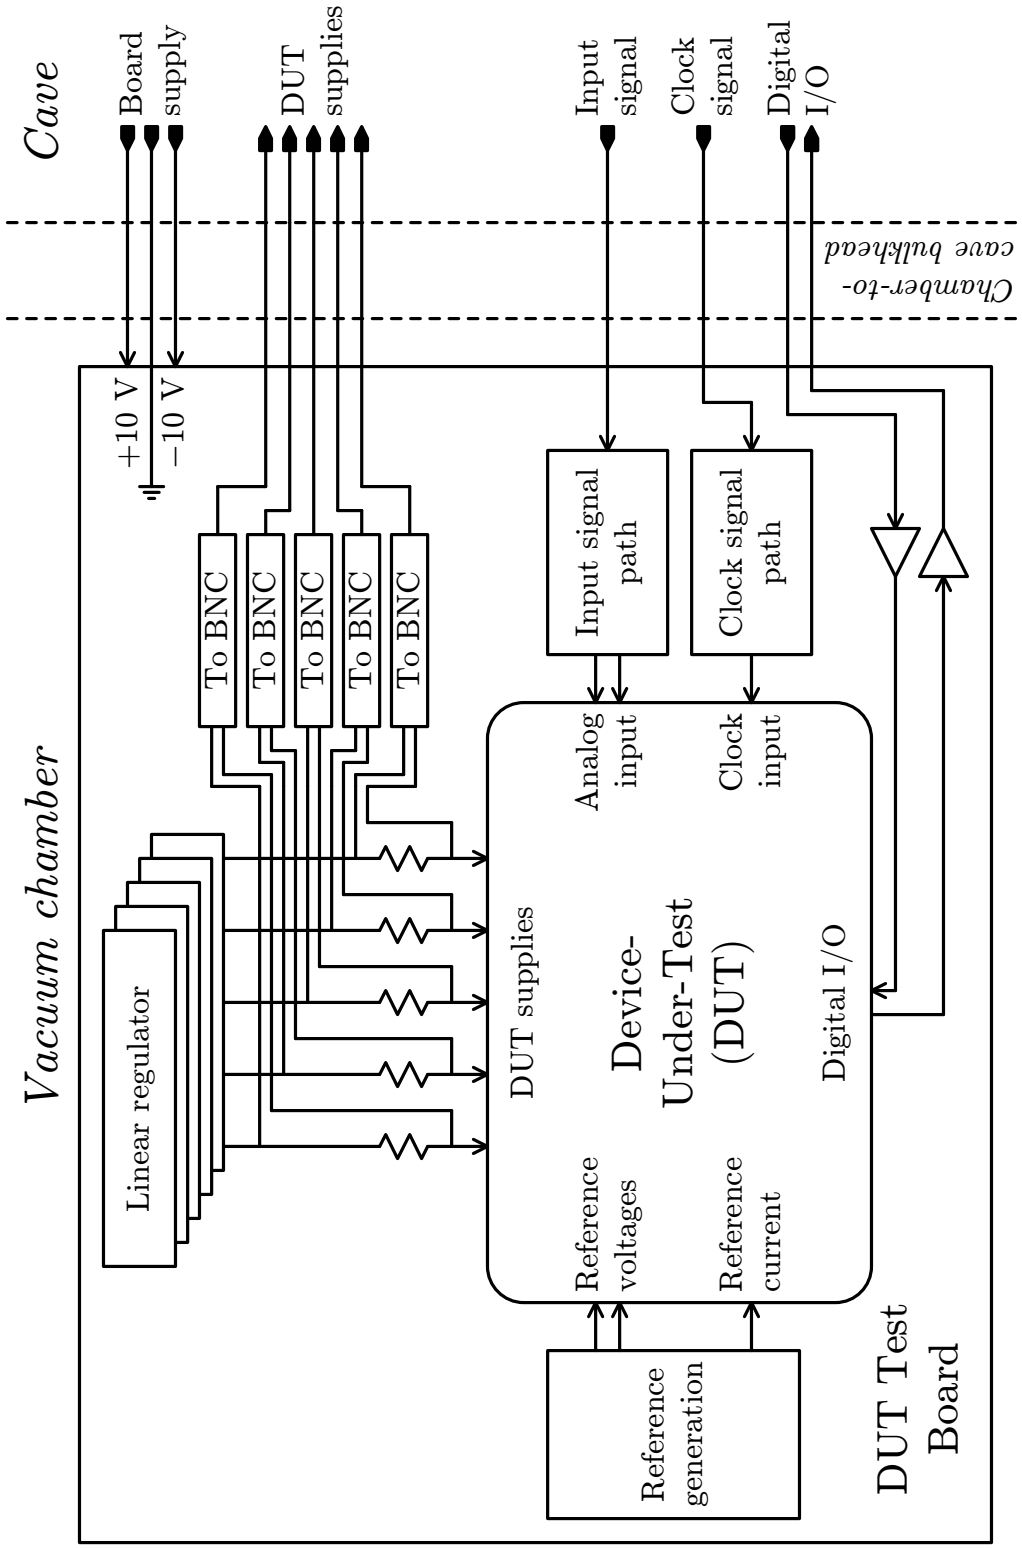

Figure C.1: Setup for 10-MeV/nucleon heavy-ion single-event testing of the DUT, showing vacuum chamber and chamber-to-cave bulkhead portions. See Figure C.2 for cave and cave roof portions. For input signal path and clock signal path, see Figures C.3 and C.4, respectively.

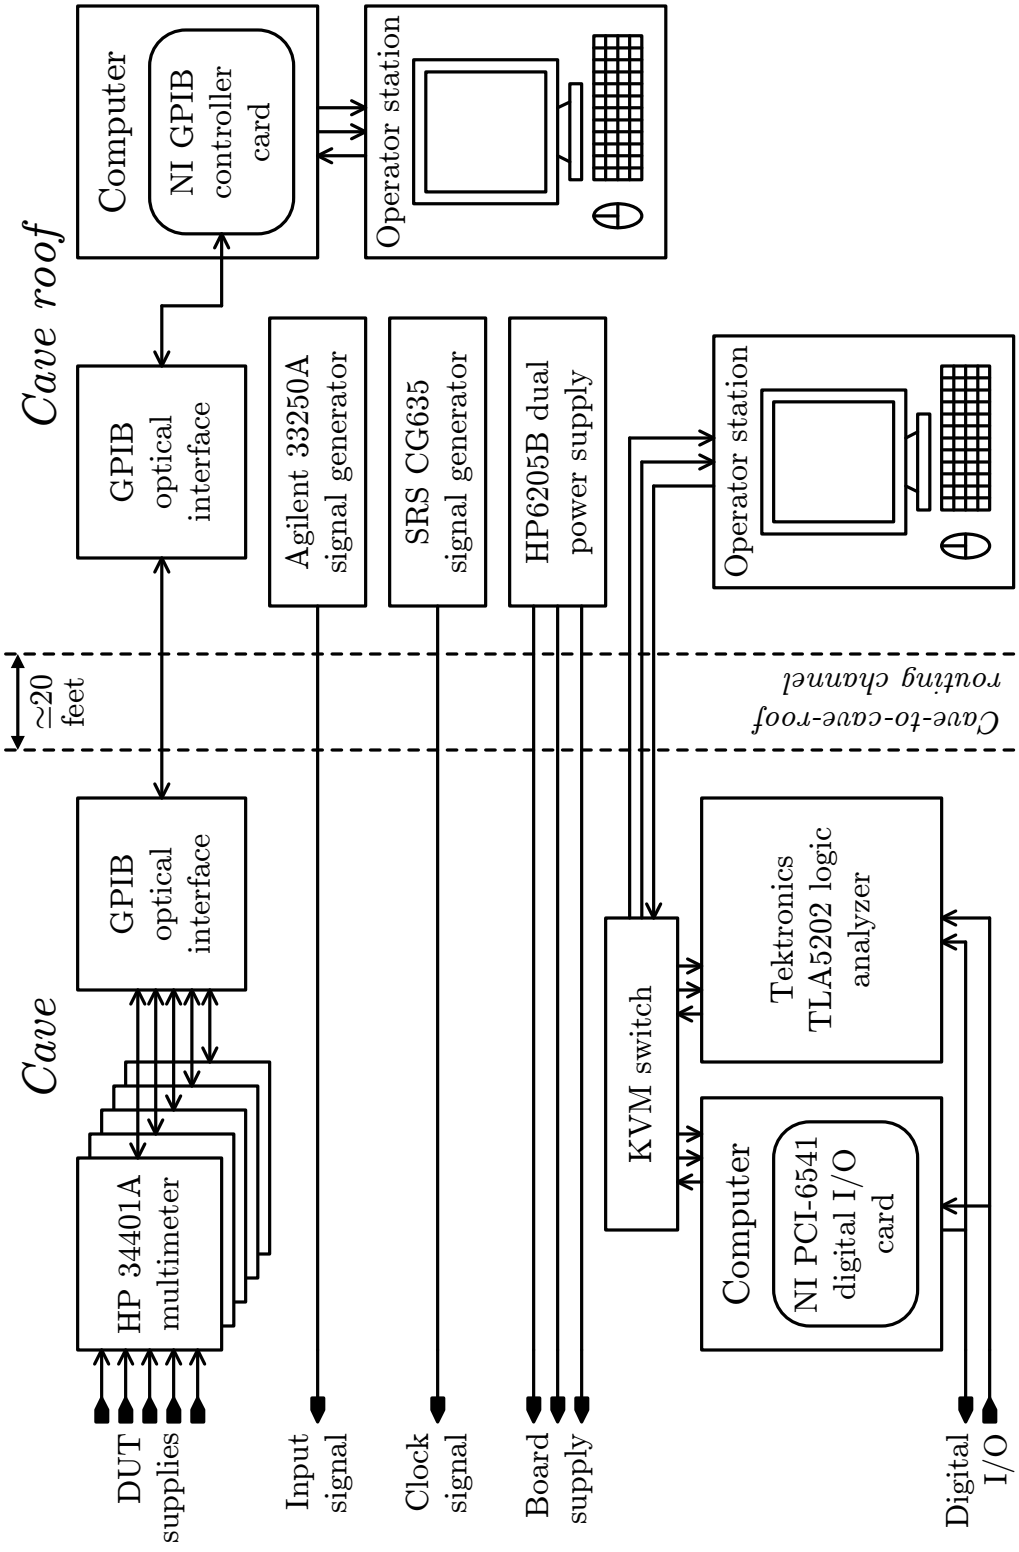

Figure C.2: Setup for 10-MeV/nucleon heavy-ion single-event testing of the DUT, showing cave and cave roof portions. See Figure C.1 for vacuum chamber and chamber-to-cave bulkhead portions.

DUT are illuminated during exposure.[7]

C.2.2 Board Supply

All power to the test board—including all power to the DUT (see Section C.2.3)—is provided by ±10-V supplies generated by an HP6205B dual power supply located on the cave roof. This design brings two primary benefits. First, it reduces the number of chamber bulkhead BNC connectors used for power transmission. Second, if latchup should occur anywhere on the test board—including the DUT—the power can be rapidly removed before component burn-out by turning off this single supply.

To reduce the effect of noise pick-up on the long cables, the ±10-V supplies are linearly regulated on the test board before use.

C.2.3 DUT Supplies

For the DUT supplies, the +10-V supply is linearly regulated to 2.5 V by five Linear Technology LT1763 linear regulators. Each regulator subsequently powers a single DUT supply. Note that each LT1763 regulator can supply up to 500 mA [*Linear Technology Corporation*, LT1763], well in excess of the current consumed by any DUT supply, and hence can provide the increased current required to sustain a latchup event.

To provide DUT supply monitoring, a small series resistor is imposed between each regulator output and DUT supply.[8] The voltage across each resistor is connected differentially to a BNC cable and then output through the chamber bulkhead for continuous monitoring by a devoted HP34401A multimeter. Since tests showed that long cable lengths between the series resistor and corresponding multimeter garnered increased noise visible in the DUT output, the multimeters are placed inside the

[7] The remaining test board electronics are commercial parts and not radiation-qualified but, being outside the beam, should not upset.

[8] The resistor values are scaled so that the voltage drop across the resistors is no more than 50 mV for any supply. In particular, the resistors have measured values of 3.02 Ω for the $V_{DD,A}$ supply, 8.41 Ω for the $V_{DD,D}$ supply, 20.17 Ω for the $V_{DD,CLK}$ supply, 30.06 Ω for the $V_{DD,IO}$ supply, and 100.02 Ω for the $V_{DD,IO}$ supply.

cave, close to the chamber.[9] A GPIB interface enables a computer on the cave roof to remotely query each multimeter in turn. To overcome GPIB cable length restrictions,[10] the GPIB signal is converted to an optical signal for the long traverse between the multimeters in the cave and the computer on the cave roof. The cave roof computer is equipped with a National Instruments (NI) GPIB controller card and runs a custom MatLAB script developed by Benjamin J. Mossawir that enables real-time display of the voltages measured by all five multimeters.

C.2.4 Input Signal Generation

To maintain input signal flexibility while guaranteeing sufficient drive strength of the DUT input, the input signal is generated by an Agilent 33250A signal generator located on the cave roof and buffered on the test board by an AD8138 fully differential operational amplifier circuit. The complete circuit is shown in Figure C.3. In addition to providing sufficient drive strength, the buffer also converts the single-ended output of the Agilent 33250A to a fully differential signal.

Due to its length, the input signal cable picks up noise that is subsequently transmitted to the DUT input by the buffer circuit. In this experiment, the noise level at the DUT input was found to be ∼450 μV (standard deviation) as measured at the DUT output, roughly 3 times larger than that measured in a quiescent laboratory environment.[11]

C.2.5 Clock Signal Generation

Similar to the input signal, the clock signal is generated by an Stanford Research Systems (SRS) CG635 clock generator located on the cave roof and driven to the test

[9] This configuration essentially places a voltmeter in parallel with each series resistor. An alternative is to remove the series resistor entirely and instead impose an ammeter in series between each regulator output and DUT supply. However, tests showed that the ammeter configuration garnered significantly greater noise, and hence the voltmeter configuration was chosen instead.

[10] According to IEEE Standard 488.1 (which governs GPIB), "caution should be taken if [an] individual cable length exceeds 4 m" [*IEEE Std 488.1-2003*, 2003, Sect. 8.4.2].

[11] Given the long cable lengths and the general noise environment at the 88-inch cyclotron facility, experimenters typically expect noise levels to be on the order of about a milliVolt (standard deviation).

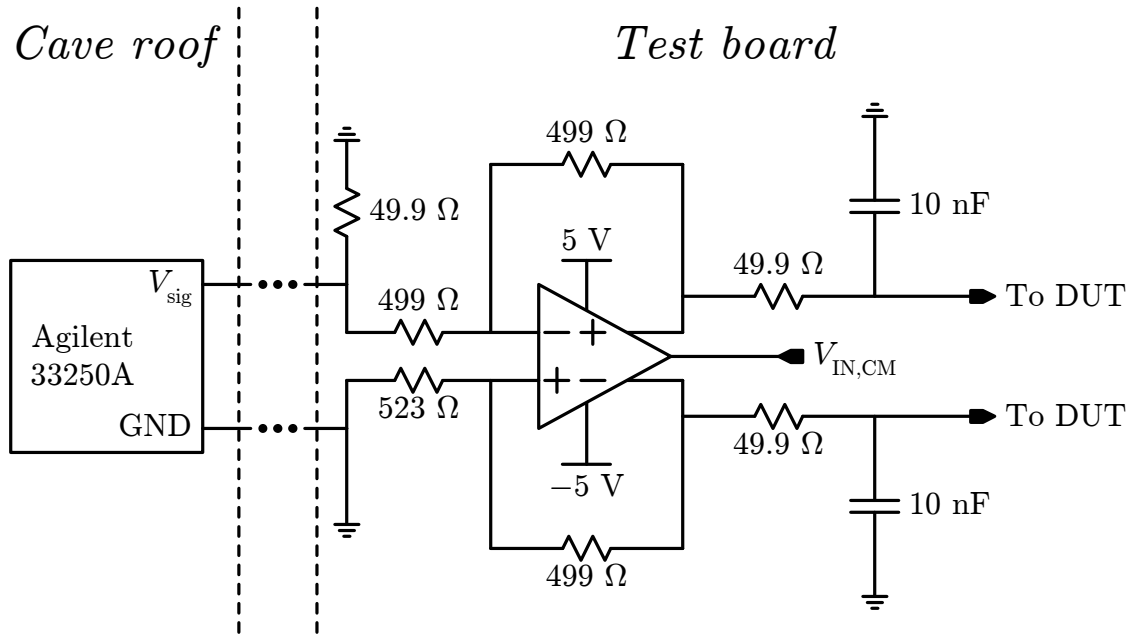

Figure C.3: Input signal path. Operational amplifier implemented by Analog Devices AD8138, capacitors by polystyrene capacitors. For $V_{\mathrm{IN,CM}}$ generation, see Figure 6.4.

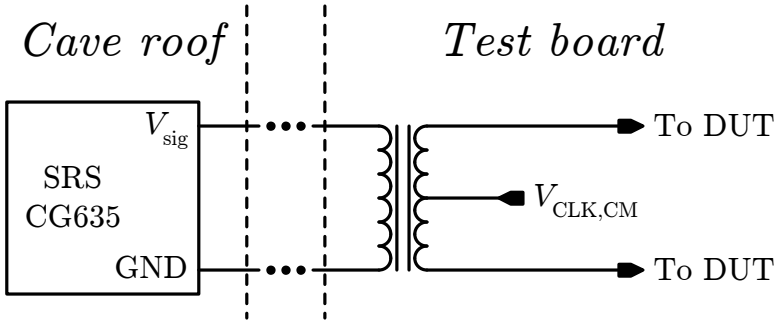

Figure C.4: Clock signal path implemented on test board. Transformer implemented by Coilcraft TTWB2010-IL. A separate instance of the circuit of Figure 6.4 (i.e., an instance in addition to that which generates $V_{\mathrm{IN,CM}}$) generates $V_{\mathrm{CLK,CM}}$.

board. Since the clock signal is buffered internally within the SVADC-1 (as depicted in Figure 5.13), it is not actively buffered on the test board.

Notably, the SVADC-1 supports fully differential clock injection via two complementary pins. One clock pin is processed as described in Section 5.3.4 and shown in Figure 5.13, while the other clock pin is buffered by inverters (to provide equal loading) but subsequently ignored. To prevent floating nodes, fully differential clocking is used in this experiment, with the single-ended output of the SRS CG635 converted to differential via a transformer circuit as shown in Figure C.4.[12]

C.2.6 Reference Generation

For reference voltage generation, both external and internal generation is used. Specifically, as per external reference voltage generation, 0.75 V and 1.75 V references are generated on board by a resistor ladder, buffered by the MAX4252 operational amplifier circuit of Figure 6.5, and driven into the $V_{\text{REF}-}$ and $V_{\text{REF}+}$ pins of the SVADC-1. In addition, as per internal reference voltage generation, the $V_{\text{REF}+}$ pin is shorted to the $V_{\text{SENSE}+}$ pin, and the $V_{\text{REF}-}$ pin is shorted to the $V_{\text{SENSE}-}$ pin, by on-board jumpers (see Figure 5.26). Internal generation prevents the $V_{\text{SENSE}+}$ and $V_{\text{SENSE}-}$ DUT nodes from floating during irradiation, while external generation helps ensure that the voltage references are well-known and thus need not be separately monitored.

For reference current generation, the resistor internal to the SVADC-1 is used: see Figure 5.27. Note that, to prevent floating nodes, the current source pad of Figure 5.27, though unused, is nonetheless loaded with a 16 kΩ resistor to ground.

[12] An alternate solution would be to use single-ended clocking and simply tie the unused clock pin to ground. However, fully differential clocking allows the ESD and buffer circuitry of both clock pins to be fully active during irradiation. As latchup can be triggered by surges encountered during normal switching [*Troutman*, 1986, p. 25], an active, switching pin should be more latchup susceptible than a static, grounded pin, enabling a more robust assessment of the latchup sensitivity of the SVADC-1.

C.2.7 Digital I/O

The digital I/O is buffered on the test board as described in Section 6.1.5. The digital signals are transmitted through two 40-pin ribbon cables that connect the test board to the chamber bulkhead. The cables are configured such that board-driven ground lines are interlaced between signal lines, reducing crosstalk.

As it is difficult to drive the digital signals over the >30-foot traverse from the chamber to the cave roof, the digital I/O is instead managed by a computer located inside the cave. This choice reduces the 40-pin ribbon cables to more-manageable lengths of ∼4 feet from the test board to the computer; to further reduce electromagnetic interference, the ribbon cables are wrapped in grounded aluminum foil. The computer is equipped with an NI PCI-6541 card and, via the program described in Section 6.1.5, both controls the calibration inputs and collects the SVADC-1 outputs. In addition to this computer, a Tektronics TLA5202 logic analyzer is placed in the cave and configured to monitor the bus, providing rapid visualization useful for real-time analysis and debugging. To enable remote operation of both pieces of equipment, the computer and logic analyzer are connected to a keyboard-video-mouse (KVM) switch that routes their human interfaces to the cave roof.

C.3 Procedure

For this experiment, the DUT is exposed to 10-MeV/nucleon heavy-ion beams of different LETs. Each exposure is called a run.

Two DUTs are tested during this experiment. The run summary for each is given in Tables C.1 and C.2, respectively. Notably, after characterization of the first DUT, it was discovered that this DUT displayed a gain error,[13] and hence a second DUT was characterized. While the results of the first DUT are thus not viable for cross section characterization, they are nonetheless valid for latchup susceptibility characterization and hence are included here. However, as it is the results of the second DUT that

[13]This gain error was subsequently confirmed to exist prior to single-event testing: it was not caused by irradiation.

Ion	LET	Input signal	Exposure	Average flux	Maximum flux	Fluence
	[MeV-cm^2/mg]		[sec]	[ions/cm^2-sec]	[ions/cm^2-sec]	[ions/cm^2]
[none]	0	0 V_{DC}	240	0	0	0
Ar	9.7	0 V_{DC}	60.04	2.54×10^2	1.16×10^3	1.53×10^4
		0 V_{DC}	60.02	5.23×10^3	8.40×10^3	3.14×10^5
		0 V_{DC}	60.04	6.47×10^4	1.44×10^5	3.89×10^6
Kr	31.28	0 V_{DC}	60.03	2.91×10^4	6.52×10^4	1.74×10^6
		0 V_{DC}	60.04	2.84×10^4	6.45×10^4	1.70×10^6
		0 V_{DC}	60.03	2.87×10^4	6.55×10^4	1.73×10^6
Ne	3.45	0 V_{DC}	60.04	2.66×10^4	5.91×10^4	1.60×10^6
		0 V_{DC}	60.03	2.60×10^4	5.83×10^4	1.56×10^6
		0 V_{DC}	60.02	2.56×10^4	4.21×10^4	1.54×10^6
B	0.87	0 V_{DC}	60.04	4.38×10^4	9.97×10^4	2.63×10^6
		0 V_{DC}	60.03	4.50×10^4	1.22×10^5	2.70×10^6
		0 V_{DC}	60.03	4.44×10^4	9.68×10^4	2.67×10^6
		0 V_{DC}	60.04	4.70×10^4	1.14×10^5	2.82×10^6
		0 V_{DC}	60.04	6.21×10^4	1.18×10^5	3.73×10^6
O	2.22	0 V_{DC}	60.02	3.43×10^4	8.13×10^4	2.06×10^6
		0 V_{DC}	60.03	3.34×10^4	8.29×10^4	2.01×10^6
		0 V_{DC}	60.03	3.28×10^4	8.11×10^4	1.97×10^6
		0 V_{DC}	60.03	3.14×10^4	3.30×10^5	1.89×10^6
		0 V_{DC}	60.03	3.14×10^4	7.89×10^4	1.88×10^6
Cu	21.33	0 V_{DC}	60.02	4.12×10^4	2.69×10^5	2.47×10^6
		0 V_{DC}	120.03	4.06×10^4	2.70×10^5	4.87×10^6
Ag	48.16	0 V_{DC}	60.02	7.54×10^3	1.42×10^4	4.52×10^6
		0 V_{DC}	300.04	7.29×10^3	1.46×10^4	2.19×10^6
Xe	58.72	0 V_{DC}	60.03	1.98×10^4	4.30×10^4	1.19×10^6
		0 V_{DC}	300.02	2.01×10^4	7.02×10^4	6.03×10^6

Table C.1: Run summary for single-event testing of the first DUT by 10-MeV/nucleon heavy ions. Runs conducted from 16:07 to 19:49 on 26 September 2006 and listed in order of execution. DC input signals given in units of V_{DC}.

Ion	LET	Input signal	Exposure	Average flux	Maximum flux	Fluence
	[MeV-cm^2/mg]		[sec]	[ions/cm^2-sec]	[ions/cm^2-sec]	[ions/cm^2]
[none]	0	1 mV_{DC}	300	0	0	0
		0.5 V_{DC}	300	0	0	0
		2.3 V_{PP}	300	0	0	0
Ne	3.45	[n/a][a]	60.03	4.03×10^4	7.56×10^4	2.42×10^6
		1 mV_{DC}	60.04	4.09×10^4	8.37×10^4	2.45×10^6
		1 mV_{DC}	240.04	4.00×10^4	8.38×10^4	9.61×10^6
		[n/a][a]	300.02	3.99×10^4	8.07×10^4	1.20×10^7
		0.5 V_{DC}	300.04	4.08×10^4	8.12×10^4	1.22×10^7
Ar	9.74	1 mV_{DC}	60.02	2.44×10^4	4.52×10^4	1.46×10^6
		1 mV_{DC}	240.04	2.70×10^4	5.01×10^4	6.48×10^6
		0.5 V_{DC}	300.02	2.83×10^4	5.32×10^4	8.48×10^6
Cu	21.33	1 mV_{DC}	60.03	4.96×10^4	2.74×10^5	2.98×10^6
		1 mV_{DC}	240.03	5.01×10^4	2.80×10^5	1.20×10^7
		0.5 V_{DC}	300.03	5.06×10^4	4.70×10^5	1.52×10^7
Kr	31.28	1 mV_{DC}	60.04	2.30×10^4	3.80×10^4	1.38×10^6
		1 mV_{DC}	240.03	2.29×10^4	5.63×10^4	5.49×10^6
		0.5 V_{DC}	300.03	2.26×10^4	5.46×10^4	6.79×10^6
Xe	58.72	[n/a][a]	3.81	3.12×10^4	7.62×10^4	1.19×10^5
		1 mV_{DC}	60.02	3.09×10^4	7.70×10^4	1.85×10^6
		1 mV_{DC}	240.02	3.15×10^4	7.90×10^4	7.56×10^6
		0.5 V_{DC}	300.04	3.13×10^4	7.98×10^4	9.38×10^6
		2.3 V_{PP}	300.03	3.18×10^4	8.13×10^4	9.55×10^6

[a] [n/a] indicates the Agilent 33250A output was not active during this run. In this case, the input buffer circuit drives an input of ~0 V_{DC} into the DUT.

Table C.2: Run summary for single-event testing of the second DUT by 10-MeV/nucleon heavy ions. Runs conducted from 22:07 on 26 September 2006 to 00:40 on 27 September 2006 and listed in order of execution. DC input signals given in units of V_{DC}. AC input signals given in units of V_{PP} (i.e., peak-to-peak amplitude).

are cited throughout this dissertation, the remainder of this discussion focuses on the testing of this part.[14]

For each LET, the DUT is first exposed to the beam for a short duration to check for latchup. If the DUT does not latchup, longer exposures for cross section measurement follow. At each LET, the DUT cross section is measured with at least two DC inputs: a 1 mV input (wherein the internal differential circuitry of the DUT is largely balanced), and a 0.5 V input (wherein the internal differential circuitry is decidedly unbalanced). In addition, an overranging, 9.99093-kHz sine wave input is included at the highest LET to check for any input-dependent latchup susceptibility.

For each run, exposure is begun by an authorized operator in the control trailer atop the cave roof. At the same time, operators at the data acquisition computer (housing the NI PCI-6541 card) and the DUT supply monitoring computer (housing the NI GPIB card) begin data acquisition. This triad of operations (beam exposure, data acquisition, and DUT supply monitoring) is synchronized by verbal coordination amongst three operators at the three separate computer stations.[15] As the run proceeds, the DUT supply monitoring computer and the HP6205B power supply are closely monitored for latchup events. Furthermore, on both the data acquisition and DUT supply monitoring computers, the appropriate data is acquired in block periodic form, that is, in the form of contiguous blocks of data spaced at regular (i.e., periodic) intervals. After each run, preliminary histograms of the first few data files are compiled to confirm data integrity. Based on these results, the next run is determined and executed.

C.3.1 Data Collection

During each run, data is acquired on both the data acquisition computer and on the DUT supply monitoring computer. In particular:

- *Data acquisition computer*

  The data acquisition computer is programmed to gather a preset number of

[14] Although the testing of the first DUT was similar.

[15] As a result, there are small start and stop time inconsistencies between the three operations. However, the length of the runs renders these inconsistencies negligible.

samples, write these samples to disk, and then repeat this gather-and-write process over a preset duration. Incoming samples are ignored during disk write. The result is periodic blocks of data. Insofar as heavy-ion-induced upset of the DUT is a stationary process, such block periodic sampling introduces no caveats to data validity. For this experiment, data files consist of 8 million samples acquired at a 5 MS/s rate, which requires about 3.2 seconds to gather and write to disk (for a duty cycle of ∼50%).

- *DUT supply monitoring computer*
  Upon request from the DUT supply monitoring computer, each monitoring multimeter acquires 20 samples and places the results on the GPIB bus. The computer then cycles between multimeters, gathering 20 samples from each. As GPIB bus transfer time can vary not only between multimeters, but even between acquisitions from the same multimeter, precise time-stamping of the data is not possible. This imprecision precludes sample-by-sample alignment of the DUT supply data with the DUT output data. However, since the multimeter data remains fairly invariant throughout each run—if not throughout the entire experiment—this lack of sample-by-sample alignment is not of concern. The multimeter data is both immediately displayed and saved for future analysis.

C.4 Results and Analysis

For both DUTs, no latchup is seen in any run. For the first DUT, at the maximum LET of 58.72 MeV-cm^2/mg, these runs include a total beam fluence of 7.22×10^6 ions/cm^2 with a 0 V DC input. For the second DUT, at the maximum LET of 58.72 MeV-cm^2/mg, these runs include a total beam fluence of 2.85×10^7 ions/cm^2 with a mixture of both DC (composed of both 1 mV and 0.5 V DC signals) and AC (composed of a 2.3 V_{PP} sinusoid at 9.99093 kHz) inputs.

For cross section analysis, a Turflinger analysis is conducted on the second DUT.[16] For this analysis, the DUT output is not self-calibrated. However, all 4 bits of the

[16] In light of the aforementioned gain error, no cross section analysis is performed for the first DUT.

Input signal	LET	Signal bins Gaussian			Noise bins Gaussian		
		μ_B	σ_B	α_B	μ_E	σ_E	α_E
	[MeV-cm^2/mg]	[code]	[code]		[code]	[code]	
1 mV_{DC}	3.45	7900.9	3.57	4	7901.0	17.72	2
	9.74				7901.3	28.45	2
	21.33				7902.4	32.31	2
	31.28				7900.8	32.26	2
	58.72				7901.5	38.27	2
0.5 V_{DC}	3.45	4188.5	3.60	4	4187.2	18.85	2
	9.74				4191.0	32.51	2
	21.33				4189.4	37.93	2
	31.28				4188.9	30.88	2
	58.72				4189.7	37.67	2

Table C.3: Properties of the Gaussians used in Turflinger analysis of second DUT. Note that analysis is done with 14-bit digital reconstruction of DUT output ($2^{14} = 16384$ total codes). For resulting soft error rates, see Section 6.3.2.3, especially Figure 6.21.

terminating ADC (i.e., of stage 6) are used, resulting in a nominal 14-bit digital reconstruction. For the results of this analysis, see Section 6.3.2.3. For completeness, the signal bin and noise bin Gaussian distributions used in the analysis are given in Table C.3.

Appendix D

SVADC-1 Single-Event Testing, 25-MeV Heavy Ion

This appendix describes the single-event testing[1] of the SVADC-1 converter by 25-MeV/nucleon heavy ions at Texas A&M University (TAMU). This experiment seeks to ascertain the susceptibility of the SVADC-1 to latchup under conditions of elevated supply voltages (i.e., 2.7 V) and temperatures as high as 131°C at an LET of 62 MeV-cm^2/mg.[2]

The experiment was conducted on 20 June 2008 at the Radiation Effects Facility at the Cyclotron Institute on the campus of Texas A&M University. The experimenters were Kirby Kruckmeyer and Tom Santiago, both with National Semiconductor Corporation. All testing was done in compliance with JESD-57, the Electronic Industries Association (EIA) standard for measurement of single-event effects from heavy-ion irradiation [*EIA/JESD57*, 1996].

D.1 Device-Under-Test

The device-under-test (DUT) is the converter portion of the SVADC-1. A single DUT is tested. The DUT is de-lidded throughout the experiment. To increase the

[1] An introduction to radiation testing, including single-event testing, is found in Appendix A.
[2] A more complete explanation of the motivation behind this testing is given in Section 6.3.2.1.

temperature of the DUT during testing, resistive spiral heaters are attached to the backside of the package. The heaters are powered by a Lakeshore 332 temperature controller. The DUT chip surface temperature is monitored by a Fluke 80T-IR infrared temperature probe, with care taken to aim the probe at a region of the chip with no metal coverage.

D.2 Setup

This experiment adopts the setup of the 10-MeV/nucleon heavy-ion testing at Lawrence Berkeley National Laboratory described in Section C.2, albeit amended to reflect both the focus on latchup, and the different venue. The amendments are listed below. Note that, similar to the 88-inch cyclotron at Lawrence Berkeley National Laboratory, the Cyclotron Institute on the campus of Texas A&M University enforces separation between the DUT and the operators during irradiation. Hence the test board is connected to the test equipment by long cables >30 feet in length.

- The test board is angled at a 48° angle to increase the effective linear energy transfer (LET) of the beam. Steeper angles would have caused shadowing of the clam shell socket over the die.
- All power supplies, including both the board and DUT supplies, are provided by a suite of benchtop supplies composed of HP6218As, HP6205Bs, and Instek PC-3030Ds. Each benchtop supply is routed through a Fluke 8050A multimeter that monitors the supply current before being cabled to the appropriate DUT supply. For the DUT supplies, this configuration does incur long cabling between each power supply and its corresponding DUT supply, however, as this testing focuses on latchup and not soft error rate measurements, the subsequent noise is permissible. The board supplies are set to ±10 V, while all DUT supplies are set to 2.7 V.
- Instead of differential clocking, single-ended clocking as per Section 6.1.3 is used. The clock signal is set to a 5-MHz square wave with 1.25 V offset and 2.5 V_{PP} amplitude.

- In general, the DUT digital outputs are not collected. However, a few outputs are monitored by a Tektronics 7404B oscilloscope to confirm that the DUT output is active both before and after irradiation.[3]

- The input signal is a 10-kHz, 2-V_{PP} sinusoid. It is still buffered on board as shown in Figure C.3.

- In many cases, different, though functionally equivalent, test equipment is used. Specifically: aside from those equipment changes already mentioned, 1) an Agilent 33120A provides the input signal (instead of an Agilent 33250A), and 2) an HP8111A provides the clock signal (instead of an Stanford Research Systems CG635).

D.3 Procedure

For this experiment, the DUT is exposed to a 25-MeV/nucleon heavy-ion beam of fixed LET. Each exposure is called a run. In contrast to the 10-MeV/nucleon heavy-ion testing of Appendix C, in this experiment it is the DUT temperature, and not the beam LET, that is varied between runs.

The heavy-ion beam is fixed to a Xenon beam: given the 48° incident angle of the DUT, this beam has an effective LET of 62 MeV-cm^2/mg. Note that this beam was the highest LET beam available in the 25-MeV/nucleon cocktail on the experiment dates.

Runs are conducted at several temperatures between room temperature and 131°C, as summarized in Table D.1. The currents of both the test board supplies and the DUT supplies are visually monitored during each run, and the current at the start and end of each run recorded. Note that the experiment includes two runs—the initial run at room temperature, and the final run at 131°C—where the temperature is held constant throughout the entire run. Whereas during the intermediate runs, the

[3]Note that the digital outputs of the SVADC-1 are nonetheless still loaded since they drive on-board digital buffers (rather, it is the on-board digital buffers that are unloaded).

DUT is being heated during the run; these runs are thus described by the temperature at either the start or end of the run.

D.4 Results and Analysis

The results of the experiment are also given in Table D.1. Note that if latchup had occurred, the DUT supply current at the end of the run would be significantly greater than that at the start of the run. While some of the DUT supplies do display current increases from the start to end of a run, these increases are small and certainly not indicative of latchup.[4] Thus the SVADC-1 displays no latchup throughout all the runs, including a worst-case run at an effective LET of 62 MeV-cm^2/mg (highest tested LET) under conditions of elevated supply (2.7 V) and temperature (131°C).

[4]There is one peculiar result in the significant drop in $V_{DD,IO}$ current during run 2. It is not clear whether this drop is due to an effect within the DUT or within the test setup: at this venue experiments are subject to large amounts of noise (aggravated by the long cable lengths between the equipment and the test boards) and possible instabilities. If the drop is indeed attributable to the DUT itself, though, then it amounts to a soft error—in particular, a single-event functional interrupt (SEFI)—and not latchup.

Run	Temperature	Beam			DUT supplies (set to 2.7 V)					Board supplies	
		Average flux	Fluence		$V_{DD,A}$	$V_{DD,CLK}$	$V_{DD,D}$	$V_{DD,IO}$	$V_{DD,REF}$	+10 V	−10 V
	[°C]	[ions/cm^2-sec]	[ions/cm^2]		[mA]	[mA]	[mA]	[mA]	[μA]	[mA]	[mA]
1	Room	4.14×10^4	1.31×10^7	Start	20.6	9.3	11.4	5.0	9.4	127.2	40.0
				End	20.6	9.3	11.4	5.0	9.4	127.3	40.0
2	65 (start)	4.61×10^4	8.41×10^6	Start	21.4	9.5	11.5	5.5	9.5	127.4	40.1
				End	21.5	9.5	11.5	1.7	9.5	127.4	40.1
3	93 (end)	3.36×10^4	1.03×10^7	Start	21.5	9.5	11.6	7.2	9.6	127.4	40.1
				End	21.5	9.5	11.6	7.5	9.6	127.9	40.1
4	121 (end)	3.42×10^4	1.25×10^7	Start	21.9	9.5	11.6	11.8	9.5	127.6	40.1
				End	22.0	9.5	11.6	12.7	9.6	127.7	40.1
5	131	4.62×10^4	1.09×10^7	Start	29.0	12.9	15.7	14.1	12.8	127.6	40.1
				End	29.0	12.9	15.7	14.1	12.8	127.6	40.1

Table D.1: Run summary for single-event testing of the DUT by 25-MeV/nucleon heavy ions. Measured currents given at both start and end of run. If measured temperature is constant during run, single temperature given; otherwise, temperature at either start or end of run given. Xenon beam of an effective LET of 62 MeV-cm^2/mg used throughout all runs.

Appendix E

SVADC-1 Single-Event Testing, Pulsed Laser

This appendix describes the single-event testing[1] of the SVADC-1 converter by pulsed laser at The Aerospace Corporation's facility in El Segundo, California. The objectives of this experiment are first, to probe the latchup susceptibility of the SVADC-1 as a prelude to more extensive heavy-ion testing, and second, to explore any interesting soft error upset signatures.

A first session of laser testing conducted on 30 June 2006 focused on the first objective, while a second session conducted on 26 and 27 October 2006 focused on the second. While this appendix addresses both sessions, it concentrates primarily on the second session, in which the SVADC-1 was probed by a 590-nm laser with beam energies as high as 4.26 nJ/pulse. The experimenters for both sessions were Charles C. Wang with Stanford University and Stephen L. LaLumondiere with The Aerospace Corporation. Both sessions of testing were coordinated by James L. Roeder of The Aerospace Corporation, principal investigator of the PARX project.

[1] An introduction to radiation testing, including single-event testing, is found in Appendix A.

E.1 Device-Under-Test

The device-under-test (DUT) is the converter portion of the SVADC-1. A single DUT is tested. The DUT is de-lidded throughout the experiment to allow direct exposure of the silicon. All testing is performed at room temperature and at standard supply voltages (i.e., with DUT supplies of 2.5 V).

E.2 Setup

This experiment adopts the setup of the 10-MeV/nucleon heavy-ion testing at Lawrence Berkeley National Laboratory described in Section C.2, albeit amended to reflect the greater interactivity between the DUT and the experimenter in laser testing. The amendments include:

- Since experimenters are stationed beside the experiment and DUT at this venue, there is no need for remote control of equipment. Hence the KVM switch between the computer and logic analyzer is omitted (instead, operators directly control both pieces of equipment in parallel), as is the GPIB interface for the multimeters monitoring the DUT supplies (instead, operators directly observe the multimeter readings). Naturally, the cables between the test board and the test equipment are thus much shorter.

- To accomplish supply monitoring, a series ammeter is imposed between the linear regulator output and DUT supply. This method is in contrast to the series-resistor-with-parallel-voltmeter technique used in Section C.2.3. In general, when long cable lengths are used, series-current monitoring is noticeably more noisy than parallel-voltmeter monitoring (see footnote 9 of Appendix C). However, the shorter cable lengths of this experiment render the noise increase negligible.

- In many cases, different, though functionally equivalent, test equipment is used. Specifically: 1) a Hameg HM8040-2 triple power supply provides board power

(instead of an HP6205B dual power supply), and 2) an Agilent 33250A provides the clock signal (instead of an Stanford Research Systems CG635).

E.2.1 Pulsed-Laser Source

This experiment uses a 590-nm, actively mode-locked, cavity-dumped dye laser to inject energy into the DUT. This laser is described more extensively by *Moss et al.* [1995]. The laser is concentrated to a spot size of approximately 2-μm diameter and provides pulses at a rate of 1 MHz with the pulse energy adjusted by a series of four optical attenuators.[2] No attempt is made to synchronize the laser pulses with the DUT clock. The laser is normally incident upon the front surface of the silicon, that is, incident at a perpendicular angle to the surface of the chip whereon the transistors are fabricated.

To probe the DUT, the test board is mounted on a motorized chuck. By manually moving the chuck in the horizontal plane of the chip, experimenters scan the laser through the DUT. An integrated, high-magnification camera system allows experimenters to see the location of the laser on the silicon in real-time (see Figure E.1 for an example).

E.3 Procedure

Laser testing is often much more ad-hoc than other types of single-event testing, with tests and procedures evolving as the experiment progresses. This section, though, focuses on the procedures used to identify regions of the DUT sensitive to soft errors.

For this objective, a DC value is input to the DUT. The laser is then scanned over regions of the SVADC-1 believed most likely to create upsets, such as operational amplifier biasing networks and switched-capacitor amplifier switches. In certain regions the $V_{\mathrm{DD,IO}}$ current increases: investigations showed that these increases

[2]Furthermore, the laser output is gated by a 93 Hz square wave of ∼50% duty cycle. That is, for half of this square wave the laser output is "on" and passes through to the DUT, whereas for the other half the laser output is "off" and no laser energy passes through to the DUT.

correspond to increased activity in the DUT output bits due to laser-induced upsets.[3] Thus as the laser is scanned over the chip, the DUT supplies are carefully monitored and, for each sensitive region identified by $V_{DD,IO}$ current increase, the location of greatest $V_{DD,IO}$ increase determined. A visual record is made of this location via a screen capture of the camera system output, and the DUT output is captured over a variety of laser energies at this location.

E.4 Results and Analysis

Before discussing the salient results, it is worthwhile to note some of the limitations of this experiment. First, due to time limitations, the entire DUT is not scanned, rather, only specific regions are probed. Furthermore, only regions with little metal coverage are probed:[4] this choice immediately omits regions such as the I/O circuitry and much of the internal digital logic that are heavily filled in the upper metals. Hence it should be recognized that the pulsed-laser testing is not as comprehensive as the heavy-ion testing of Appendices C and D. Second, while the procedure described in Section E.3 identifies scanned regions that are sensitive to upset, it does not necessarily identify the remaining scanned regions as being resistant to upset. While these remaining regions could indeed be upset-resistant, on the other hand it could be that local changes in metal coverage simply prevent sufficient energy deposition into the substrate in these regions. Finally, such local changes in metal coverage also limit comparative studies of different sensitive regions, for example, upset "thresholds" at different regions cannot be compared without more detailed studies of the effect of metal coverage on the injected energy.

Nonetheless, the pulsed-laser testing enables identification of many upset-sensitive locations. Many of these locations are expected, such as the aforementioned operational amplifier bias transistors and switched-capacitor amplifier switches. A perhaps less obvious and more interesting result, though, is shown in Figure E.1; the

[3] Notably, the $V_{DD,IO}$ current returns to its unirradiated value upon removal of the laser source: the increases are thus not self-sustaining and hence not latchup events.

[4] As metal reflects the laser beam, in regions with high metal coverage the laser energy injected into the silicon is likely negligible.

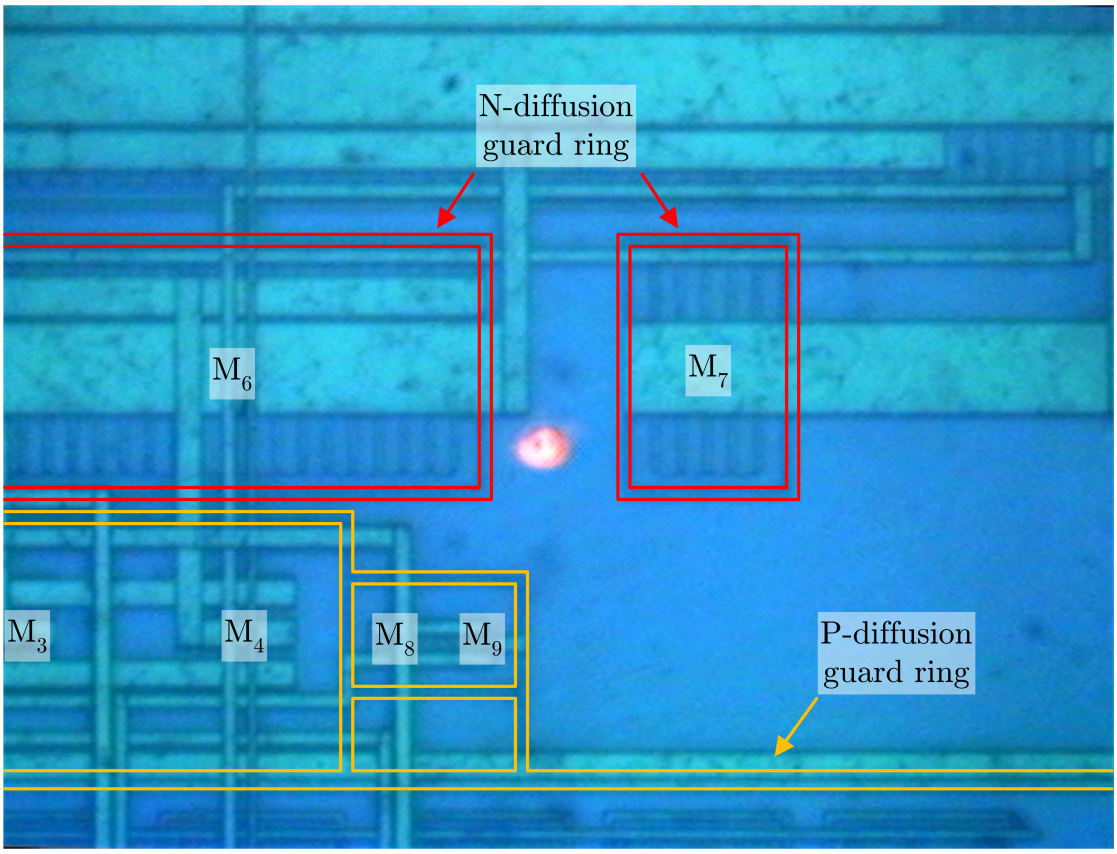

Figure E.1: An upset-sensitive location discovered by laser probing as captured by the integrated camera system. Laser beam appears as bright pink spot in center of figure. Nearby devices and guard rings labeled. Location occurs within operational amplifier of calibrated stage residue amplifier: for corresponding schematic, see Figure E.2.

identified devices are shown in the schematic of Figure E.2. Here, the laser probes a region of bulk silicon between the transistors of an operational amplifier. Note that no active devices are illuminated and, furthermore, that the nearby active devices are fully guard-ringed. However, laser energy injected into this region nonetheless causes noticeable upsets. Most likely, the energy deposited by the laser migrates to and is absorbed by the nearby guard rings. This results in local fluctuations to the supply and ground lines which in turn cause upsets. Hence, in addition to regions near active devices, regions of otherwise-fallow silicon can be upset sensitive. A more robust design, then, should consider imposing additional guard-rings around these

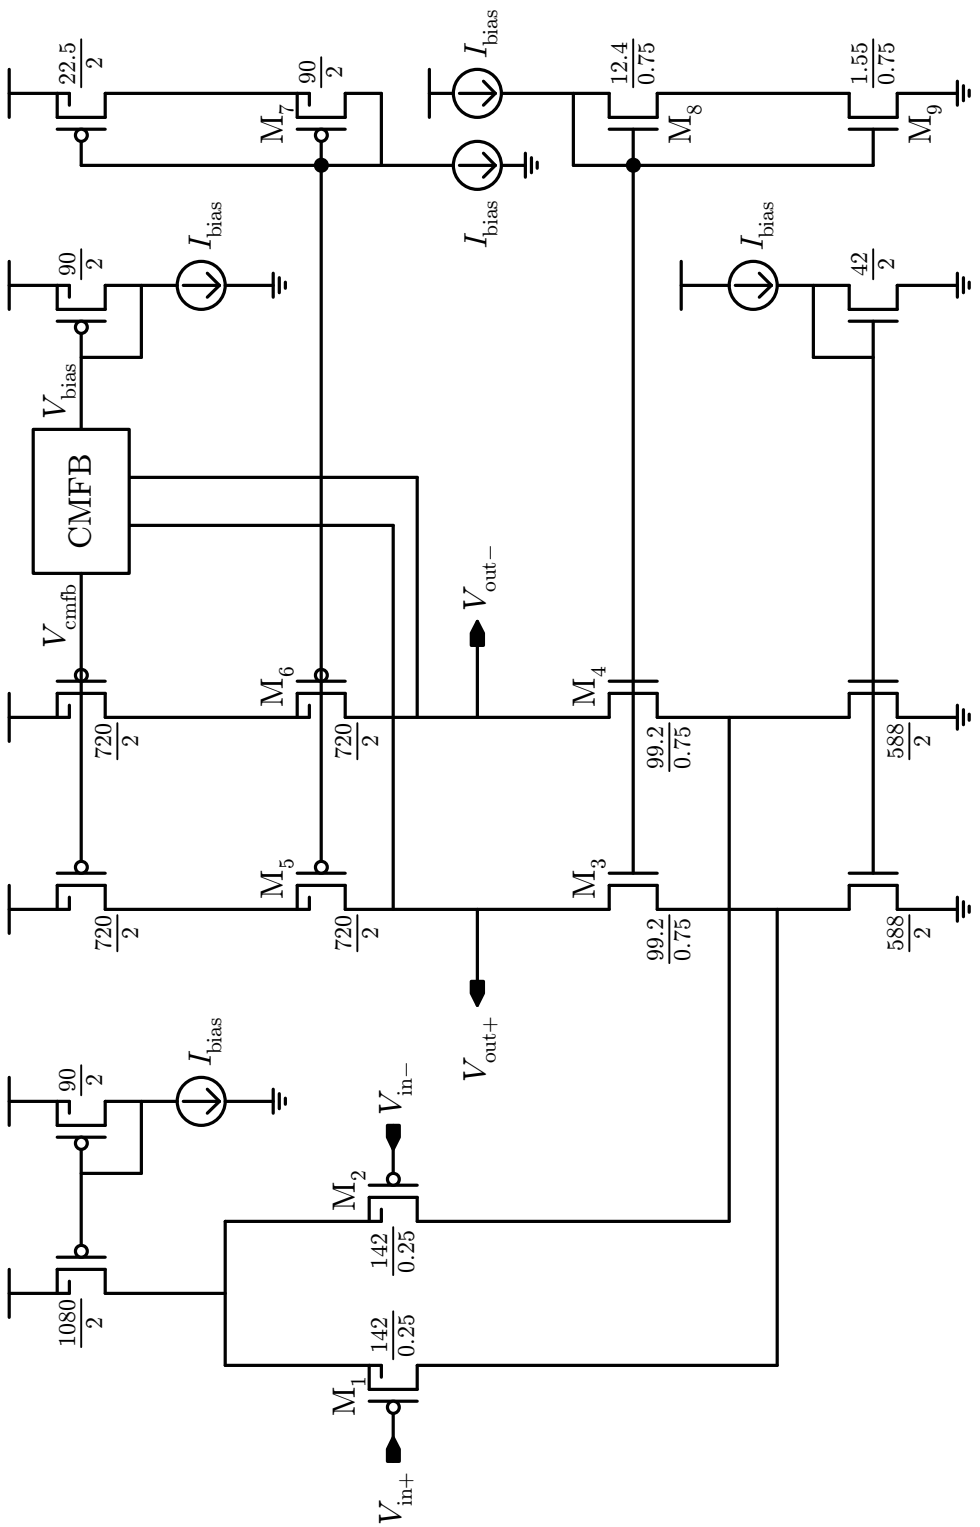

Figure E.2: Corresponding schematic for Figure E.1 with nearby devices labeled. For further information (such as I_{bias} value and supply connections), consult Figure 5.29.

fallow regions, and furthermore connecting these guard rings to ground near the main supply metal lines, distant from the local transistor supplies.

Finally, it should be noted that the DUT displays no latchup during all laser testing, including probing with laser energies as high as 4.29 nJ/pulse.

Appendix F

NMOS Total-Dose Testing, ^{60}Co γ-ray

This appendix describes total-dose testing[1] of the NMOS test transistors on the SVADC-1 chip. These test devices are seen at the bottom of the chip micrograph of Figure 5.40. The set includes ten NMOS transistors, conceptually organized into five pairs, each pair composed of the same device drawn in either standard or enclosed terminal fashion. The transistors are culled from the converter portion of the SVADC-1, where they serve in the CMOS switches of the switched-capacitor circuits.

In this experiment, the transistors are irradiated by γ-rays from a ^{60}Co source in logarithmic steps up to 2 Mrad(Si) under varying irradiation biases. After irradiation, each device is characterized by key parameters including a I_{D}-V_{GS} curve (from which threshold voltage and transconductance are derived), leakage current, and channel resistance. In addition, devices are characterized following an unbiased, 65-hour anneal at room temperature. In total, four decades of devices are tested, with each decade representing a different irradiation bias condition.

The experiment was conducted from 10 through 13 October 2006, and on 16 October 2006, at the ^{60}Co radiation test facility at The Aerospace Corporation in El Segundo, California. The experimenters were Charles C. Wang with Stanford

[1] An introduction to radiation testing, including total-dose testing, is found in Appendix A.

University, and Everett E. King, Mark P. Zakrzewski, and James L. Roeder with The Aerospace Corporation. Stephanie Brown with The Aerospace Corporation also provided support, including bonding of the test chips in their respective packages. The experiment was coordinated by James L. Roeder of The Aerospace Corporation, principal investigator of the PARX project.

F.1 Device-Under-Test

The devices-under-test (DUTs) consist of five pairs of devices, each of which represents the same nominal NMOS device drawn in either standard or enclosed terminal fashion. This choice enables comparisons between the two layout styles. A summary of the DUTs is given in Table F.1. The DUTs are drawn from actual devices used within the SVADC-1 converter,[2] with particular selections made to reflect a wide breadth of device sizings. Note that both low voltage (LV) and high voltage (HV) devices are included.[3]

Each DUT is enclosed in a standardized frame to reduce variation between devices. The frame consists of independent gate, source, drain, and bulk pads to enable independent testing of each DUT. A schematic of the frame is given in Figure F.1(a), a corresponding layout is shown in Figure F.1(b). Note that the frame provides ESD-protection of the DUT gate via diodes to bulk. Figure F.1 also shows the source and drain naming convention for the DUTs. While technically the source and drain terminals of a MOSFET depend on the bias placed thereupon, to simplify matters, for the enclosed devices the enclosed terminal is referred to as the source and the non-enclosed terminal the drain. This convention is extended to the standard devices as well: the standard device terminal equivalent to the enclosed device terminal is referred to as the source, the other as the drain. To reduce ambiguity, for all further figures an enclosed terminal device (with the enclosed terminal indicating the source) is used to represent the DUT.

[2]With the exception of devices 3 and 4. These devices are included as variations of devices 5 and 6—having the same drawn width, but a different number of fingers—to isolate possible differences due to the latter.

[3]The differences between these device types described in Sections 2.3.3 and 5.1.

Device number	Type	Device dimensions[a]					Enclosed terminal	Converter location
		W_{drawn}	f	W	L	t_{ox}		
		[μm]		[μm]	[μm]	[nm]		
1	LV	1.3	1	1.3	0.25	5.3	[none]	[not used]
2	LV	1.3	1	1.3	0.25	5.3	Source	CMOS switch type a[b]
3	LV	3	1	3	0.25	5.3	[none]	[not used]
4	LV	3	1	3	0.25	5.3	Source	[not used]
5	LV	3	2	6	0.25	5.3	[none]	[not used]
6	LV	3	2	6	0.25	5.3	Source	CMOS switch type c[b]
7	LV	8.5	4	34	0.25	5.3	[none]	[not used]
8	LV	8.5	4	34	0.25	5.3	Source	CMOS switch type e[b]
9	HV	2.5	2	5	0.4	6.4	[none]	[not used]
10	HV	2.5	2	5	0.4	6.4	Source	Bootstrapped switch[c]

[a] Device dimensions given by drawn channel width (W_{drawn}), number of fingers (f), total channel width (W), channel length (L), and typical gate oxide thickness (t_{ox}).
[b] CMOS switch types described in Table 5.1.
[c] Eight of these devices in parallel implement the track-and-hold amplifier bootstrapped switch M_S shown in Figure 5.18.

Table F.1: Test NMOS transistors included on the SVADC-1 chip.

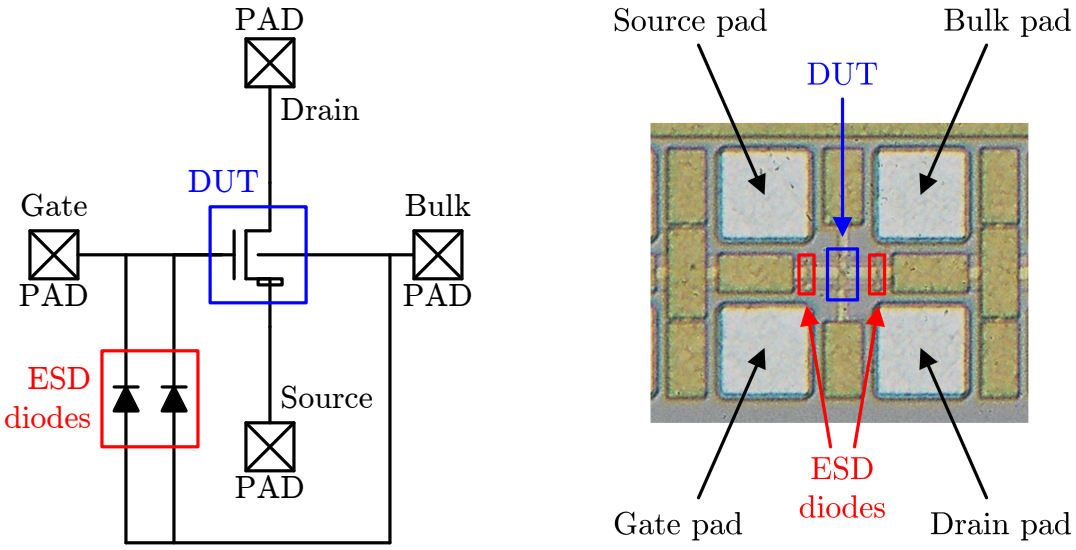

(a) Frame schematic. (b) Frame layout.

Figure F.1: DUT frame (a) schematic and (b) layout.

By convention, a group of all ten DUTs of Table F.1 is referred to as a decade. For this experiment, each decade is represented by two instances of the SVADC-1 chip. Each instance is bonded into a 24-pin dual-inline package (DIP), with the first instance bonded to provide access to devices 1 through 6, and the second instance bonded to provide access to devices 7 through 10. These packages are chosen since 1) they are known to be compatible with the semiconductor parameter analyzer test fixtures used in this experiment, and 2) zero-insertion force (ZIF) sockets for these packages are readily available.

In all, four decades of DUTs are tested in this experiment. All packages are de-lidded throughout the experiment. All testing is performed at room temperature.

F.2 Irradiation

During radiation testing, each of the four decades of DUTs is subjected to a different irradiation bias. These biases are given in Table F.2. Naturally, worst-case irradiation bias is included. Also included are an all-ground irradiation bias (which should

Decade number	Device biasing				Description
	Gate	Source	Drain	Bulk	
	[V]	[V]	[V]	[V]	
1	0	0	0	0	All-ground decade
2	V_{DD}	0	0	0	Worst-case decade
3	V_{DD}	0	1.75	0	Drain-high decade
4	V_{DD}	1.75	0	0	Source-high decade

Table F.2: Irradiation biases of DUTs, by decade. All devices DC biased. V_{DD} is 2.5 V for LV, and 3.3 V for HV, devices.

experience minimal radiation damage), and drain-high and source-high irradiation biases. The latter two stress the inherent asymmetry of the enclosed terminal device channel: a large electric field occurs across the gate oxide near the 0 V terminal, whereas a small field occurs near the 1.75 V terminal. The 1.75 V voltage is chosen since this voltage is the maximum of the expected signal range for the device in the SVADC-1 converter.[4]

A custom irradiation bias board, designed and constructed by The Aerospace Corporation with assistance from Charles C. Wang of Stanford University, houses all four decades of DUTs simultaneously during irradiation. The board is powered by three supplies: an HP6214B DC power supply outputting 2.5 V, an HP6214B DC power supply outputting 3.3 V, and a Tektronics TM500 series power supply outputting 1.75 V.[5] Each supply powers all the DUT terminals on the irradiation bias board requiring the specified voltage.

The DUTs are irradiated by γ-rays from a ^{60}Co source at the γ-ray facility of The Aerospace Corporation in El Segundo, California, in logarithmic dose steps up to a total dose of 2 Mrad(Si). The complete irradiation schedule is given in Table F.3.

[4] An alternative is to use 0.75 V, the minimum of the expected input signal range. However, assuming the other channel terminal is connected to 0 V, a 1.75 V terminal voltage generates a larger contrast in the local oxide electric field than a 0.75 V terminal voltage, better magnifying any underlying asymmetries.

[5] The Tektronics supply providing 1.75 V failed between the 100 krad(Si) and 200 krad(Si) irradiations and was replaced by a Kikusui PAB 25-1TR power supply.

Date	Start time	End time	Exposure	Dose rate	Irradiated dose	Accumulated dose	Assigned dose
			[min]	[rad(Si)/sec]	[krad(Si)]	[krad(Si)]	[krad(Si)]
10 Oct 2006	08:56	08:58	2.5	6.63	0.99	0.99	1
10 Oct 2006	13:13	13:23	10	6.63	3.98	4.97	5
10 Oct 2006	14:51	15:03	12.5	6.63	4.97	9.95	10
11 Oct 2006	08:24	08:50	25	6.63	9.95	19.90	20
11 Oct 2006	14:22	14:29	6.87	72.80	30.01	49.90	50
12 Oct 2006	08:30	08:41	11.45	72.80	50.02	99.92	100
12 Oct 2006	13:22	13:45	22.89	72.80	99.99	199.91	200
12 Oct 2006	16:44	17:53	68.68	72.80	300.01	499.92	500
13 Oct 2006	08:17	10:15	114.46	72.80	499.99	999.91	1000
13 Oct 2006	13:33	17:22	228.92	72.80	999.98	1999.99	2000

Table F.3: Irradiation schedule for total-dose testing of the DUTs by ^{60}Co γ-rays. All four decades of DUTs are irradiated simultaneously.

Dose step	At power supply			On bias board		
	1.75 V	2.5 V	3.3 V	1.75 V	2.5 V	3.3 V
[krad(Si)]	[V]	[V]	[V]	[V]	[V]	[V]
1	1.915	2.499	3.300	1.751	2.496	3.300
5	1.943	2.496	3.300	1.747	2.493	3.300
10	1.933	2.496	3.300	1.736	2.493	3.300
20	1.922	2.496	3.300	1.747	2.492	3.300
50	1.884	2.496	3.300	1.682	2.493	3.300
100	1.951	2.496	3.300	1.755	2.493	3.300
200	1.941	2.497	3.300	1.743	2.493	3.300
500	1.940	2.496	3.300	1.743	2.491	3.300
1000	1.940	2.496	3.300	1.742	2.491	3.300
2000	1.940	2.496	3.300	1.743	2.491	3.300

Table F.4: Measurements taken before irradiations to confirm DUT functionality. Note that, as the 1.75 V supply draws current (being connected to device channels), the 1.75 V power supply is tuned to a higher voltage to compensate for resistive voltage drops along cabling. Supplies are set prior to 1 krad(Si) dose, afterwards, they are checked but not adjusted.

Notably, two dose rates are used: a slower rate for the lower dose steps (to increase the dose time and provide margin for chamber shutter speeds), and a faster rate for the higher dose steps (to maintain reasonable dose times at large doses).[6] To confirm proper biasing, prior to exposure the voltage of each supply is checked both at the supply and on the board. These measurements are summarized in Table F.4.

[6] Also notably, irradiation took place over several days. The Military Standard [*MIL-STD-883G*, 2006, Method 1019.7] prescribes the maximum duration allowed from the end of one exposure to the start of the next to be two hours. However, in this experiment, the use of longer durations is justified in that, although annealing may take place between dose steps, given the exponential growth in dose at each step, any annealed radiation damage is likely reinstated in the next exposure. For the sake of completeness it should be noted that between dose steps, except when it is being characterized, each DUT is left in an unbiased state at room temperature.

F.3 Characterization

Following irradiation, each DUT is characterized via a suite of measurements. Overall, the breadth of the suite, as well as the fineness of each measurement in the suite, is ultimately limited by the allowed measurement time: the suite is designed so that all four decades of DUTs can be characterized within 2 hours.

F.3.1 Equipment

Two test stations execute the measurement suite. The first station features an Agilent 4156C Semiconductor Parameter Analyzer with an installed Agilent 41501B Expander operated in conjunction with an HP 16442A Test Fixture. This station characterizes the DUTs of the all-ground and worst-case irradiation bias decades. The second station features an Agilent 4155C Semiconductor Parameter Analyzer with an installed HP 41501A Expander operated in conjunction with an Agilent 16442A Test Fixture. This station characterizes the DUTs of the drain-high and source-high irradiation bias decades. To speed data acquisition, the stations are operated in parallel by separate operators.

The semiconductor parameter analyzers allow the operator to set the voltage or current of each DUT terminal and measure the resulting current or voltage. To maintain test accuracy, the measurement suite is programmed into the analyzer and executed on each DUT. Following execution, the measured results are stored in a binary file format custom to the analyzers for future processing.

F.3.2 Measurements

Each DUT is characterized by three measurements: an I_{D}-V_{GS} curve (which provides a general device description, including transconductance and threshold voltage), a leakage-current measurement (which describes the off-state of the device when used as a switch), and a channel-resistance measurement (which describes the on-state of the device when used as a switch). Note that, while quantities may be swept over the course of measurements, all measurements are fundamentally of DC values: no AC

characterization is undertaken in this experiment.

Many of these measurements require sweeps of device voltages or currents. For convenience, this appendix adopts the notation $V_{\mathrm{X}} = [a\,, b\,, c]$ to denote that the parameter V_{X} is swept through the values a, b, and c. For longer sweeps, the notation $V_{\mathrm{Y}} = [d : e : f]$ is used to denote sweeps of the parameter V_{Y} from d to f (inclusive) via increments of e. It is notable that two-parameter sweeps are implemented as an outer sweep followed by an inner sweep on the semiconductor parameter analyzers: for each outer sweep value the entire inner sweep is executed,[7] which often leads to data collection under extraneous bias conditions.

F.3.2.1 I_{D}-V_{GS} Curve

The setup for the I_{D}-V_{GS} curve measurement is depicted in Figure F.2. The drain-to-source voltage is set to a low value (0.1 V, an oft-used value) and the gate voltage is swept from a negative voltage to the rail V_{DD}. The drain current is measured at each gate voltage. While lower gate voltages better capture the sub-threshold regime of the DUT, the ESD diodes on each DUT gate impose a lower limit: -0.5 V is chosen as an acceptable compromise.[8] Parameters such as transconductance and threshold voltage are subsequently derived from the I_{D}-V_{GS} curve; see Sections F.4.1 and F.4.2, respectively.

F.3.2.2 Leakage Current

To measure leakage current, the gate voltage is set to 0 V and maximum electric field imposed across the DUT channel by setting the source to 0 V and the drain to the rail. To reflect the asymmetry of the enclosed terminal devices, the opposite field—wherein the source is set to the rail and the drain to 0 V—is also measured. The final measurement setup is shown in Figure F.3. For each channel condition, the drain current is measured although, given the rise in leakage currents at 200 krad(Si), from that dose step onwards the currents of all four terminals are measured.

[7] Similar to a for loop within a for loop in C programming.

[8] Indeed, the ESD diodes did not show significant current until the 2 Mrad(Si) dose step, and then only at the -0.5 V gate voltage step.

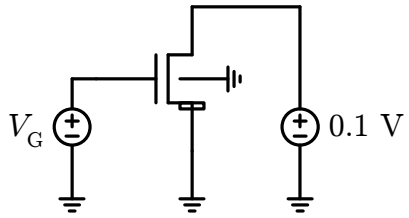

Sweep:
$V_G = [-0.5 : 0.025 : V_{DD}]$ V
Measure:
I_D

Figure F.2: I_D-V_{GS} curve measurement setup. V_{DD} is 2.5 V for LV, and 3.3 V for HV, devices.

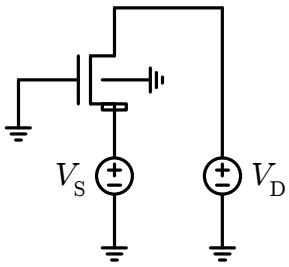

Sweep:
$V_S = [0\ ,\ V_{DD}]$ V (outer)
$V_D = [0\ ,\ V_{DD}]$ V (inner)
Measure:
I_D (for doses $\leq$200 krad(Si))
I_D, I_S, I_G, I_B (for doses >200 krad(Si))

Figure F.3: Leakage current measurement setup. V_{DD} is 2.5 V for LV, and 3.3 V for HV, devices.

F.3.2.3 Channel Resistance

To measure channel resistance, it is desirable to imitate the condition of the devices when they are used as switches in the SVADC-1 converter. The ideal measurement is illustrated in Figure F.4(a): the gate voltage is set to the rail and the channel voltage swept over the expected input signal range of 0.75 V to 1.75 V. A small differential voltage—say, 10 mV—is imposed across the channel and the resulting channel current measured. On the semiconductor parameter analyzers, though, Figure F.4(a) is implemented with outer and inner sweeps on the DUT source and drain as shown in Figure F.4(b). In addition, the range of channel voltages is expanded to 0.65 V to 1.85 V. Finally, given the asymmetric enclosed terminal devices, two measurements are taken to capture the channel resistance when V_{DS} is either positive or negative.

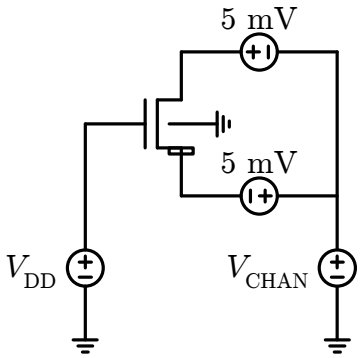

Sweep:
$V_{CHAN} = [0.75 : 0.1 : 1.75]$ V
Measure:
I_D

(a) Ideal setup.

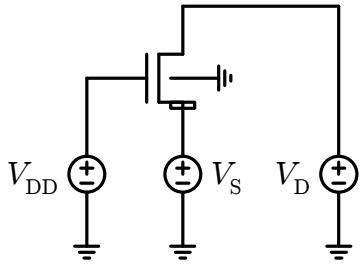

For $V_{DS} > 0$:
Sweep:
$V_D = [0.655 : 0.1 : 1.855]$ V (outer)
$V_S = [0.645 : 0.1 : 1.845]$ V (inner)
Measure:
I_D

For $V_{DS} < 0$:
Sweep:
$V_D = [0.645 : 0.1 : 1.845]$ V (outer)
$V_S = [0.655 : 0.1 : 1.855]$ V (inner)
Measure:
I_D

(b) Implemented setup.

Figure F.4: Channel resistance measurement setup, including (a) ideal setup and (b) implemented setup.

Measurement	Set quantities[a]				Sweeps		Measured quantities
	Gate	Source		Drain	Outer sweep	Inner sweep	
	[V]	[V]		[V]	[V]	[V]	
I_D-V_{GS} curve	Outer	0		0.1	$[-0.5 : 0.025 : V_{DD}]$	[n/a]	I_D
Leakage current[b]							
For doses $\leq$200 krad(Si)	0	Outer		Inner	$[0\,, V_{DD}]$	$[0\,, V_{DD}]$	I_D
For doses $>$200 krad(Si)	0	Outer		Inner	$[0\,, V_{DD}]$	$[0\,, V_{DD}]$	I_D, I_S, I_G, I_B
Channel resistance[c]							
For $V_{DS} > 0$ V	V_{DD}	Outer		Inner	$[0.645 : 0.1 : 1.845]$	$[0.655 : 0.1 : 1.855]$	I_D
For $V_{DS} < 0$ V	V_{DD}	Outer		Inner	$[0.655 : 0.1 : 1.855]$	$[0.645 : 0.1 : 1.845]$	I_D

[a] Set quantities set to either indicated value or swept as per outer or inner sweep. Bulk voltage set to 0 V for all measurements.
[b] Leakage current measurement expanded after 200 krad(Si) dose step.
[c] Both channel resistance measurements executed for all DUTs at all dose steps.

Table F.5: Measurement suite for characterizing DUT performance. V_{DD} is 2.5 V for LV, and 3.3 V for HV, devices.

F.3.3 Summary

The measurement suite is summarized in Table F.5. Note that, at each dose step, each DUT is assessed by an I_D-V_{GS} curve, a leakage-current measurement (the measured quantities of this measurement changing at 200 krad(Si)), and both channel-resistance measurements (i.e., for both $V_{DS} > 0$ V and $V_{DS} < 0$ V).

The complete characterization schedule for the radiation testing of this experiment is given in Table F.6. The schedule includes an anneal dataset, taken after the DUTs have been left in an unbiased state at room temperature for 65 hours. Note that the characterization of the all-ground and worst-case irradiation bias decades—that is, the results quoted in this dissertation—both complete within 2 hours after the end of irradiation. The exception is the 1 Mrad(Si) dose step, where characterization was delayed by a forced evacuation due to a fire alarm.

F.4 Metric Definitions

The definitions of leakage current and channel resistance are readily apparent from their descriptions in Sections F.3.2.2 and F.3.2.3, respectively. The definitions of metrics derived from the I_D-V_{GS} curve, though, bear more explanation.

F.4.1 Transconductance Definition

The transconductance is computed by taking the derivative of the I_D-V_{GS} curve. To reduce noise, the I_D-V_{GS} is first decimated by a factor of two (increasing the V_{GS} increment to 50 mV). The derivative is then computed as:

$$g_m(V_{GS}) = \frac{I_D(V_{GS} + \delta) - I_D(V_{GS} - \delta)}{2\delta} \tag{F.1}$$

where $\delta = 50$ mV.[9] The maximum transconductance over V_{GS}, that is, $g_{m,max}$, is often quoted to characterize the entire curve.

[9] The derivative of Equation (F.1) is the definition used by the semiconductor parameter analyzers themselves in their own transconductance computations. However, an alternate definition can be

Dose step [krad(Si)]	Date	Irradiation stop time	Irradiation bias	Characterization		Delay[a]
				Start time	End time	
1	10 Oct 2006	08:58	All-ground	09:24	10:14	01:16
			Worst-case	10:17	10:51	01:53
			Drain-high	09:22	11:06	02:08
			Source-high	10:04	11:01	02:03
5	10 Oct 2006	13:23	All-ground	13:32	14:00	00:37
			Worst-case	14:03	14:33	01:10
			Drain-high	13:32	14:00	00:37
			Source-high	14:02	14:34	01:11
10	10 Oct 2006	15:03	All-ground	15:12	15:41	00:38
			Worst-case	15:43	16:12	01:09
			Drain-high	15:12	15:45	00:42
			Source-high	15:46	16:16	01:13
20	11 Oct 2006	08:50	All-ground	09:12	09:45	00:55
			Worst-case	09:47	10:25	01:35
			Drain-high	09:35	10:49	01:59
			Source-high	11:03	12:40	03:50
50	11 Oct 2006	14:29	All-ground	14:42	15:06	00:37
			Worst-case	15:08	15:37	01:08
			Drain-high	14:44	15:52	01:23
			Source-high	16:31	17:18	02:49
100	12 Oct 2006	08:41	All-ground	08:53	09:18	00:37
			Worst-case	09:21	09:48	01:07
			Drain-high	08:53	09:47	01:06
			Source-high	09:59	10:56	02:15

Continued on next page...

...Continued from previous page.

Dose step [krad(Si)]	Date	Irradiation stop time	Irradiation bias	Characterization		Delay[a]
				Start time	End time	
200	12 Oct 2006	13:45	All-ground	13:57	14:23	00:38
			Worst-case	14:25	14:53	01:08
			Drain-high	13:57	14:47	01:02
			Source-high	14:52	15:55	02:10
500	12 Oct 2006	17:53	All-ground	18:01	18:24	00:31
			Worst-case	18:27	18:55	01:02
			Drain-high	18:03	18:55	01:02
			Source-high	18:59	19:27	01:34
1	13 Oct 2006	10:15	All-ground	11:06	11:30	01:15
			Worst-case	12:08	12:39	02:24
			Drain-high	11:07	12:03	01:48
			Source-high	12:44	13:22	03:07
2	13 Oct 2006	17:22	All-ground	17:31	17:54	00:32
			Worst-case	17:58	18:20	00:58
			Drain-high	17:31	18:11	00:49
			Source-high	18:13	18:41	01:19
Anneal	16 Oct 2006		All-ground	08:13	08:46	63:24
			Worst-case	08:52	09:24	64:02
			Drain-high	10:16	10:51	65:29
			Source-high	10:59	11:37	66:15

[a] Delay is the time between the end of irradiation and the end of characterization. Given in format of [hours]:[minutes].

Table F.6: Characterization schedule for total-dose testing of the DUTs by ^{60}Co γ-rays.

F.4.2 Threshold Voltage Definition

While the threshold voltage is a fundamental quantity—being defined as the gate voltage at the onset of channel inversion [*Pierret*, 1990, p. 70]—it is difficult to measure directly. Several methods have been proposed, all of which assume a transistor model and then propose measurements to isolate the threshold voltage of that model [*Ortiz-Conde et al.*, 2002].[10] This dissertation adopts the extrapolated threshold voltage.

The extrapolated threshold voltage is computed by drawing a tangent line to the I_D-V_{GS} curve at the point of maximal slope (i.e., at $g_{m,max}$) and measuring the x-intercept of that tangent line. This calculation is illustrated in Figure F.5. A small correction factor based on V_{DS} is then applied to give the final threshold voltage. Specifically, if $V_{GS}(g_{m,max})$ and $I_D(g_{m,max})$ are the voltage and current at $g_{m,max}$, respectively, then the extrapolated threshold voltage V_{TH} is:

$$V_{TH} = V_{GS}(g_{m,max}) - \frac{I_D(g_{m,max})}{g_{m,max}} - \frac{V_{DS}}{2} \tag{F.2}$$

The extrapolated threshold voltage has many benefits. First, it is fairly simple to compute, but has nonetheless been empirically shown to maintain comparable accuracy to many other, more computationally intensive, methods [*Ortiz-Conde et al.*, 2002]. Second, it is the method used by the semiconductor parameter analyzers themselves, providing a cross-check on the results from data processing. Finally, the extrapolated threshold voltage is the standard method adopted by The Aerospace Corporation (and indeed in much of the radiation testing community in general),

used which relies on directly adjacent points:

$$g_m(V_{GS}) = \frac{I_D(V_{GS}) - I_D(V_{GS} - \delta)}{\delta}$$

Attempts were made to compute the transconductance using either this adjacent-point derivative or the split-point derivative of Equation (F.1), executed on the I_D-V_{GS} curve at either 25 mV or 50 mV increments, and with the possible inclusion of median filtering (of filter orders 3 through 7). These attempts showed that the method of Section F.4.1 is at least comparable (in not in some cases superior) to these many others.

[10] Indeed, often the fidelity of the computed threshold voltage is as much a function of the fidelity of the device to the model as to the actual threshold voltage itself.

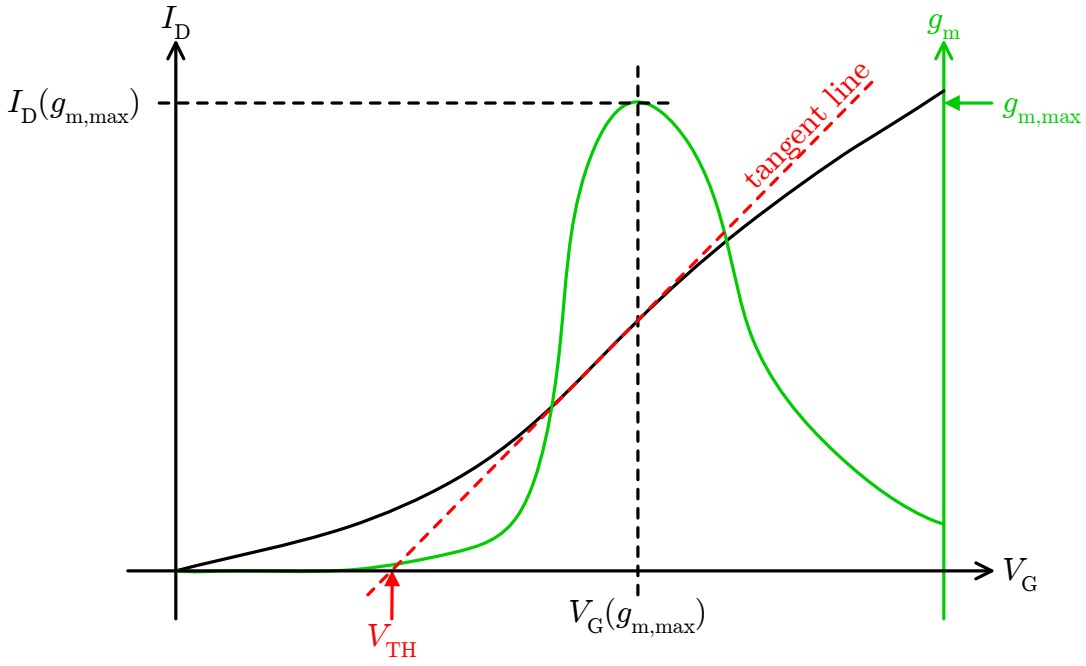

Figure F.5: Illustration of the derivation of extrapolated threshold voltage. Note that illustration shows V_{TH} derivation without V_{DS} correction.

enabling better comparison with other, published results.

F.5 Results and Analysis

The key results of the radiation evolution of the transconductance, threshold voltage, and leakage current are presented in the dissertation proper, specifically, see Sections 2.3.3.3; 2.3.3.1; and 2.3.3.2 and 5.2.2.2, respectively. This appendix, then, only addresses the results of the channel resistance.

The radiation evolution of the channel resistance R_{CHAN} under worst-case irradiation bias is shown in Figure F.6 for the $W/L = 6/0.25$ device, drawn in both standard (left) and enclosed terminal (right) fashion. The performance of this device is indicative of all other LV devices. The R_{CHAN} is plotted for varying channel conditions. Note that the R_{CHAN} shows little change between the $V_{DS} < 0$ or $V_{DS} > 0$ conditions. However, it shows strong variation with channel voltage.

In general, the R_{CHAN} at lower channel voltages—indicated by the 0.75 V and 1.25 V channel voltages—is more pertinent since, in the SVADC-1 converter, R_{CHAN} dominates the CMOS switch on-resistance over this range. Two comments regarding the lower channel voltage performance are in order. First, in this regime the R_{CHAN} of the standard and enclosed devices is very similar, validating the use of standard device models for enclosed terminal devices in simulation.[11] Second, this similarity extends to the radiation evolution of both devices, which is additionally very steady: both devices show little R_{CHAN} variation with radiation.

While the device performance over the lower channel voltages is thus well-behaved, at the higher channel voltages—indicated by the 1.75 V channel voltage—greater differences are manifest: not only are the channel resistances of the standard and enclosed devices different, but they also experience much greater shift with radiation. It is likely that the latter arises largely from increased sensitivity to radiation-induced threshold voltage shifts at such low V_{GS}. Lending credence to this interpretation is Figure F.7, which repeats Figure F.6 for the HV, $W/L = 5/0.4$ device: the higher 3.3 V rail translates to higher V_{GS}, resulting in little change in R_{CHAN} with radiation even at the 1.75 V channel voltage.

All in all, in the context of the CMOS switch application for which they are intended, the channel resistances of the tested devices show good accord between standard and enclosed terminal devices, and little change in resistance with dose.

[11] Specifically, the R_{CHAN} at 0.75 V is 254 Ω and 272 Ω for the standard and enclosed terminal devices, respectively (for a 6.8% difference), and the R_{CHAN} at 1.25 V is 470 Ω and 504 Ω, respectively (for a 7.2% difference). The worst-case difference of R_{CHAN} between standard and enclosed terminal device pairs occurs for the LV, $W/L = 1.3/0.25$ devices (i.e., devices 1 and 2), where the device channel is narrowest and corner effects from the enclosure thus have greater impact. For these devices the R_{CHAN} at 0.75 V is 1.20 kΩ and 1.51 kΩ for the standard and enclosed terminal devices, respectively (for a 25.9% difference), and the R_{CHAN} at 1.25 V is 2.19 kΩ and 2.83 kΩ, respectively (for a 29.2% difference).

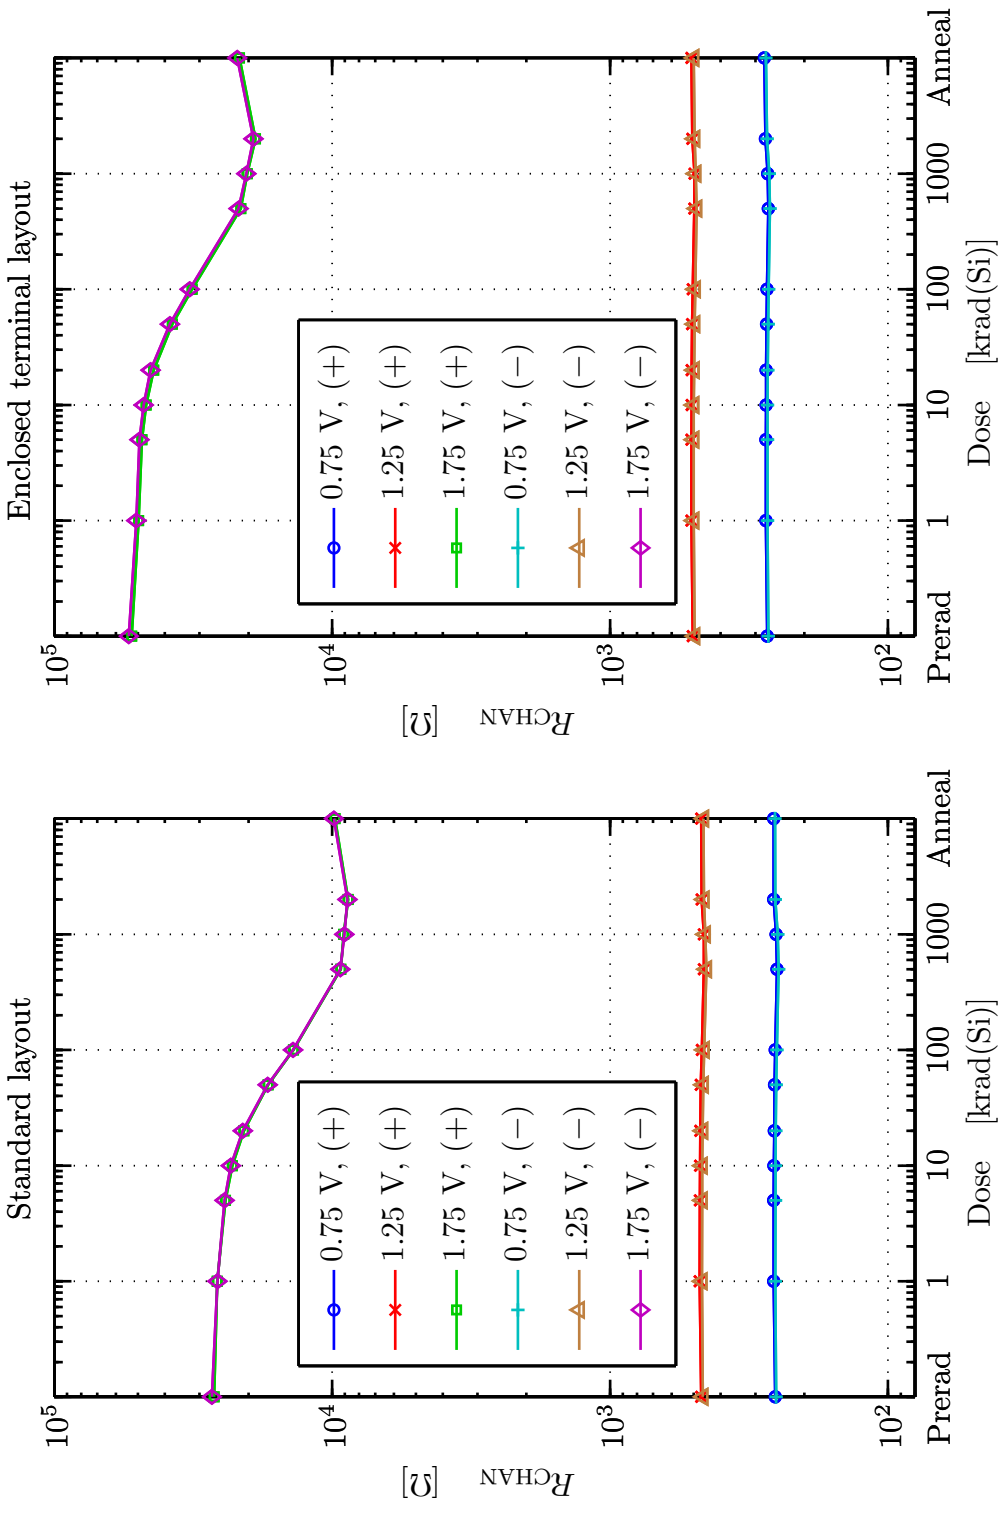

Figure F.6: Measured channel resistance R_{CHAN} versus total dose for an LV $W/L = 6/0.25$ NMOS device from BiCMOS8iED. Select channel voltages shown, corresponding to low (0.75 V), mid (1.25 V), and high (1.75 V) voltages of the SVADC-1 converter signal range, and under $V_{\text{DS}} = 10$ mV and $V_{\text{DS}} = -10$ mV conditions (labeled as (+) and (−), respectively, in the legend). Device irradiated under worst-case bias. Markers indicate measured data points.

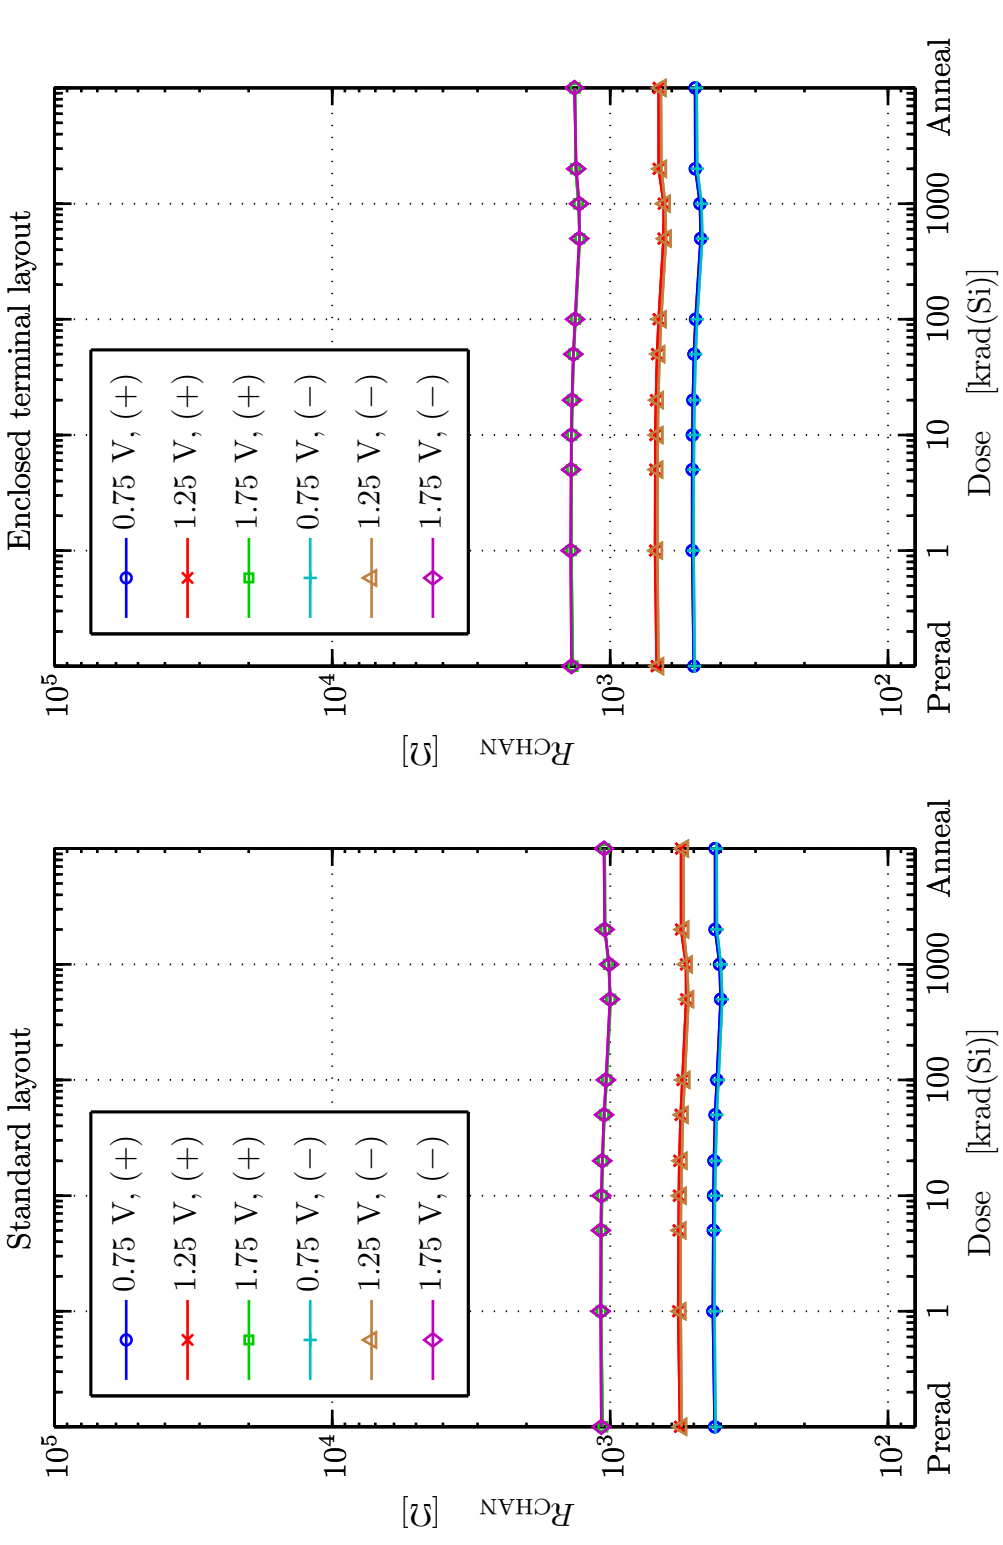

Figure F.7: Measured channel resistance R_{CHAN} versus total dose for an HV $W/L = 5/0.4$ NMOS device from BiCMOS8iED. Select channel voltages shown, corresponding to low (0.75 V), mid (1.25 V), and high (1.75 V) voltages of the SVADC-1 converter signal range, and under $V_{DS} = 10$ mV and $V_{DS} = -10$ mV conditions (labeled as (+) and (−), respectively, in the legend). Device irradiated under worst-case bias. Markers indicate measured data points.

Appendix G

Power Estimate of Plasma Wave Instrument

This appendix presents a back-of-the-envelope calculation of the power consumption of a plasma wave instrument that incorporates the SVADC-1 and the SVLNAE-3, the former being the analog-to-digital converter of this dissertation, the latter the low-noise amplifier and anti-aliasing filter designed by fellow Ph.D. student Benjamin J. Mossawir [*Mossawir et al.*, 2006; *Mossawir*, in preparation]. For convenience, this instrument is referred to as the μPWI here. While the μPWI thus accounts for many practical issues, it should be emphasized that it is intended as a mock-up for power consumption estimation: it is *not* intended as a formal schematic for a plasma wave instrument.

The μPWI captures 3 electric field channels and 3 magnetic field channels, all with wideband receivers. As is shown in this appendix, it consumes an estimated 5.48 W from a spacecraft bus power supply of $\pm$28 V (a standard value). The power consumption of the μPWI is hence comparable to that of the FAST fields instrument (shown in Figure 1.9 and described in Section 1.2.2), which consumes 5.52 W. It is notable, though, that the composition of the two instruments is different: whereas the FAST fields instrument employs a mixture of different receiver architectures, the μPWI uses only wideband receivers and relies on subsequent digital signal processing to balance concerns of data storage and telemetry.

This appendix begins by describing the architecture of the μPWI. All components used in the μPWI are explicitly radiation-hard, or have radiation-hard equivalents, up to a total dose of at least 100 krad(Si). Naturally, the architecture focuses only on those components which contribute significantly to power consumption. The appendix then estimates the total power consumption of the μPWI, taking into account the supply voltage down-conversions required assuming the spacecraft bus provides power via $\pm$28 V supplies.

G.1 Instrument Description

The architecture of the μPWI is shown in Figure G.1. It is composed of 6 wideband receivers—3 for capturing electric field signals, and 3 for capturing magnetic field signals—each of which is designed to capture signals over a 100-Hz to 1-MHz bandwidth at a spurious-free dynamic range (SFDR) of at least 90 dB (assuming 100-Hz/bin spectral resolution). In each wideband receiver, the signals are sampled at 5 MS/s.

The digitized signals from each wideband receiver are processed by an instrument director. The instrument director manages the receivers and interfaces with the spacecraft bus. It also handles data from the housekeeping system. The housekeeping system provides measurements of sensor signals (such as supply voltages and temperatures) made throughout the instrument that report the instrument status.

G.1.1 Wideband Receiver

The architecture of an electric field wideband receiver is shown in Figure G.2. This receiver is a more detailed version of the receiver of Figure 1.11. The analog input signal is captured by an electric dipole antenna and processed by a signal path of a low-noise amplifier (LNA), anti-aliasing filter (AAF), and buffer in preparation for digitization by an analog-to-digital converter (ADC). The digitized signal is then handled by a field-programmable gate array (FPGA), which performs any necessary signal processing and data formatting. The FPGA also manages the receiver and

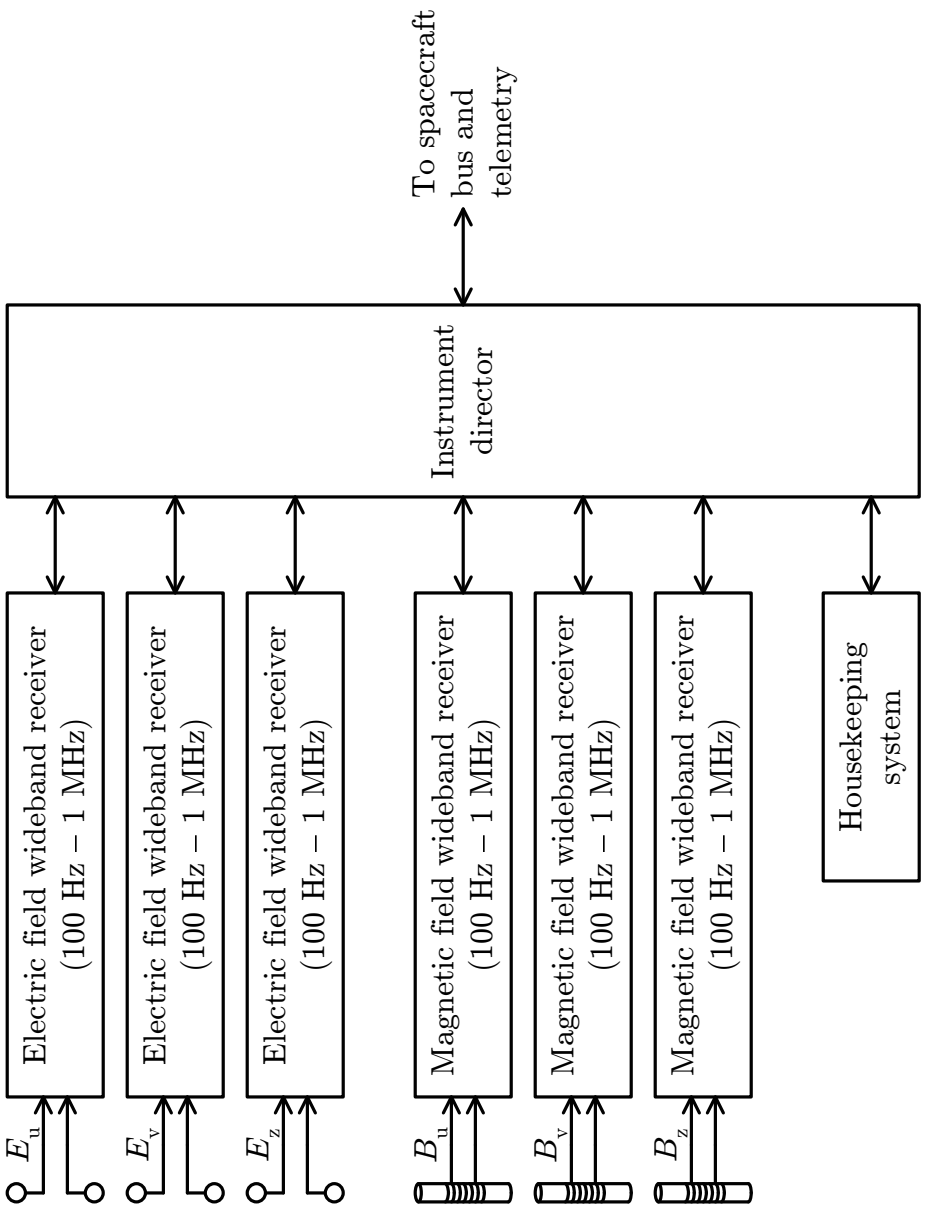

Figure G.1: Architecture of the μPWI instrument used for power estimation.

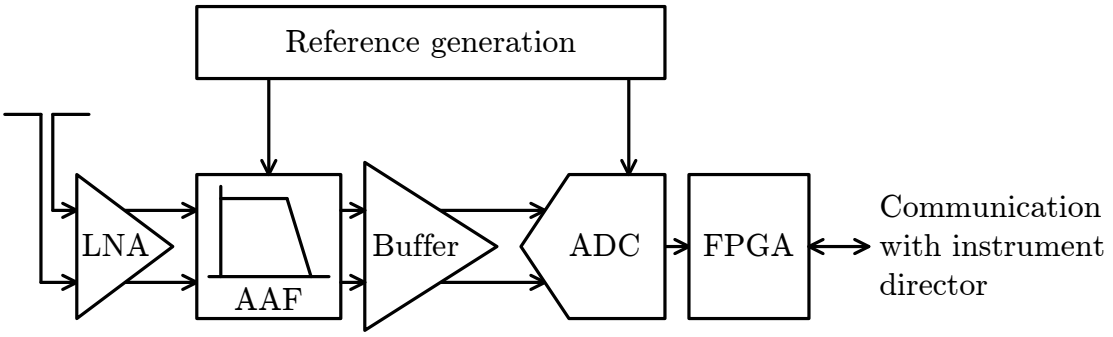

Figure G.2: Architecture of the electric field wideband receiver of the μPWI. Low-noise amplifier denoted as LNA, anti-aliasing filter denoted as AAF, analog-to-digital converter denoted as ADC, field-programmable gate array denoted as FPGA.

facilitates communication between the receiver and the instrument director. To describe the signal path blocks of Figure G.2 in more detail:

- *Low-noise amplifier (LNA)*

  The LNA is implemented by the SVLNAE-3 chip developed by Benjamin J. Mossawir [*Mossawir et al.*, 2006]. This LNA is a fully differential, radiation-hard (rated to 1 Mrad(Si)) amplifier designed to interface with electric dipole antennas.[1] The key specifications of the SVLNAE-3 LNA are given in Figure 1.11.[2]

- *Anti-aliasing filter (AAF)*

  The AAF is also implemented within the SVLNAE-3 chip [*Mossawir*, in preparation]. This AAF is a fully differential, radiation-hard (rated to at least 100 krad(Si)[3]) low-pass filter with a cutoff of 1080 MHz.[4] The key specifications

[1] In particular, the SVLNAE-3 LNA has a high input impedance and externally adjustable gain. The former enables the LNA to capture the electric antenna signal in a voltmeter-like fashion, reducing signal loss. The latter enables the LNA to accommodate a wide variety of antenna lengths, important since the received voltage V depends on the antenna length L_{eff} (specifically, $V = E(L_{\text{eff}})$ where E is the electric field signal captured by the antenna).

[2] In addition to the SVLNAE-3 chip, the LNA block of Figure G.2 also includes protection circuitry between the antenna and LNA. However, since this protective circuitry is primarily composed of passive or inactive devices (such as protection diodes), its power contribution is negligible compared to the SVLNAE-3.

[3] Final total-dose rating pending.

[4] While the SVLNAE-3 AAF cutoff frequency can be set to smaller bandwidths, for this study only the full-bandwidth 1080-MHz cutoff frequency is considered.

of the AAF are given in Figure 1.11.

- *Buffer*
  To properly drive the switched-capacitor input of the ADC, a buffer is included between the AAF and ADC. For purposes of power estimation, it is assumed that the buffer is built around a National Semiconductor LM6172 dual operational amplifier: not only do the two operational amplifiers allow buffering of the fully differential signal path, but the LM6172 is also available in a radiation-hard version rated to 300 krad(Si) [*National Semiconductor Corporation*, MNLM6172AM-X-RH].

- *Analog-to-digital converter (ADC)*
  The ADC is implemented by the SVADC-1 chip described in this dissertation. This ADC is a fully differential, radiation-hard (rated to 1 Mrad(Si)) pipeline ADC clocked at 5 MS/s. The key specifications of the ADC are given in Figure 1.11. The SVADC-1 chip includes the analog portion of the SVADC-1 converter, as described in Chapter 5 and shown in Figure 5.40.

- *Field-programmable gate array (FPGA)*
  The FPGA is implemented by an Actel RTAX250S [*Actel Corporation*, RTAX-S].[5] This FPGA is radiation-hard (rated to 200 krad(Si)) and clocked at 5 MHz. As mentioned, it includes the receiver control logic and any additional signal processing and data formatting as needed. It also includes the digital portion of the SVADC-1 converter, as detailed in Section 6.4.1.

In addition to the signal path, the receiver of Figure G.2 also includes reference generation:

- *Reference generation*
  The SVLNAE-3 and SVADC-1 both require two voltage references apiece. To generate and drive these references, it is assumed that one LM6172 is used for each.

[5] A more complete description of the Actel RTAX250S can be found in footnote 29 of Chapter 6.

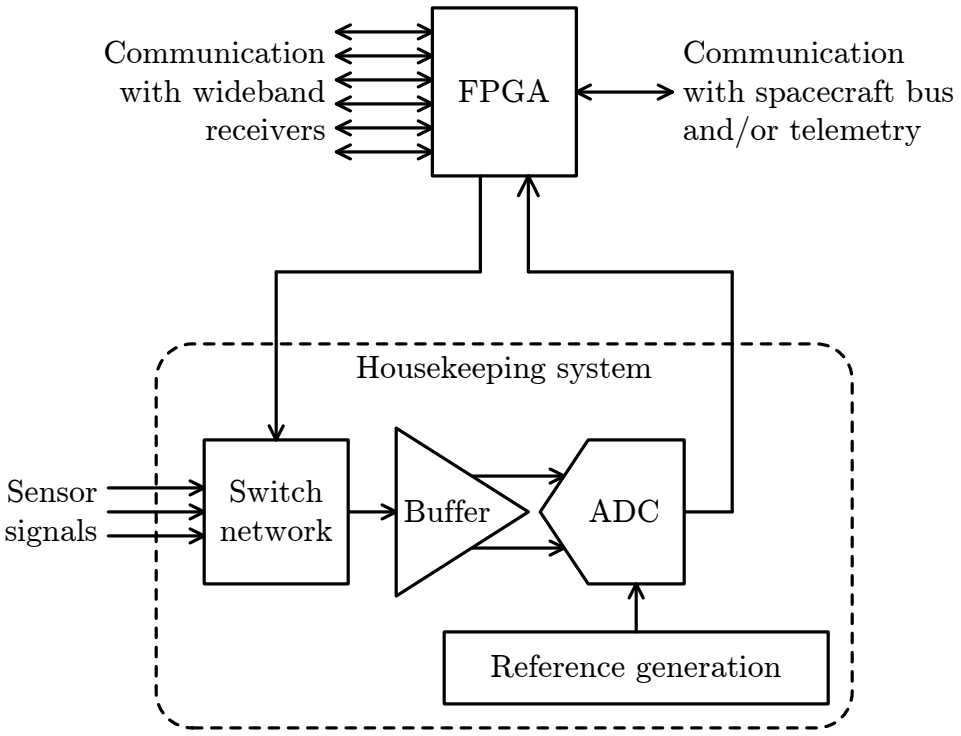

Figure G.3: Architecture of the instrument director and housekeeping system of the μPWI.

The architecture of a magnetic field wideband receiver is similar to that of the electric field wideband receiver of Figure G.2, differing only in its LNA design. However, it is expected that the power consumption of the electric and magnetic field LNAs is similar. Thus, for these power estimates, the architecture of Figure G.2 is assumed for both electric and magnetic receivers.

G.1.2 Instrument Director and Housekeeping System

The architecture of the instrument director and housekeeping system is shown in Figure G.3.

The heart of the instrument director is the FPGA, which directs the entire μPWI and handles the μPWI-to-spacecraft-bus interface. The former task includes managing the data collection of each of the wideband receivers—possibly performing additional signal processing of the wideband receiver data—and handling the

housekeeping system. The latter task includes interpretation of spacecraft bus commands and any formatting of the outgoing data—both from the wideband receiver and from housekeeping—as needed. The FPGA is implemented as:

- *Field-programmable gate array (FPGA)*
  Similar to the electric field wideband receiver, the FPGA is implemented by a radiation-hard Actel RTAX250S clocked at 5 MHz. In addition to the aforementioned tasks, the FPGA also includes the digital portion of the SVADC-1 converter (since the SVADC-1 implements the housekeeping ADC, see below).

The housekeeping system is shown in detail in Figure G.3. While the housekeeping system handles a large number of sensor signals, the report rate of these signals need not be high. Hence it is assumed that the housekeeping system employs only a single ADC: a large switch network routes the desired sensor signal to this ADC for digitization. The housekeeping system is implemented as:

- *Switch network*
  The switch network itself can be implemented in a low power fashion, such as by analog multiplexers or latching relays. The power consumption of this portion of the switching network is hence not included in this estimate. However, depending on the size of the instrument, the sensor signals may also need to be buffered throughout the switch network. For these estimates, it is assumed that this buffering is implemented by four LM6172s.

- *Buffer*
  To provide proper drive, a buffer precedes the ADC. Again, it is assumed this buffer is implemented by an LM6172.

- *Analog-to-digital converter (ADC)*
  The ADC is implemented by the SVADC-1 chip described in this dissertation. For this estimate, it is assumed that the ADC is continuously clocked at 5 MS/s, although the ADC can be clocked at lower conversion rates and for short bursts for lower power consumption.

- *Reference generation*
  The SVADC-1 requires two voltage references. It is assumed that an LM6172 is used in generating and driving these references.

G.2 Power Estimate

When estimating the power of plasma wave instruments, it is important to take into account the power supplies provided by the spacecraft bus. These power supplies are often of much larger voltage than required by the instrument electronics, necessitating voltage down-conversions within the instrument. These down-conversions often consume a non-negligible portion of the instrument power budget.

G.2.1 FAST Fields Instrument

The power consumption of the FAST fields instrument is summarized in Table G.1. For purposes of comparison with the μPWI, only the power consumptions of the low-frequency analog signal conditioning (LF analog), high-frequency analog signal conditioning (HF analog), swept frequency analyzer (SFA), and ADCs (A/D) are considered. These portions total to 4.14 W at component, that is, assuming the local, lower supplies. Taking into account power regulation, and assuming an overall efficiency of 75%,[6] the total power consumption at the spacecraft bus is 5.52 W.

G.2.2 μPWI Instrument

For the μPWI, it is assumed that the spacecraft bus provides power at $\pm$28 V, a commonly used value. Supply down-conversion is then performed by a combination of switching and linear regulators. Switching regulators provide down-conversion at respectable efficiencies (between 65% and 90%) but can be noisy as switching transients propagate into the down-converted voltage. Thus switching regulators do

[6]The 75% efficiency is derived by noting that the total regulated power of the signal processing component of the FAST fields instrument is 7.82 W, whereas the total unregulated power is 10.43 W: (7.82 W/10.43 W) 100% = 75%.

Component	Sub Component	Number	Mass	Power	Duty cycle	Orbit averaged power
	Radial boom/sensors	4	3.1 kg			
	Radial BEB	4	0.5 kg	0.5 W		
Electric Field	Axial boom/sensors	2	1.6 kg			
BEB/sensors	Axial BEB	1	0.5 kg	0.5 W		
	Totals (regulated power)			2.5 W		
	Total (unregulated power)		18.1 kg	2.8 W	0.25	0.70 W
	Senor		0.87 kg	0.45 W		
Search coil	Boom		1.87 kg			
	Totals (unregulated power)		2.65 kg	0.60 W	0.25	0.15 W
	Sensor		0.63 kg			
Fluxgate	Boom		1.78 kg			
(not including	Dirver electronics		0.90 kg	1.51 W		
ACS sensor.)	Totals (unregulated power)		3.31 kg	2.01 W	1.00	2.01 W
Signal	LF analog and A/D converters		1.76 kg	2.22 W		
processing	HF analog and SFA		0.88 kg	1.92 W	1.00 x	
	Digital signal proc.		0.40 kg	1.28 W	0.35 W	
	High speed burst memory		0.55 kg	0.65 W	and	
	Wave-particle correlator		0.80 kg	1.75 W	0.25 x	
	Totals (regulated power)			7.82 W	10.1 W	
	Totals (unregulated power)		4.39 kg	10.43 W		2.88 W
Totals			28.45 kg	15.9 W		5.74 W

Table G.1: Mass and power consumption of the FAST fields instrument. Reproduced in total from *Ergun et al.* [2001].

not directly supply sensitive, high-fidelity analog electronics in the μPWI. Instead, for these circuits the switching regulator output is additionally down-converted by a linear regulator before feeding the desired components. Unfortunately, linear regulators are not particularly efficient. While this efficiency can be improved by decreasing the separation between the linear regulator's input and output voltage, this separation is often restricted to at least a few hundreds of milliVolts to prevent dropout.

For the μPWI, it is assumed that both switching and linear regulation is implemented by radiation-hard regulators rated to at least 100 krad(Si), such as those from International Rectifier and ST Microelectronics, respectively. For switching regulators, an efficiency of 65% is assumed:[7] while the typical efficiency for many switching regulators is higher, a conservative value is assumed to account for effects such as non-optimal loading. For linear regulators, a dropout voltage of at least 2 V is assumed: again, while most linear regulators have lower dropout voltages, a conservative value is assumed.[8]

The power consumption estimate for a single wideband receiver of the μPWI is summarized in Table G.2, and the same for the instrument director and housekeeping system is summarized in Table G.3. In general, typical power consumption is used for all components. To capture down-conversion inefficiency, the tables present power at the local, lower supplies (listed as "power" under "at component") and then input-refer all power consumptions to a $\pm$28 V spacecraft bus supply through both linear and switching regulators (listed as "input power" under "linear regulation", and "input power" under "switching regulation", respectively). As can be seen, the referral through the regulators is significant, more than doubling the power consumption.

From the results of Tables G.2 and G.3, the total power consumption estimate of the μPWI at component (that is, at the local, lower supplies) is:

$$\text{At local supplies: } P_{\mu\text{PWI}} = 6P_{\text{receiver}} + P_{\text{director \& housekeeping}} = 2.35 \text{ W} \qquad \text{(G.1)}$$

[7] For switching regulators, "efficiency" refers to input-power to output-power efficiency.

[8] For linear regulators, it is assumed that the regulator provides the current required at its regulated output without consuming any additional power.

and the total power consumption estimate at the spacecraft bus ±28 V supplies is:

$$\text{At } \pm 28 \text{ V supplies: } P_{\mu\text{PWI}} = 6P_{\text{receiver}} + P_{\text{director \& housekeeping}} = 5.48 \text{ W} \qquad \text{(G.2)}$$

Recalling the corresponding powers of 4.14 W and 5.52 W for the FAST fields instrument, it is clear that the power consumption of the μPWI is comparable. Naturally, it should be kept in mind that the FAST fields instrument is a measured power consumption, whereas the μPWI is an estimate: if actually constructed, the μPWI would doubtless contain additional circuitry not included in this estimate. However, given the conservative overdesign incorporated into the μPWI estimate, it is likely that, even with the additional circuitry, the power consumption of an actual μPWI would nonetheless be competitive with that of the FAST fields instrument.

Component	At component		Linear regulation[a]		Switching regulation[b]	
	Supplies	Power	Regulation	Input power	Regulation	Input power
Signal path						
LNA (SVLNAE-3)	2.5 V	2 mW	5 V → 2.5 V	4 mW	28 V → 5 V	6.2 mW
AAF (SVLNAE-3)	2.5 V	48 mW	5 V → 2.5 V	96 mW	28 V → 5 V	147.7 mW
Buffer (LM6172)	±5 V	46 mW	±7 V → ±5 V	64.4 mW	±28 V → ±7 V	99.1 mW
ADC (SVADC-1)	2.5 V	60 mW	5 V → 2.5 V	120 mW	28 V → 5 V	184.6 mW
FPGA (Actel RTAX250S)[c]						
Core	1.5 V	68.2 mW	[No linear regulation]		28 V → 1.5 V	113.7 mW
I/O	2.5 V	7.4 mW	[No linear regulation]		28 V → 2.5 V	12.3 mW
Reference generation						
Drivers (2 × LM6172)	±5 V	92 mW	±7 V → ±5 V	128.8 mW	±28 V → ±7 V	198.2 mW
Totals		323.6 mW		488.8 mW		761.7 mW

[a] Linear regulator input power computed assuming regulator provides required current without any additional power consumption.

[b] Switching regulator input power computed assuming 65% input-to-output power efficiency for all conversions.

[c] For estimation purposes, FPGA power computed assuming 100% gate utilization, 40 input pins, and 40 output pins. Actual utilization and pin count expected to be lower. FPGA static power (which constitutes the vast majority of the FPGA power consumption at these rates) wholly attributed to FPGA core as this choice results in greater overall power consumption.

Table G.2: Power estimate for a single wideband receiver of the μPWI. All supplies single-ended (with 0-V return) unless noted otherwise. Power consumption is tabulated after referral to the local component supply, to the linear regulator input, and to the switching regulator input. Total power consumption at each level is also included.

Component	At component		Linear regulation[a]		Switching regulation[b]	
	Supplies	Power	Regulation	Input power	Regulation	Input power
FPGA (Actel RTAX250S)[c]						
Core	1.5 V	68.2 mW	[No linear regulation]		28 V → 1.5 V	113.7 mW
I/O	2.5 V	7.4 mW	[No linear regulation]		28 V → 2.5 V	12.3 mW
Housekeeping system						
Switching network (4 × LM6172)	±5 V	184 mW	±7 V → ±5 V	257.6 mW	±28 V → ±7 V	396.3 mW
Buffer (LM6172)	±5 V	46 mW	±7 V → ±5 V	64.4 mW	±28 V → ±7 V	99.1 mW
ADC (SVADC-1)	2.5 V	60 mW	5 V → 2.5 V	120 mW	28 V → 5 V	184.6 mW
Reference generation (LM6172)	±5 V	92 mW	±7 V → ±5 V	128.8 mW	±28 V → ±7 V	99.1 mW
Totals		411.6 mW		582.0 mW		905.1 mW

[a] Linear regulator input power computed assuming regulator provides required current without any additional power consumption.
[b] Switching regulator input power computed assuming 65% input-to-output power efficiency for all conversions.
[c] For estimation purposes, FPGA power computed assuming 100% gate utilization, 40 input pins, and 40 output pins. Actual utilization and pin count expected to be lower. FPGA static power (which constitutes the vast majority of the FPGA power consumption at these rates) wholly attributed to FPGA core as this choice results in greater overall power consumption.

Table G.3: Power estimate for the instrument director and housekeeping system of the μPWI. All supplies single-ended (with 0-V return) unless noted otherwise. Power consumption is tabulated after referral to the local component supply, to the linear regulator input, and to the switching regulator input. Total power consumption at each level is also included.

Appendix H

Distribution of nhSFDR due to Noise

This appendix derives the probability distribution function (PDF) of the nonharmonic spurious-free dynamic range (nhSFDR) achieved by a signal $y[n] = s[n] + x[n]$ where $s[n]$ is a single-tone sinusoid and $x[n]$ a noise signal. The spectrum is computed via an N-point FFT as defined in Equation (3.21).[1] Without loss of generality, this appendix assumes that N is even.[2] $s[n]$ is assumed to be of an on-bin frequency and amplitude A_{in}. $x[n]$ is assumed to be an independent, identically distributed (i.i.d.) Gaussian process of mean 0 and variance σ^2, that is, $x[n]$ is described by the PDF:

$$f_{x[n]}(x) = N(0, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-x^2/(2\sigma^2)} \tag{H.1}$$

where the notation $N(a, b)$ describes a Gaussian distribution of mean a and variance b. $s[n]$ and $x[n]$ are both assumed to be real-valued and independent of each other.

Section H.1 is the central proof of this appendix, and derives the statistics of the nhSFDR for the system just described. The final PDF and CDF of the nhSFDR are

[1] Note that in this case the nhSFDR is the same as the SFDR. However, in this appendix the term nhSFDR is used to maintain consistent terminology with Section 3.3.3.

[2] In the following derivations, the cases when the FFT bin index k is 0 or $N/2$ are often distinguished from the rest. Since $k \in \mathbb{Z}$, when N is even both these cases exist, whereas when N is odd, only the $k = 0$ case exists. Hence even N is the more general.

given in Equations (H.25) and (H.26), respectively. To maintain the overall flow, a number of propositions in Section H.1 are claimed without proof, including essential statements of the independence of various random variables. These statements are therefore revisited and proven in full in Section H.2.

An intermediate result of Section H.1 is the PDF of the noise power in an FFT bin, that is, the PDF of $|X[k]|^2$. Section H.3, then, completes this appendix by noting some statistical properties of $|X[k]|^2$, including its moments and an interesting geometric interpretation of its construction in the complex plane.

H.1 Derivation of nhSFDR Statistics

This section presents the derivation of the statistics of the nhSFDR. To sketch the proof, let $X[k]$ be the FFT of the noise signal $x[n]$. To begin, derive the PDF of $|X[k]|^2$, that is, the distribution of the noise power in just the k-th FFT bin. Extend this result to then consider the PDF of the noise floor created by the $|X[k]|^2$ in aggregate. As the nhSFDR is a function of a deterministic signal power compared to this probabilistic noise floor, the nhSFDR PDF can be derived from the noise floor PDF. The following three subsections essentially trace these steps.

H.1.1 PDF of $|X[k]|^2$

To derive the PDF of $|X[k]|^2$, first derive the PDF of $|\hat{X}[k]|^2$, where $\hat{X}[k]$ is the unscaled FFT of $x[n]$:

$$\hat{X}[k] = \sum_{n=0}^{N-1} x[n] e^{-i2\pi kn/N} \quad , \quad k = 0, 1, \ldots, N-1 \tag{H.2}$$

Comparing with the FFT definition of Equation (3.21):

$$X[k] = \frac{1}{N} \sum_{n=0}^{N-1} x[n] e^{-i2\pi kn/N} \quad , \quad k = 0, 1, \ldots, N-1 \tag{H.3}$$

it is clear that:

$$X[k] = \frac{1}{N}\hat{X}[k] \tag{H.4}$$

Begin by concentrating on $\hat{X}[k]$. Expand by Euler's identity [*Smith*, 2008, p. 15]:

$$\hat{X}[k] = \sum_{n=0}^{N-1} x[n]\cos\left(\frac{2\pi kn}{N}\right) + i\sum_{m=0}^{N-1} x[m]\sin\left(\frac{2\pi km}{N}\right) \tag{H.5}$$

Since $x[n] \in \mathbb{R}$:

$$|\hat{X}[k]|^2 = \left\{\sum_{n=0}^{N-1} x[n]\cos\left(\frac{2\pi kn}{N}\right)\right\}^2 + \left\{\sum_{m=0}^{N-1} x[m]\sin\left(\frac{2\pi km}{N}\right)\right\}^2 \tag{H.6}$$

$|\hat{X}[k]|^2$ is thus composed of the square of a sum of cosine-weighted independent Gaussians, and the square of a sum of sine-weighted independent Gaussians. That is, if:

$$Z_{\mathrm{c}}[k] = \sum_{n=0}^{N-1} x[n]\cos\left(\frac{2\pi kn}{N}\right) \tag{H.7}$$

$$Z_{\mathrm{s}}[k] = \sum_{n=0}^{N-1} x[n]\sin\left(\frac{2\pi kn}{N}\right) \tag{H.8}$$

then:

$$|\hat{X}[k]|^2 = Z_{\mathrm{c}}^2[k] + Z_{\mathrm{s}}^2[k] \tag{H.9}$$

Hence $|\hat{X}[k]|^2$ can be constructed from $Z_{\mathrm{c}}^2[k]$ and $Z_{\mathrm{s}}^2[k]$.

To describe the PDFs of $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$, recall that a $N(0,\sigma^2)$ distribution weighted by a is $N(0,a^2\sigma^2)$ distributed [*Leon-Garcia*, 1994, p. 123]. And recall that the distribution of a sum of independent Gaussians is also a Gaussian whose mean is the sum of the means, and whose variance is the sum of the variances [*Leon-Garcia*, 1994, p. 272]. Hence $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$ are $N(0,\sigma_c^2)$ and $N(0,\sigma_s^2)$ Gaussian distributed, respectively, where it can be shown that (see Section H.2.3.1):

$$\sigma_c^2 = \sigma^2\sum_{n=0}^{N-1}\cos^2\left(\frac{2\pi kn}{N}\right) = \sigma^2\begin{cases} N & , \quad k = 0 \text{ or } k = N/2 \\ N/2 & , \quad \text{otherwise} \end{cases} \tag{H.10}$$

$$\sigma_s^2 = \sigma^2 \sum_{n=0}^{N-1} \sin^2\left(\frac{2\pi kn}{N}\right) = \sigma^2 \begin{cases} 0 & , \quad k = 0 \text{ or } k = N/2 \\ N/2 & , \quad \text{otherwise} \end{cases} \tag{H.11}$$

Furthermore, it can be shown that $Z_c[k]$ and $Z_s[k]$ are independent (see Section H.2.1).

In combining $Z_c[k]$ and $Z_s[k]$ as per Equation (H.9), two cases arise. In the first, $k = 0$ or $k = N/2$ and $Z_c[k]$ and $Z_s[k]$ are independent Gaussians of different variances. This case is called the edge bins case. In the second, encompassing all other k, $Z_c[k]$ and $Z_s[k]$ are independent Gaussians though described by the same distribution. This case is called the inner bins case. Addressing each case in turn:

Edge bins case

When $k = 0$ or $k = N/2$, it is useful to backtrack and note that Equation (H.5) reduces to:

$$\hat{X}[0] = \sum_{n=0}^{N-1} x[n] \qquad \text{and} \qquad \hat{X}[N/2] = \sum_{n=0}^{N-1} (-1)^n x[n] \tag{H.12}$$

As a $N(0, \sigma^2)$ distribution is still $N(0, \sigma^2)$ after a gain of -1, and as $\hat{X}[0]$ and $\hat{X}[N/2]$ are the sums of independent Gaussians, both $\hat{X}[0]$ and $\hat{X}[N/2]$ are $N(0, N\sigma^2)$ distributed. Since they are real-valued, their squared magnitudes are the same as their squares, and hence are described by the PDF [*Leon-Garcia*, 1994, p. 125]:

$$f_{|\hat{X}[0]|^2}(x) = f_{|\hat{X}[N/2]|^2}(x) = \begin{cases} \dfrac{1}{\sqrt{2\pi N\sigma^2 x}} e^{-x/(2N\sigma^2)} & , \quad x > 0 \\ 0 & , \quad x < 0 \end{cases} \tag{H.13}$$

As expected, $f_{|\hat{X}[0]|^2}(x)$ and $f_{|\hat{X}[N/2]|^2}(x)$ are only one-sided.

Inner bins case

Here $Z_c[k]$ and $Z_s[k]$ are independent Gaussians with equal means 0 and equal variances $N\sigma^2/2$. In this case, it can be shown that $|\hat{X}[k]| = \sqrt{Z_c^2[k] + Z_s^2[k]}$ is

Rayleigh distributed as [*Papoulis*, 1991, pp. 139–140]:

$$f_{|\hat{X}[k]|}(x) = \begin{cases} \dfrac{2}{N\sigma^2} x e^{-x^2/(N\sigma^2)} & , \quad x \geq 0 \\ 0 & , \quad x < 0 \end{cases} \tag{H.14}$$

To find the PDF of $|\hat{X}[k]|^2$, recall that finding the PDF of the random variable Y, where Y is a general function of a random variable X as per $Y = g(X)$, requires first determining all m solutions of $x = g^{-1}(y)$. Labeling these solutions x_m, the PDF of Y is then [*Papoulis*, 1991, p. 93]:

$$f_Y(y) = \sum_m \frac{f_X(x_m)}{|g'(x_m)|} \tag{H.15}$$

In this case, $x = g^{-1}(y)$ has only one solution, $x_1 = \sqrt{y}$, since $f_{|\hat{X}[k]|^2}(x)$ is one-sided. Hence:

$$f_{|\hat{X}[k]|^2}(x) = \begin{cases} \dfrac{1}{N\sigma^2} e^{-x/(N\sigma^2)} & , \quad x > 0 \\ 0 & , \quad x < 0 \end{cases} \tag{H.16}$$

Equations (H.13) and (H.16) give the complete PDF of $|\hat{X}[k]|^2$. To derive the PDF of $|X[k]|^2$ (that is, of the the scaled FFT), note that Equation (H.4) implies:

$$|X[k]|^2 = \frac{1}{N^2} |\hat{X}[k]|^2 \tag{H.17}$$

Again applying Equation (H.15), the PDF of $|X[k]|^2$ is, for $x > 0$:

$$f_{|X[k]|^2}(x) = \begin{cases} \sqrt{\dfrac{N}{2\pi\sigma^2 x}}\, e^{-Nx/(2\sigma^2)} & , \quad k = 0 \text{ or } k = N/2 \\ \dfrac{N}{\sigma^2} e^{-Nx/\sigma^2} & , \quad \text{otherwise} \end{cases} \tag{H.18}$$

Naturally, $f_{|X[k]|^2}(x) = 0$ for $x < 0$. Equation (H.18) thus gives the PDF of the noise power in each FFT bin.[3]

[3] There is a subtlety here in that, as stated, the PDF of $|X[k]|^2$ is not defined at $x = 0$. Certainly for the edge bin case of Equation (H.18) this omission makes sense, as the function goes to ∞ at

The cumulative probability function (CDF) of $|X[k]|^2$ is the running integral of the PDF. It can be shown that the CDF is, for $x > 0$:

$$F_{|X[k]|^2}(x) = \begin{cases} \operatorname{erf}\left(\sqrt{\dfrac{Nx}{2\sigma^2}}\right) & , \quad k = 0 \text{ or } k = N/2 \\ 1 - e^{-Nx/\sigma^2} & , \quad \text{otherwise} \end{cases} \tag{H.19}$$

where $\operatorname{erf}(\cdot)$ is the error function [*Spiegel*, 1971, p. 212]:

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-u^2} \, \mathrm{d}u \tag{H.20}$$

Naturally, $F_{|X[k]|^2}(x) = 0$ for $x < 0$. Equation (H.19) thus describes the probability that the noise power $|X[k]|^2$ is less than some value x.

H.1.2 PDF of Noise Floor

Equations (H.18) and (H.19) describe the noise power of a single FFT bin $X[k]$. In concert the $X[k]$ establish the noise floor of the spectrum. Since $x[n] \in \mathbb{R}$, the noise floor is symmetric about $N/2$ [*Oppenheim et al.*, 1999, p. 576]. Furthermore, the DC bin is typically omitted when defining noise floor. Hence, in considering the noise floor, it is sufficient to consider only the $|X[k]|^2$ for $0 < k < N/2$: such $|X[k]|^2$ are called positive inner bins.

From Equation (H.18), it is clear that the positive inner bins are identically distributed. It can be shown that they are also independent (see Section H.2.2). Hence the probability that the noise floor lies below a level x is:

$$\begin{aligned} \mathrm{P}\{[\text{noise floor}] < x\} &= \mathrm{P}\{[\text{all positive inner bins}] < x\} \\ &= \mathrm{P}\{[\text{an inner bin}] < x\}^{\eta} \end{aligned} \tag{H.21}$$

where η is the number of positive inner bins considered.[4] In conjunction with Equation

$x = 0$. As the author is an engineer and not a mathematician, though, for the remainder of this appendix this subtlety is ignored.

[4]This generalization allows situations where the noise floor is defined over only a fraction of the total positive inner bins.

(H.19), the CDF of the noise floor is then:

$$F_{\mathrm{NF}}(x) = \mathrm{P}\{[\text{noise floor}] < x\} = \begin{cases} \left(1 - e^{-Nx/\sigma^2}\right)^{\eta} & , \quad x > 0 \\ 0 & , \quad x < 0 \end{cases} \tag{H.22}$$

Hence the PDF of the noise floor is:

$$f_{\mathrm{NF}}(x) = \begin{cases} \eta e^{-Nx/\sigma^2} \left(\frac{N}{\sigma^2}\right) \left(1 - e^{-Nx/\sigma^2}\right)^{\eta-1} & , \quad x > 0 \\ 0 & , \quad x < 0 \end{cases} \tag{H.23}$$

As expected, the PDF and CDF of the noise floor are both one-sided.

H.1.3 PDF of nhSFDR

The nhSFDR is defined in Equation (3.28) as the ratio of the signal power to the power of the largest noise floor bin, where the noise floor is considered only over the positive inner bins. The nhSFDR is typically given in dB. Recall that $s[n]$ is a single-tone sinusoid of on-bin frequency and amplitude A_{in}: $s[n]$ is thus confined to a single FFT bin in which it displays power $\alpha = A_{\mathrm{in}}^2/4$. Hence, if the random variable Y describes the SFDR and the random variable X describes the noise floor, then Y and X are related as:

$$Y = 10 \log_{10}\left(\frac{\alpha}{X}\right) \tag{H.24}$$

Applying Equation (H.15), and recognizing that there is only a single solution $x_1 = \alpha 10^{-y/10}$ over the domain of interest, it can be shown that the nhSFDR PDF is:

$$f_{\mathrm{nhSFDR}}(x) = \left(\frac{N}{\sigma^2}\right) \left[\frac{\eta\, \alpha \ln(10)}{10}\right] \left[\frac{e^{-(N/\sigma^2)\alpha 10^{-x/10}}}{10^{x/10}}\right] \left[1 - e^{-(N/\sigma^2)\alpha 10^{-x/10}}\right]^{\eta-1} \tag{H.25}$$

which is the desired result.

An example of $f_{\mathrm{nhSFDR}}(x)$ is plotted in Figure H.1. Note that $f_{\mathrm{nhSFDR}}(x)$ is not symmetric. Also included in Figure H.1 is a plot of $f_{\mathrm{nhSFDR}}(x)$ as measured from numerical simulations: the measured $f_{\mathrm{nhSFDR}}(x)$ represents the (scaled) histogram

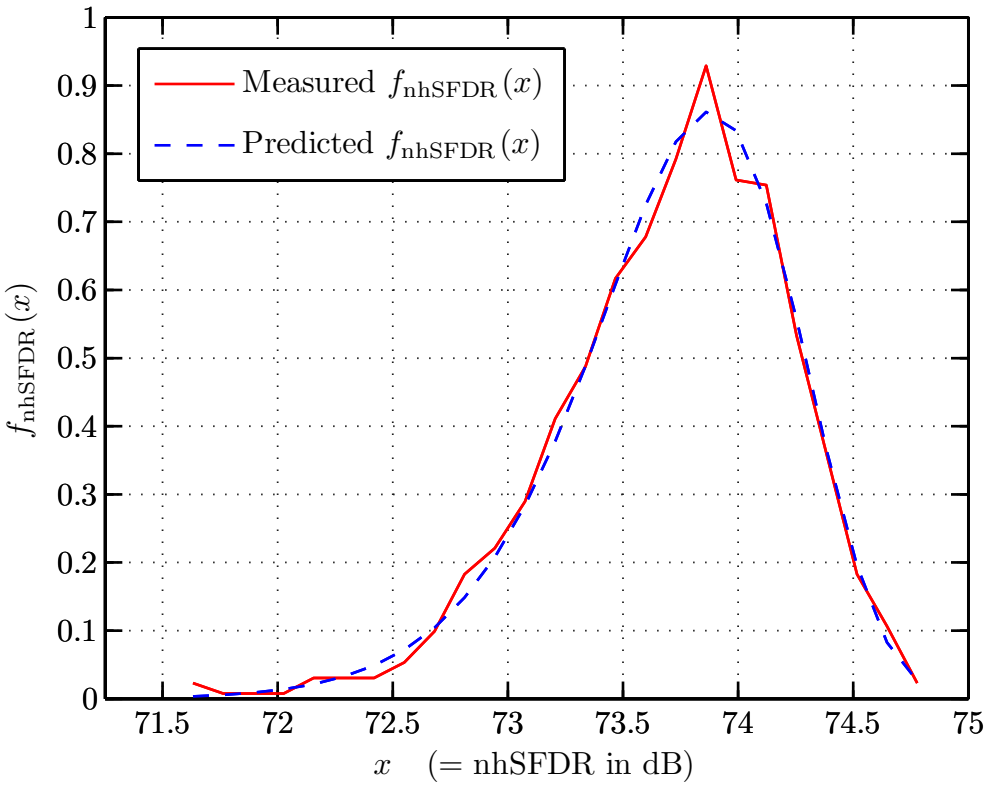

Figure H.1: Comparison of measured and predicted $f_{\text{nhSFDR}}(x)$. Measured $f_{\text{nhSFDR}}(x)$ derived from 1000 Monte Carlo runs, predicted $f_{\text{nhSFDR}}(x)$ as per Equation (H.25). Conditions: $A_{\text{in}} = 1$ and $\sigma = 7.07 \times 10^{-4}$ (corresponding to 60-dB SNR), and $N = 50,000$ and $\eta = 24,998$ (that is, η includes all the positive inner bins except the signal bin).

results of the nhSFDR from 1000 Monte Carlo runs. All in all, the numerical and theoretical results accord well.

The CDF of the nhSFDR is:

$$F_{\text{nhSFDR}}(x) = 1 - \left[1 - e^{-(N/\sigma^2)\alpha 10^{-x/10}}\right]^{\eta} \tag{H.26}$$

$F_{\text{nhSFDR}}(\text{x})$ is the probability that the nhSFDR is at most x. In practice, it is often more useful to consider the probability that the nhSFDR is at least x. This quantity

is simply the complement probability:

$$\mathrm{P}\{\mathrm{nhSFDR} > x\} = 1 - \mathrm{P}\{\mathrm{nhSFDR} \leq x\} = \left[1 - e^{-(N/\sigma^2)\alpha 10^{-x/10}}\right]^{\eta} \tag{H.27}$$

Given Equation (H.27), it is possible to derive the $p\%$-point nhSFDR_p: $p\%$ of the time, the nhSFDR is at least nhSFDR_p. It can be shown that:

$$\mathrm{nhSFDR}_p = 10\log_{10}\left(-\left(\frac{N}{\sigma^2}\right)\frac{\alpha}{\ln(1-(c/100)^{1/\eta})}\right) \tag{H.28}$$

From a design perspective, Equation (H.28) guarantees an nhSFDR performance of at least nhSFDR_p to a certainty level of $p\%$. Equation (H.28) is thus the result quoted in Equation (3.37).

H.2 Lemmas

This section addresses several claims omitted from Section H.1 for the sake of clarity, including the independence of $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$, the independence of the noise power in two different positive inner bins, and a variety of computations involving sums of sinusoids.

H.2.1 Independence of $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$

One of the key points in deriving the PDF of $|X[k]|^2$ is that $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$ are independent. This fact is especially important in the inner bins case, where it is used to show that $|X[k]|$ is Rayleigh distributed.

It can be shown that if two random variables X and Y are jointly Gaussian and uncorrelated, then they are also independent [*Papoulis*, 1991, p. 197]. Thus, to show the independence of $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$, it suffices to show that:

$Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$ are jointly Gaussian

A collection of Gaussian random variables is jointly Gaussian if and only if any linear combination of the collection is also Gaussian [*Papoulis*, 1991, p. 197].

Consider, then, the general linear combination $Z[k] = aZ_c[k] + bZ_s[k]$ for any constants $a, b \in \mathbb{R}$:

$$\begin{aligned} Z[k] &= a\sum_{n=0}^{N-1} x[n]\cos\left(\frac{2\pi kn}{N}\right) + b\sum_{m=0}^{N-1} x[m]\sin\left(\frac{2\pi km}{N}\right) \\ &= \sum_{n=0}^{N-1}\left[a\cos\left(\frac{2\pi kn}{N}\right) + b\sin\left(\frac{2\pi kn}{N}\right)\right] x[n] \end{aligned} \tag{H.29}$$

As $Z[k]$ is a weighted-sum of independent Gaussians $x[n]$, $Z[k]$ is also Gaussian. Hence $Z_c[k]$ and $Z_s[k]$ are jointly Gaussian.

$Z_c[k]$ and $Z_s[k]$ are uncorrelated

Consider $\mathrm{E}\{Z_c[k]Z_s[k]\}$, the correlation of $Z_c[k]$ and $Z_s[k]$:

$$\mathrm{E}\{Z_c[k]Z_s[k]\} = \mathrm{E}\left\{\left[\sum_{n=0}^{N-1} x[n]\cos\left(\frac{2\pi kn}{N}\right)\right]\left[\sum_{m=0}^{N-1} x[m]\sin\left(\frac{2\pi km}{N}\right)\right]\right\} \tag{H.30}$$

Recall that $\mathrm{E}\{X+Y\} = \mathrm{E}\{X\} + \mathrm{E}\{Y\}$ regardless of the interdependence of X and Y [*Leon-Garcia*, 1994, p. 232]. Since the $x[n]$ are independent:

$$\mathrm{E}\{x[n]x[m]\} = \begin{cases} 1 & , \quad n = m \\ 0 & , \quad \text{otherwise} \end{cases} \tag{H.31}$$

and the cross terms in Equation (H.30) cancel, leaving only:

$$\begin{aligned} \mathrm{E}\{Z_c[k]Z_s[k]\} &= \mathrm{E}\left\{\sum_{n=0}^{N-1} x^2[n]\cos\left(\frac{2\pi kn}{N}\right)\sin\left(\frac{2\pi kn}{N}\right)\right\} \\ &= \mathrm{E}\{x^2[n]\}\sum_{n=0}^{N-1}\cos\left(\frac{2\pi kn}{N}\right)\sin\left(\frac{2\pi kn}{N}\right) \end{aligned} \tag{H.32}$$

where the final step occurs since $\mathrm{E}\{aX\} = a\mathrm{E}\{X\}$ where a is a constant [*Leon-Garcia*, 1994, p. 132]. It can be shown that (see Section H.2.3.1):

$$\sum_{n=0}^{N-1} \cos\left(\frac{2\pi kn}{N}\right) \sin\left(\frac{2\pi kn}{N}\right) = 0 \tag{H.33}$$

and hence:

$$\mathrm{E}\{Z_\mathrm{c}[k] Z_\mathrm{s}[k]\} = 0 \tag{H.34}$$

and $Z_\mathrm{c}[k]$ and $Z_\mathrm{s}[k]$ are uncorrelated.

Being jointly Gaussian and uncorrelated, $Z_\mathrm{c}[k]$ and $Z_\mathrm{s}[k]$ are independent.

H.2.2 Independence of $|X[k]|^2$ and $|X[l]|^2$

The independence of two different positive inner bins $|X[k]|^2$ and $|X[l]|^2$ is important as it renders the CDF of the noise floor separable into the CDFs of the individual bins. This section proves the independence of $|X[k]|^2$ and $|X[l]|^2$, for $0 < k < N/2$, $0 < l < N/2$, and $k \neq l$.[5]

To do so, decompose $|X[k]|^2$ and $|X[l]|^2$ as per Equation (H.9) into the random variables $Z_\mathrm{c}[k]$ and $Z_\mathrm{s}[k]$ (which construct $|X[k]|^2$), and $Z_\mathrm{c}[l]$ and $Z_\mathrm{s}[l]$ (which construct $|X[l]|^2$). It then suffices to show that these four random variables—namely, $Z_\mathrm{c}[k]$, $Z_\mathrm{s}[k]$, $Z_\mathrm{c}[l]$ and $Z_\mathrm{s}[l]$—are independent of each other.

The independence of $Z_\mathrm{c}[k]$ and $Z_\mathrm{s}[k]$, and of $Z_\mathrm{c}[l]$ and $Z_\mathrm{s}[l]$, is established in Section H.2.1. Hence only the independence of the cross-pairs remains to be shown, that is, the independence between $Z_\mathrm{c}[k]$ and $Z_\mathrm{c}[l]$, between $Z_\mathrm{c}[k]$ and $Z_\mathrm{s}[l]$, between $Z_\mathrm{s}[k]$ and $Z_\mathrm{c}[l]$, and between $Z_\mathrm{s}[k]$ and $Z_\mathrm{s}[l]$. Similar to Section H.2.1, the proof is done by showing that the cross-pairs are jointly Gaussian and uncorrelated.

The cross-pairs are jointly Gaussian

To show a cross-pair jointly Gaussian, it suffices to show that any linear

[5] The bounds on k and l are important here. For example, since $x[n] \in \mathbb{R}$, the power spectrum is symmetric. In particular, each positive frequency bin is equal to a negative frequency bin: these two bins are thus decidedly not independent. However, over just the positive inner bins, pairs of bins are indeed independent.

combination of the cross-pair is Gaussian. Consider the first cross-pair:

$$aZ_{\mathrm{c}}[k] + bZ_{\mathrm{c}}[l] = \sum_{n=0}^{N-1} \left[a\cos\left(\frac{2\pi kn}{N}\right) + b\cos\left(\frac{2\pi ln}{N}\right) \right] x[n] \tag{H.35}$$

As Equation (H.35) is a weighted-sum of independent Gaussians $x[n]$, it is also Gaussian. The extension to the other cross-pairs is obvious and yields similar results. Hence the cross-pairs are jointly Gaussian.

The cross-pairs are uncorrelated

By similarity to Section H.2.1, the correlation of the cross-pairs is:

$$\mathrm{E}\{Z_{\mathrm{c}}[k]Z_{\mathrm{c}}[l]\} = \mathrm{E}\{x^2[n]\} \sum_{n=0}^{N-1} \cos\left(\frac{2\pi kn}{N}\right) \cos\left(\frac{2\pi ln}{N}\right) \tag{H.36}$$

$$\mathrm{E}\{Z_{\mathrm{c}}[k]Z_{\mathrm{s}}[l]\} = \mathrm{E}\{x^2[n]\} \sum_{n=0}^{N-1} \cos\left(\frac{2\pi kn}{N}\right) \sin\left(\frac{2\pi ln}{N}\right) \tag{H.37}$$

$$\mathrm{E}\{Z_{\mathrm{s}}[k]Z_{\mathrm{c}}[l]\} = \mathrm{E}\{x^2[n]\} \sum_{n=0}^{N-1} \sin\left(\frac{2\pi kn}{N}\right) \cos\left(\frac{2\pi ln}{N}\right) \tag{H.38}$$

$$\mathrm{E}\{Z_{\mathrm{s}}[k]Z_{\mathrm{s}}[l]\} = \mathrm{E}\{x^2[n]\} \sum_{n=0}^{N-1} \sin\left(\frac{2\pi kn}{N}\right) \sin\left(\frac{2\pi ln}{N}\right) \tag{H.39}$$

It can be shown that, under the established conditions on k and l, the summation (see Section H.2.3.2):

$$\sum_{n=0}^{N-1} \cos\left(\frac{2\pi kn}{N} + \phi\right) \cos\left(\frac{2\pi ln}{N} + \gamma\right) = 0 \tag{H.40}$$

and hence:

$$\mathrm{E}\{Z_{\mathrm{c}}[k]Z_{\mathrm{c}}[l]\} = \mathrm{E}\{Z_{\mathrm{c}}[k]Z_{\mathrm{s}}[l]\} = \mathrm{E}\{Z_{\mathrm{s}}[k]Z_{\mathrm{c}}[l]\} = \mathrm{E}\{Z_{\mathrm{s}}[k]Z_{\mathrm{s}}[l]\} = 0 \tag{H.41}$$

Thus, the cross-pairs are uncorrelated.

As they are all pairwise jointly Gaussian and uncorrelated, the quartet of random

variables $Z_{\mathrm{c}}[k]$, $Z_{\mathrm{s}}[k]$, $Z_{\mathrm{c}}[l]$, and $Z_{\mathrm{s}}[l]$ are all independent of each other. Therefore a function of $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$ is independent of a function of $Z_{\mathrm{c}}[l]$ and $Z_{\mathrm{s}}[l]$: in particular, $|X[k]|^2$ is independent of $|X[l]|^2$.

H.2.3 Sums of Sinusoids

Throughout these derivations, reference is made to various sums of sines and cosines. Closed form expressions for these sums are derived here. The task is divided into two cases: sums involving two sinusoids of the same frequency, and sums involving two sinusoids of different frequencies.

H.2.3.1 Of the Same Frequency

Such sums are invoked in computing σ_c^2 and σ_s^2 (the variances of $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$, respectively) and in determining $\mathrm{E}\{Z_{\mathrm{c}}[k]Z_{\mathrm{s}}[k]\}$ (the correlation of $Z_{\mathrm{c}}[k]$ and $Z_{\mathrm{s}}[k]$). To tackle these sums, first consider the simpler sum $S_{\cos}$:

$$S_{\cos} = \sum_{n=0}^{N-1} \cos\left(\frac{4\pi kn}{N}\right) = \frac{1}{2}\left[\sum_{n=0}^{N-1} e^{i4\pi kn/N} + \sum_{m=0}^{N-1} e^{-i4\pi km/N}\right] \tag{H.42}$$

where the latter equality is by Euler's identity. Recall that, for $z \in \mathbb{C}$, the geometric series of z has closed form expression [*Smith*, 2008, pp. 99–100]:

$$\sum_{n=0}^{N-1} z^n = \begin{cases} \dfrac{1-z^N}{1-z} & , \quad z \neq 1 \\ N & , \quad z = 1 \end{cases} \tag{H.43}$$

In the case of Equation (H.42), $z = e^{\pm i4\pi k/N}$. Note that $e^{\pm i4\pi k/N} = 1$ when $(2k/N) \in \mathbb{Z}$. Given that $k \in [0, 1, \ldots, N{-}1]$, this condition only occurs when $k = 0$ or when $k = N/2$, assuming N is even. Hence:

$$S_{\cos} = \frac{1}{2}\begin{cases} 2N & , \quad k = 0 \text{ or } k = N/2 \\ \dfrac{1-e^{i4\pi k}}{1-e^{i4\pi k/N}} + \dfrac{1-e^{-i4\pi k}}{1-e^{-i4\pi k/N}} & , \quad \text{otherwise} \end{cases} \tag{H.44}$$

Through a series of algebraic manipulations, it can be shown that:

$$S_{\cos} = \begin{cases} N & , \quad k = 0 \text{ or } k = N/2 \\ \left[\dfrac{\sin(2\pi k)}{\sin(2\pi k/N)}\right] \cos\left(\dfrac{2\pi k(N-1)}{N}\right) & , \quad \text{otherwise} \end{cases} \tag{H.45}$$

Consider the latter case of $k \neq 0$ and $k \neq N/2$. Since $k \in \mathbb{Z}$, thus $\sin(2\pi k) = 0$. Furthermore, since $k \neq 0$ and $k \neq N/2$, thus $\sin(2\pi k/N) \neq 0$. Hence, Equation (H.45) simplifies to:

$$S_{\cos} = \begin{cases} N & , \quad k = 0 \text{ or } k = N/2 \\ 0 & , \quad \text{otherwise} \end{cases} \tag{H.46}$$

By a similar derivation, it can be shown that:

$$S_{\sin} = \sum_{n=0}^{N-1} \sin\left(\frac{4\pi kn}{N}\right) = 0 \tag{H.47}$$

Consider, now, the aforementioned sums of interest. For Equation (H.10), which describes σ_c^2:

$$\begin{aligned} \sum_{n=0}^{N-1} \cos^2\left(\frac{2\pi kn}{N}\right) &= \frac{1}{2} \sum_{n=0}^{N-1} \left[1 + \cos\left(\frac{4\pi kn}{N}\right)\right] = \frac{N}{2} + \frac{1}{2} S_{\cos} \\ &= \begin{cases} N & , \quad k = 0 \text{ or } k = N/2 \\ N/2 & , \quad \text{otherwise} \end{cases} \end{aligned} \tag{H.48}$$

For Equation (H.11), which describes σ_s^2:

$$\begin{aligned} \sum_{n=0}^{N-1} \sin^2\left(\frac{2\pi kn}{N}\right) &= \frac{1}{2} \sum_{n=0}^{N-1} \left[1 - \cos\left(\frac{4\pi kn}{N}\right)\right] = \frac{N}{2} - \frac{1}{2} S_{\cos} \\ &= \begin{cases} 0 & , \quad k = 0 \text{ or } k = N/2 \\ N/2 & , \quad \text{otherwise} \end{cases} \end{aligned} \tag{H.49}$$

And finally, for Equation (H.33), which describes $\mathrm{E}\{Z_\mathrm{c}[k]Z_\mathrm{s}[k]\}$:

$$\sum_{n=0}^{N-1}\cos\left(\frac{2\pi kn}{N}\right)\sin\left(\frac{2\pi kn}{N}\right)=\frac{1}{2}\sum_{n=0}^{N-1}\sin\left(\frac{4\pi kn}{N}\right)=\frac{1}{2}S_{\sin}=0 \tag{H.50}$$

H.2.3.2 Of Different Frequencies

This case is invoked in considering the correlation of the cross-pairs of $Z_\mathrm{c}[k]$, $Z_\mathrm{s}[k]$, $Z_\mathrm{c}[l]$, and $Z_\mathrm{s}[l]$. In particular, the generalized sum:

$$S_{\cos,\cos}=\sum_{n=0}^{N-1}\cos\left(\frac{2\pi kn}{N}+\phi\right)\cos\left(\frac{2\pi ln}{N}+\gamma\right) \tag{H.51}$$

is considered, under the conditions that $0<k<N/2$ and $0<l<N/2$, and $k\neq l$. Expanding by Euler's identity, invoking the closed form of a geometric series, and applying some algebraic elbow-grease, it can be shown that:

$$\begin{aligned}S_{\cos,\cos}=&\frac{1}{4}e^{i(\phi+\gamma)+i\pi(k+l)(1-1/N)}\left\{\frac{\sin(\pi(k+l))}{\sin(\pi(k+l)/N)}\right\}\\&+\frac{1}{4}e^{i(\phi-\gamma)+i\pi(k-l)(1-1/N)}\left\{\frac{\sin(\pi(k-l))}{\sin(\pi(k-l)/N)}\right\}\\&+\frac{1}{4}e^{i[-(\phi+\gamma)]+i\pi[-(k+l)](1-1/N)}\left\{\frac{\sin(\pi[-(k+l)])}{\sin(\pi[-(k+l)]/N)}\right\}\\&+\frac{1}{4}e^{i[-(\phi-\gamma)]+i\pi[-(k-l)](1-1/N)}\left\{\frac{\sin(\pi[-(k-l)])}{\sin(\pi[-(k-l)]/N)}\right\}\end{aligned} \tag{H.52}$$

Each term of Equation (H.52) contains a fraction of sines of the form:

$$R_{\sin}=\frac{\sin(\pi a)}{\sin(\pi a/N)} \tag{H.53}$$

where $a\in\mathbb{Z}$. The numerator of $R_{\sin}$ is always 0. The denominator of $R_{\sin}$ is zero only when $a\in\{\ldots,-2N,-N,0,N,2N,\ldots\}$. However, to within a sign, a is composed of the sum and the difference of k and l. Since $0<k<N/2$ and $0<l<N/2$, and since $k\neq l$, both the sum and the difference are bounded. Specifically, it can be seen that $(k+l)\in[3,N-3]$ and $(k-l)\in[1,(N/2)-2]$ (assuming, without loss of generality,

that $k > l$). Hence $a \notin \{\ldots, -2N, -N, 0, N, 2N, \ldots\}$ and the denominator of $R_{\sin}$ is nonzero. Thus, given the restrictions on k and l:

$$R_{\sin} = 0 \tag{H.54}$$

which implies:

$$S_{\cos,\cos} = 0 \tag{H.55}$$

Equation (H.55) is thus the result quoted in Equation (H.40). Note that this result is anticipated by the orthogonality of the FFT bins.

H.3 Properties of $|X[k]|^2$

This appendix concludes with some interesting properties of $|X[k]|^2$, the noise power in an FFT bin, including the moments and an interesting geometric interpretation.

H.3.1 Moments

The m-th moment of $|X[k]|^2$ is given by [*Widrow and Kollár*, 2008, pp. 33–35]:

$$\mathrm{E}\left\{\left(|X[k]|^2\right)^m\right\} = \int_{-\infty}^{\infty} x^m f_{|X[k]|^2}(x)\,\mathrm{d}x \tag{H.56}$$

It can be shown that:

$$\mathrm{E}\left\{\left(|X[k]|^2\right)^m\right\} = \left(\frac{\sigma^2}{N}\right)^m \begin{cases} \dfrac{2^m}{\sqrt{\pi}}\,\Gamma\left(\dfrac{1}{2}+m\right) & , \quad k = 0 \text{ or } k = N/2 \\ \Gamma(1+m) & , \quad \text{otherwise} \end{cases} \tag{H.57}$$

where $\Gamma(\cdot)$ is the Gamma function [*Spiegel*, 1971, p. 210]:

$$\Gamma(m) = \int_0^{\infty} x^{m-1} e^{-x}\,\mathrm{d}x \tag{H.58}$$

Utilizing some properties of the Gamma function,[6] it can be shown that:

$$\mathrm{E}\left\{\left(|X[k]|^2\right)^m\right\} = \left(\frac{\sigma^2}{N}\right)^m \begin{cases} \left(\frac{1}{2^m}\right)\frac{(2m)!}{m!} & , \quad k=0 \text{ or } k=N/2 \\ m! & , \quad \text{otherwise} \end{cases} \tag{H.59}$$

An interesting result of Equation (H.59) is that the first moment of $|X[k]|^2$ is always σ^2/N, regardless of k. This fact is quoted in Equation (3.34). Other moments, though, are k-dependent, in particular, differing between the edge and inner bins.

H.3.2 A Geometric View

As a point of interest, since addition of complex numbers can be visualized as vector addition in the complex plane, $X[k]$ can be interpreted geometrically as the result of a two-dimensional random walk. In this walk, the direction of each step is predetermined (being given by the $e^{-i2\pi kn/N}$) but the amplitude is random (being given by the $x[n]$). $|X[k]|$ is then the distance from the origin at the end of N steps. From Equation (H.14), and applying Equation (H.15) to the distributions of $\hat{X}[0]$ and $\hat{X}[N/2]$, this distance is distributed as, for $x>0$:

$$f_{|\hat{X}[k]|}(x) = \begin{cases} \sqrt{\frac{2}{\pi N\sigma^2}}\, e^{-x^2/(2N\sigma^2)} & , \quad k=0 \text{ or } k=N/2 \\ \frac{2}{N\sigma^2} x e^{-x^2/(N\sigma^2)} & , \quad \text{otherwise} \end{cases} \tag{H.60}$$

Of course, $f_{|\hat{X}[k]|}(x)=0$ for $x<0$.

The geometric view also anticipates that the $k=0$ and $k=N/2$ bins differ from the rest: for these bins, the predetermined directions are such that the walk reduces to just a one-dimensional random walk along the real axis.

[6] In general, for $a>0$, the Gamma function is recursive where $\Gamma(a+1)=a\Gamma(a)$ [*Spiegel*, 1971, p. 213]. For the case of $a\in\mathbb{N}$, it can further be shown that $\Gamma(a+1)=a!$ [*Spiegel*, 1971, p. 213]. Finally, for the noninteger cases of Equation (H.57), the recursive property is combined with the fact that $\Gamma(1/2)=\sqrt{\pi}$ [*Spiegel*, 1971, p. 214].

Appendix I

$\Sigma\Delta$-Modulators

This appendix presents a basic introduction to $\Sigma\Delta$-modulators, also known as $\Delta\Sigma$-modulators.[1] These oversampling converters improve the signal-to-noise ratio (SNR) of a coarse quantizer by embedding it in a feedback loop. Classically, the loop is designed to high-pass filter the quantization noise yet allow the input signal to pass through largely unchanged. Thus over the low frequencies where the quantization noise is greatly attenuated, the SNR can thus be quite high.

The objective of this appendix is to convey an intuitive grasp of $\Sigma\Delta$-modulators. Readers desirous of a more rigorous understanding are recommended to the "Yellow Book" [*Norsworthy et al.*, 1997], which provides a very good introduction to the entirety of modulator design from architecture to circuits to decimation. Indeed, the Yellow Book is referenced heavily throughout this appendix.

This appendix begins with the noise-shaping feedback coder—a more intuitive noise-shaping coder—then expands this architecture to form the $\Sigma\Delta$-modulator. After a discussion of the general $\Sigma\Delta$-modulator, this appendix moves to some practical concerns in $\Sigma\Delta$-modulator design and addresses how the SNR of the modulator can be expanded by various architectural choices.

[1] The etymology of both these terms is discussed in footnote 8 of this appendix.

I.1 Z-Transforms

The discussions of this appendix concentrate on the quantizing aspect of ΣΔ-modulators and assume sampled-time throughout.[2] The terms "analog-to-digital converter" (ADC) and "digital-to-analog converter" (DAC) thus refer to conversions of discrete-time signals between continuous-value and discrete-value representations.

To analyze discrete-time systems, this appendix uses the z-transform. The z-transform of the discrete-time sequence $x[n]$ is defined as [*Oppenheim et al.*, 1999, p. 95]:

$$X(z) = \sum_{n=-\infty}^{\infty} x[n]z^{-n} \tag{I.1}$$

where $z \in \mathbb{C}$. Note that the discrete-time Fourier transform (DTFT) of $x[n]$ is given by [*Oppenheim et al.*, 1997, p. 361]:

$$X(\tilde{f}) = \sum_{n=-\infty}^{\infty} x[n]e^{-i2\pi n\tilde{f}} \tag{I.2}$$

where $\tilde{f} \in \mathbb{R}$ is the discrete-time frequency. Hence the z-transform can be seen as a generalization of the DTFT.[3] From Equations (I.1) and (I.2), the relationship between

[2] Many modulators are implemented using switched-capacitor circuits for the internal filters, justifying the use of a sampled-time analysis. However, an alternative is to use continuous-time circuits, such as active-RC filters, for the internal filters. Such "continuous-time modulators" can provide many benefits, not the least of which is that the operational amplifiers of such modulators often require less bandwidth to process signals of the same input bandwidth: switched-capacitor amplifiers often requiring many time constants for proper settling. These continuous-time modulators can thus often achieve higher OSRs than switched-capacitor modulators in the same technology, or alternately maintain the same OSR but at lower power [*Kulchycki et al.*, 2008]. However, the design of continuous-time modulators often involves transforming the continuous-time system into an equivalent discrete-time system to better facilitate computer simulation [*Kulchycki*, 2007, Chap. 4]. Hence the study of sampled-time modulators is often applicable to continuous-time modulators as well.

[3] Although it is often more popular to envision the relationship in reverse: the DTFT is the z-transform evaluated along the unit circle in the complex plane.

the z-transform variable z and the DTFT variable $\tilde{f}$ is simply:[4]

$$z = e^{i2\pi\tilde{f}} \tag{I.3}$$

In considering $\Sigma\Delta$-modulators, it is often useful to relate z to the continuous-time Fourier transform variable f; for example, this transformation is used in the IDR derivations of Section I.3.2. Recalling that $\tilde{f} = f/f_S$, where f_S is the sampling frequency (see footnote 18 of Chapter 3), the relation of z to f is:

$$z = e^{i2\pi f/f_S} \tag{I.4}$$

I.2 Theory

To introduce the theory of $\Sigma\Delta$-modulators, this appendix opts to start with the noise-shaping feedback coder [*Cutler*, 1954] then extend this architecture to form a $\Sigma\Delta$-modulator.[5]

I.2.1 Noise-Shaping Feedback Coder

The noise-shaping feedback coder is shown in Figure I.1(a). In this architecture, the sub-ADC quantizer error is explicitly computed and the result, after filtering by $H(z)$, subtracted from the input $X(z)$ via negative feedback. Note that $H(z)$ is assumed strictly causal to force a net delay around the loop.

Essentially, the feedback loop of Figure I.1(a) attempts to predict and then cancel the quantizer error. As an example, consider the case when $H(z)$ is a simple delay $H(z) = z^{-1}$. Then the feedback loop predicts that the quantizer error at time $n+1$

[4]The relationship between the z-transform and the DTFT in discrete-time systems analysis is thus similar to that between the Laplace transform and the continuous-time Fourier transform in continuous-time systems analysis. Indeed, much like Laplace, the z-transform is often used to conduct stability analysis.

[5]Two origins are oft traced for the $\Sigma\Delta$-modulator: the noise-shaping feedback coder ([*Cutler*, 1954]) and the Δ-modulator ([*de Jager*, 1952]). As the former path is perhaps more intuitive, it is chosen here. Readers nonetheless interested in the latter path (which is perhaps the more historical) are referred to the original derivation by *Inose et al.* [1962], or to footnote 8 of this appendix which briefly sketches said derivation.

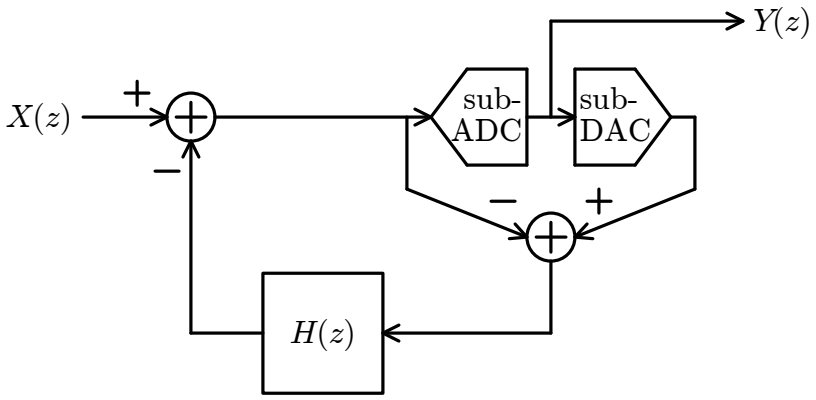

(a) Noise-shaping feedback coder.

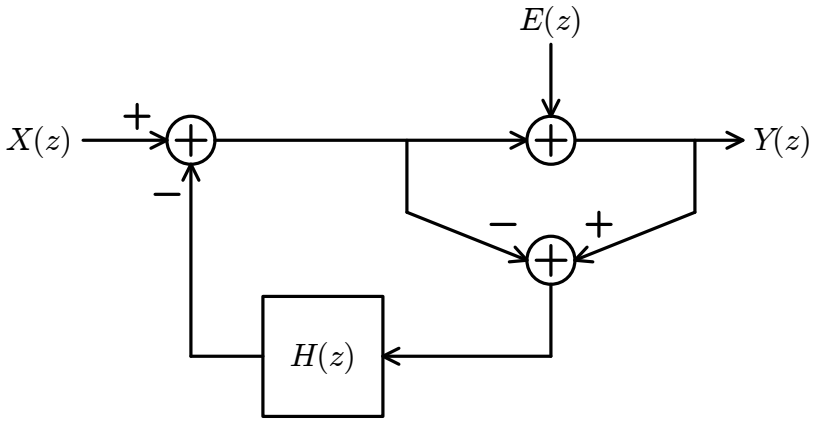

(b) Noise-shaping feedback coder under PQN model.

Figure I.1: Noise-shaping feedback coder. (a) With sub-analog-to-digital converter (sub-ADC) and sub-digital-to-analog converter (sub-DAC) shown explicitly. (b) Assuming the PQN model for the sub-ADC.

is the same as it was at time n. Intuitively, this prediction holds for slowly varying $X(z)$; indeed, it is expected that the lower the frequencies of $X(z)$, the better the prediction and, consequently, the better the cancellation. From a design perspective, then, to preserve the cancellation over a given input signal bandwidth the loop should be run at a frequency much greater than said input signal bandwidth.

This intuition can be formalized by invoking the PQN model wherein the sub-ADC is modeled as an input-independent, additive white noise $E(z)$ (see Section 3.1.2) as shown in Figure I.1(b). Solving this system for $Y(z)$:

$$Y(z) = X(z) - H(z)E(z) + E(z) = X(z) + E(z)\left[1 - H(z)\right] \tag{I.5}$$

Hence the feedback loop filters the quantizer error $E(z)$ by $1 - H(z)$, whereas the input $X(z)$ passes through unchanged. This characteristic lends the coder its name: the "noise" $E(z)$ is said to be "shaped" by the feedback loop. If $H(z)$ is chosen to be a simple delay z^{-1} (as is classically done), then the shaping $1 - H(z)$ represents a differentiation of $E(z)$ that attenuates $E(z)$ at low frequencies. If $X(z)$ is consequently confined to these low frequencies of $E(z)$ attenuation, then the resulting SNR over the bandwidth of $X(z)$ can be quite high. To guarantee this relationship between $X(z)$ and $E(z)$ the system is oversampled.

Naturally, other $H(z)$ can be chosen, but the key concept is that $E(z)$ is ultimately attenuated over the bandwidth of $X(z)$ through a combination of oversampling and noise-shaping.

I.2.2 Extension to ΣΔ-Modulators

As feedback systems are sensitive to variations in the feedback path [*Lee*, 1998, pp. 391–394], in practice it is desirable to simplify the feedback path of the noise-shaping feedback coder as much as possible. Figure I.2 demonstrates such a simplification. First, the filter $H(z)$ is pushed through the negative feedback junction to generate Figure I.2(b). Then one global feedback path is reduced to a local feedback path

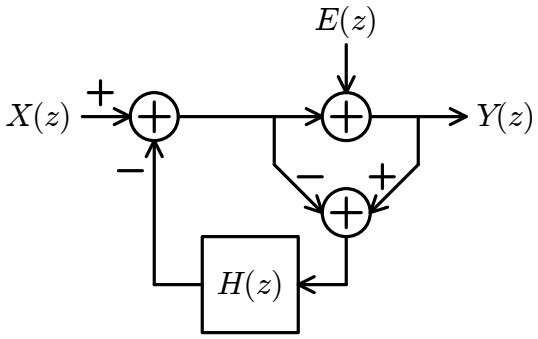

(a) Noise-shaping feedback coder.

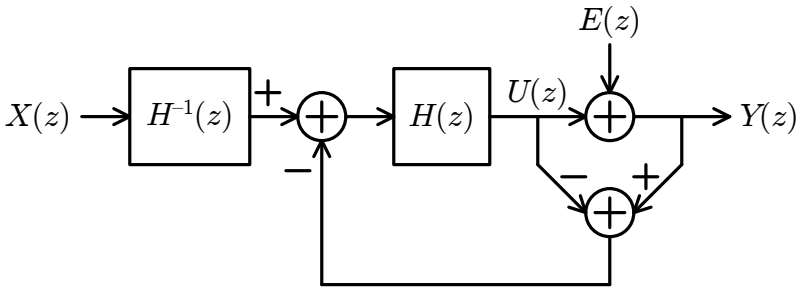

(b) After reduction of feedback filter to feedforward filters.

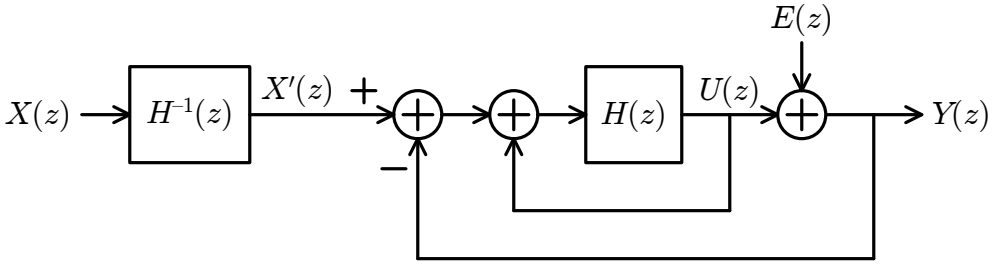

(c) After reduction of one global feedback path to a local feedback path.

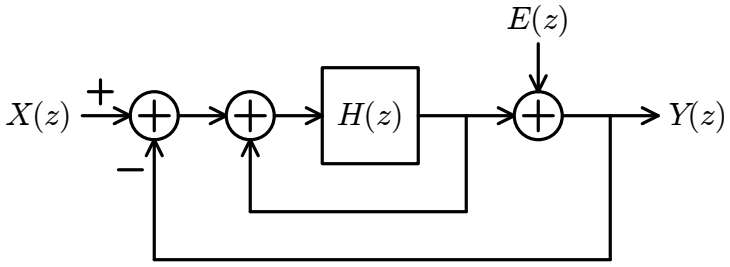

(d) After removal of input prefilter: ΣΔ-modulator.

Figure I.2: Transformation of a noise-shaping feedback coder into a ΣΔ-modulator.

as per Figure I.2(c).[6] While the system of Figure I.2(c) indeed reduces all feedback to simple unity-gain paths, it requires prefiltering of $X(z)$ by $H^{-1}(z)$. Typically, this prefiltering is simply ignored and $X'(z)$ instead assumed to be the input.[7] The resulting system is shown in Figure I.2(d) and has output:

$$Y(z) = X(z)H(z) + E(z)\left[1 - H(z)\right] \tag{I.6}$$

Comparing Equation (I.6) to Equation (I.5), in this new architecture the filtering of $E(z)$ is unchanged although the input $X(z)$ is now also filtered. While this fact does constrain the design of $H(z)$, often it is not limiting: for example, if $H(z) = z^{-1}$ then the input filtering amounts to only a delay. For historical reasons, the architecture of Figure I.2(d) is called a $\Sigma\Delta$-modulator.[8]

[6] This operation is simply the reorganization:

$$U(z) = H(z)\left[H^{-1}(z)X(z) - E(z)\right] = H(z)\left[H^{-1}(z)X(z) - Y(z) + U(z)\right]$$

Figure I.2(b) represents the first equality, Figure I.2(c) the second.

[7] Indeed, if $H(z) = z^{-1}$ as in the classical noise-shaping feedback coder, then the prefiltering is anticausal and hence not even realizable.

[8] The term $\Sigma\Delta$-modulator arises from a historical derivation of the $\Sigma\Delta$-modulator from a Δ-modulator. The Δ-modulator is a differential pulse code modulator (DPCM) [*de Jager*, 1952]. DPCMs are similar to noise-shaping feedback coders in that they rely on prediction via feedback. But whereas a noise-shaping feedback coder attempts to predict and then cancel the quantizer error, DPCMs instead attempt to predict the input signal itself. This prediction is formed directly from the past outputs of the DPCM. The DPCM then quantizes the difference between the input signal and this prediction. Conceptually, if the input is slowly varying, then it is well correlated with itself sample to sample and accurate predictions can be made. Thus, if the encoder and decoder agree on the prediction technique, then an efficient encoding would be to simply transmit the error of this prediction: such a strategy exploits the redundancy of the input signal to yield better encoding efficiency. Naturally, ever more accurate predictions yield ever better efficiencies. The Δ-modulator, then, is a DPCM with a particularly simple predictor: traditionally, just an integration of the Δ-modulator output [*de Jager*, 1952]. The Δ-modulator is thus appealing for its simplicity. However, it has many drawbacks, one of which is that (it can be shown) the Δ-modulator ultimately quantizes the derivative of its input signal instead of the input signal itself [*Inose et al.*, 1962]. However, *Inose et al.* [1962] proposed overcoming many of these deficiencies by prefacing the Δ-modulator with an integrator. The resulting system quantizes the input signal directly and, after some algebraic manipulations, yields the encoder architecture of Figure I.2(d). The result can thus be thought of as an integration (Σ) followed by a Δ-modulator, yielding the name $\Sigma\Delta$-modulator. Oddly enough, [*Inose et al.*, 1962] themselves actually called the result a $\Delta\Sigma$-modulator, since the architecture consists of a subtraction (Δ) followed by an integration (Σ) (the latter arising when $H(z) = z^{-1}$ in Figure I.2(d), for example). In modern parlance both terms are used interchangeably [*Gray*, 1997, p. 60]. This dissertation adopts the term $\Sigma\Delta$-modulator; certainly this name is the preference at

I.2.3 General ΣΔ-Modulators

A generalized ΣΔ-modulator is shown in Figure I.3(a).[9] Here the feedback path is given by $F(z)$ although, for reasons previously discussed, typically $F(z)$ is made as simple as possible. Assuming the PQN model for the sub-ADC as shown in Figure I.3(b), the ΣΔ-modulator output is:

$$Y(z) = E(z) + [X(z) - Y(z)F(z)]\,A(z) \tag{I.7}$$

Rearranging:

$$Y(z) = \underbrace{\frac{1}{1 + A(z)F(z)}}_{H_{\mathrm{E}}(z)} E(z) + \underbrace{\frac{A(z)}{1 + A(z)F(z)}}_{H_{\mathrm{X}}(z)} X(z) \tag{I.8}$$

Hence the quantizer error $E(z)$ and the input signal $X(z)$ experience different (though interlinked) transfer functions $H_{\mathrm{E}}(z)$ and $H_{\mathrm{X}}(z)$, respectively.

That $H_{\mathrm{E}}(z)$ and $H_{\mathrm{X}}(z)$ are different is key to ΣΔ-modulators as it allows different filtering of $E(z)$ and $X(z)$. Classically, the modulator is designed such that $E(z)$ is high-pass filtered while $X(z)$ remains essentially unchanged. In this case, $E(z)$ is attenuated by $H_{\mathrm{E}}(z)$ at low frequencies. If $X(z)$ is subsequently confined to a bandwidth $f \in [0, f_{\mathrm{B}}]$ over these low frequencies, then the SNR of $Y(z)$ over $[0, f_{\mathrm{B}}]$ can be quite high. Such confinement typically requires oversampling $X(z)$.

The high-pass filtering of $E(z)$, though, also amplifies $E(z)$ at higher frequencies. To remove this excess high-frequency noise, $Y(z)$ is low-pass filtered in the digital domain. During this digital filtering, it is also often simultaneously decimated down to the Nyquist rate $2f_{\mathrm{B}}$ [*Norsworthy and Crochiere*, 1997].

Stanford University as evidenced by many recent publications [*Vleugels et al.*, 2001; *Nam et al.*, 2005; *Kulchycki et al.*, 2008].

[9] Note that the generalized ΣΔ-modulator of Figure I.3 is different from the noise-shaping feedback coder of Figure I.1 in that, while the noise-shaping feedback coder operates on the sub-ADC quantizer error explicitly, the ΣΔ-modulator operates only on the sub-ADC output.

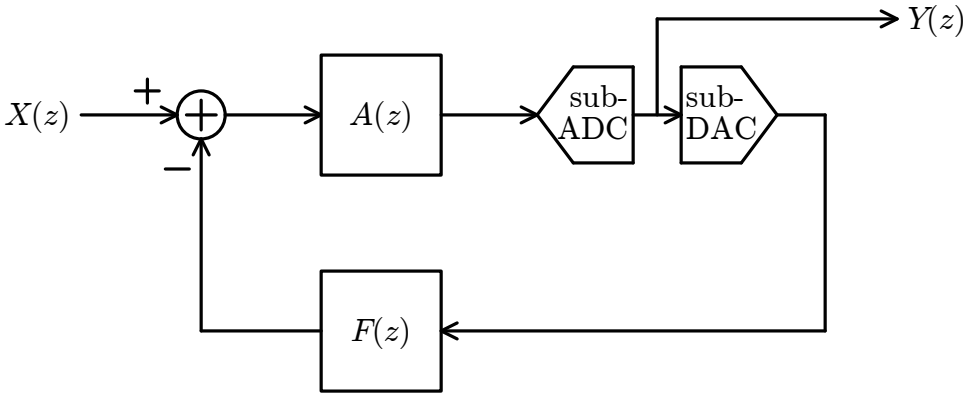

(a) General ΣΔ-modulator.

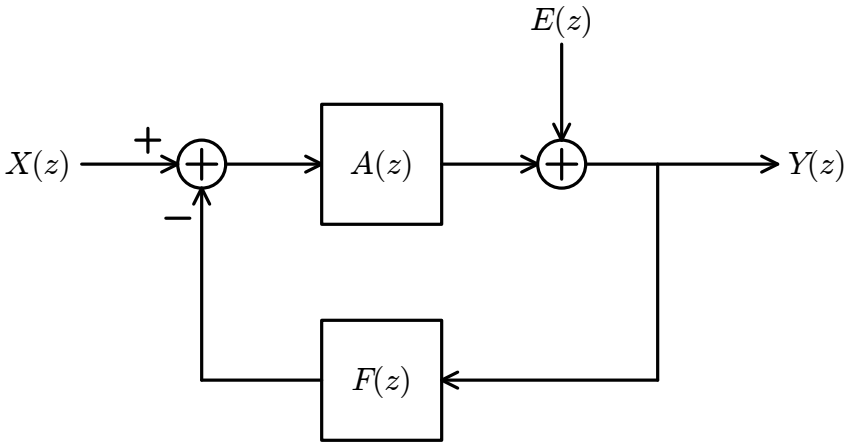

(b) General ΣΔ-modulator under PQN model.

Figure I.3: General ΣΔ-modulator (a) architecture and (b) assuming the PQN model for the sub-ADC.

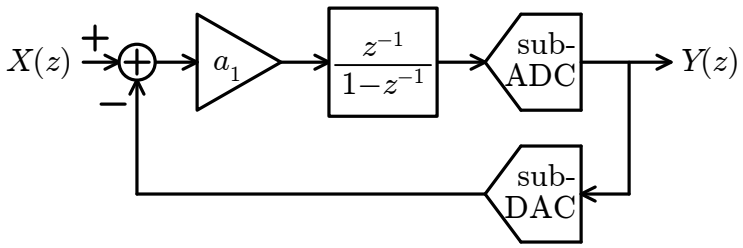

(a) 1st-order noise-differencing, ΣΔ-modulator.

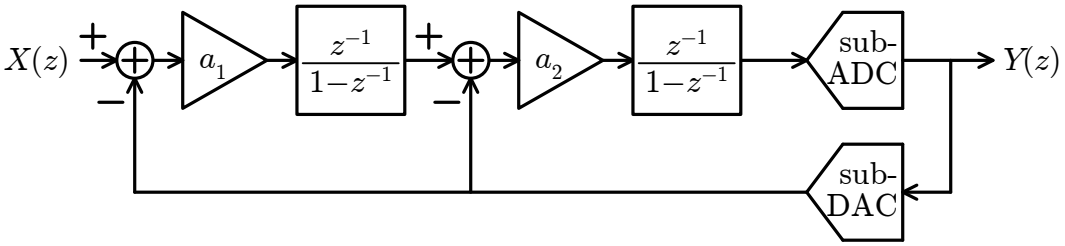

(b) 2nd-order noise-differencing, ΣΔ-modulator.

Figure I.4: Examples of general (a) 1st-order and (b) 2nd-order noise-differencing ΣΔ-modulators.

I.3 Practice

With the theory in hand, this appendix now turns to the design of particular modulator architectures, with a focus on baseband modulators[10,11]

I.3.1 Noise-Differencing Examples

To concretize the theory of the previous section, Figure I.4 shows the architectures of generalized 1st-order and 2nd-order ΣΔ-modulators. The feedforward paths of these modulators are implemented by a series of integrators, while the feedback paths are reduced to simple unity gain paths. The operation of each modulator can be described

[10]That is, modulators that convert input signal components over $f \in [0, f_B]$, where $2f_B < f_S$, with high SNR.

[11]It should be noted that ΣΔ-modulators are not restricted to high-pass $H_E(z)$ that provide high SNR at baseband frequencies. Indeed, it is possible to design ΣΔ-modulators with band-notch $H_E(z)$ that garner high SNR over a narrow bandwidth centered at a nonzero frequency. Such "bandpass modulators" often find utility in wireless communications applications, for example [*Jantzi et al.*, 1997].

by assuming the PQN model for the sub-ADC. Doing so for the architecture of Figure I.4(a), and assuming $a_1 = 1$, it can be shown that the output of the 1st-order $\Sigma\Delta$-modulator is:

$$Y(z) = z^{-1}X(z) + \left(1 - z^{-1}\right) E(z) \tag{I.9}$$

As modulators are conventionally named by their $H_{\mathrm{E}}(z)$, Figure I.4(a) depicts a 1st-order noise-differencing $\Sigma\Delta$-modulator. Repeating this process for the architecture of Figure I.4(b), and assuming $a_1 = 1/2$ and $a_2 = 2$, it can be shown that the output of the 2nd-order $\Sigma\Delta$-modulator is:

$$Y(z) = z^{-2}X(z) + \left(1 - z^{-1}\right)^2 E(z) \tag{I.10}$$

And hence Figure I.4(b) depicts a 2nd-order noise-differencing $\Sigma\Delta$-modulator.[12]

To illustrate the operation of a $\Sigma\Delta$-modulator, Figure I.5 shows the simulated output power spectrum of the 2nd-order, noise-differencing $\Sigma\Delta$-modulator while processing an input sinusoid. While some circuit noise is included in the simulation, the majority of the noise in Figure I.5 is quantization noise. As a guide, Figure I.5 also includes a red dashed line that indicates the expected quantization noise level of the sub-ADC alone (i.e., without the $\Sigma\Delta$-modulator loop) assuming the PQN model. The high-pass shaping of this quantization noise is evident in Figure I.5(a). This shaping gains the quantization noise at high frequencies, but attenuates it at low frequencies: the attenuation is more clearly seen in Figure I.5(b), which focuses on the low-frequency regime. If the output spectrum is subsequently digitally filtered by a low-pass filter that selects a regime of attenuation, and if the input sinusoid is confined to this regime, then the resulting SNR can be very high.

[12] For Equations (I.9) and (I.10), the values of a_1 and a_2 are chosen to reproduce these classical noise-differencing transfer functions. However, in practice the coefficients are often altered based on other architectural concerns. For example, in a 2nd-order, noise-differencing $\Sigma\Delta$-modulator with a single-bit sub-ADC, often $a_1 = a_2 = 1/2$ to make the integrators the same and thus simplify the design (note that, given a single-bit sub-ADC, the gain of the final integrator can be set to any convenient value since only its sign is relevant). And compensating for circuit limitations, such as the linear output range of the integrators, can often result in altered coefficients.

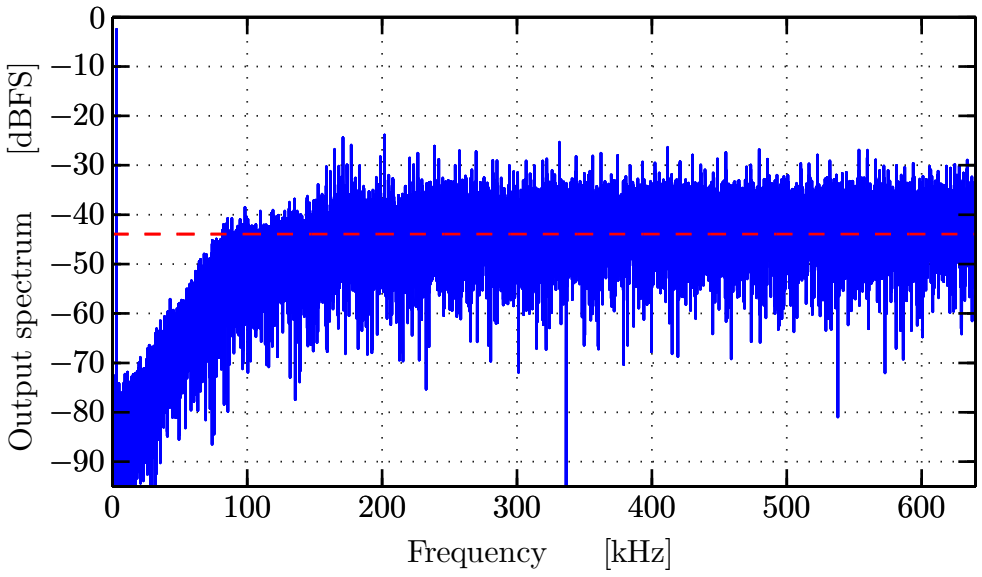

(a) Full spectrum.

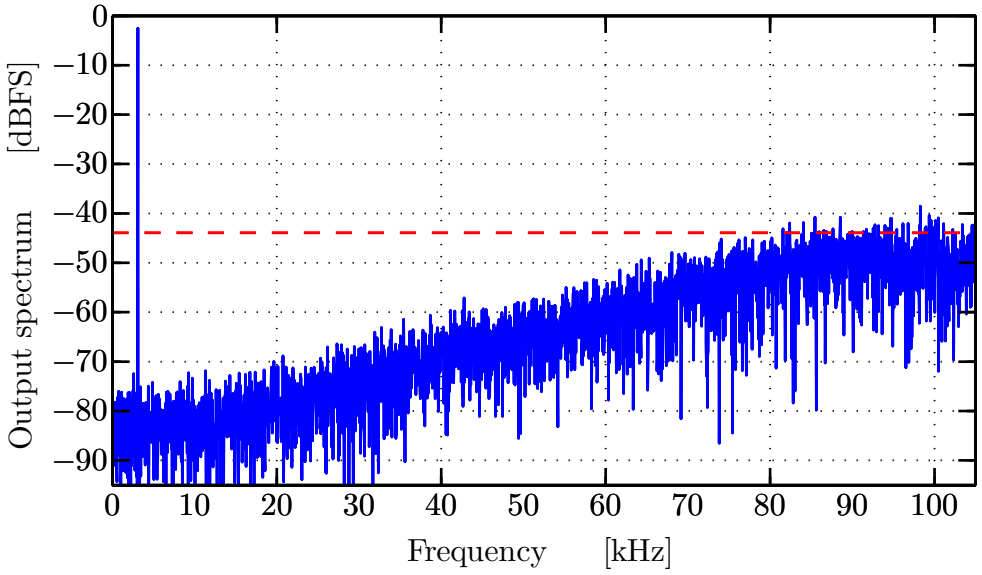

(b) Zoomed spectrum.

Figure I.5: Simulated output spectrum for a 2nd-order noise-differencing ΣΔ-modulator as per Figure I.4(b) with $a_1 = 1/2$, $a_2 = 2$, and a single-bit sub-ADC. Modulator processes 20-kHz input sinusoid at 1.28 MS/s; 32,768-point FFT (∼40 Hz/bin) shown. Spectrum (a) in full and (b) zoomed to low frequencies. Red dashed line gives expected quantization noise floor of same sub-ADC with oversampling, but not noise-shaping, under PQN model.

I.3.2 Performance Prediction

To quantify the degree of SNR improvement, compare the IDR achieved by three architectures: 1) an M-bit ADC sampling at the Nyquist rate, 2) an oversampled M-bit ADC, and 3) an M-bit ADC embedded in an L-th order, noise-differencing $\Sigma\Delta$-modulator.

1. *Nyquist-rate M-bit ADC*
   Assume the M-bit ADC has the quantizer characteristic of Figure 3.6. The maximum input sinusoid before overrange then has an amplitude of $V_{\mathrm{FS}}/2$ for a maximum signal power of $V_{\mathrm{FS}}^2/8$. Under the PQN model, the quantization noise has power $\Delta^2/12$ where $\Delta = V_{\mathrm{FS}}/2^M$. The IDR is thus:

$$\mathrm{IDR}_{\mathrm{standard}} = \frac{V_{\mathrm{FS}}^2/8}{\Delta^2/12} = \frac{3}{2}2^{2M} \tag{I.11}$$

   and in dB:

$$\mathrm{IDR}_{\mathrm{standard}}\ [\mathrm{dB}] \simeq 1.76 + 6.02M \tag{I.12}$$

2. *Oversampled M-bit ADC*
   If the M-bit ADC is oversampled, then while the quantization noise power $\Delta^2/12$ is evenly spread over the entire spectrum, only a small portion of that noise lies over the input bandwidth. Assume the input is confined to the bandwidth $f \in [0, f_{\mathrm{B}}]$, where $2f_{\mathrm{B}} < f_{\mathrm{S}}$ and f_{S} is the sampling frequency. The total inband quantization noise is then:

$$Q_{\mathrm{oversampled}} = \int_{-f_{\mathrm{B}}}^{f_{\mathrm{B}}} \frac{Q_{\mathrm{noise}}}{f_{\mathrm{S}}}\,\mathrm{d}f = \left(\frac{\Delta^2}{12}\right)\frac{2f_{\mathrm{B}}}{f_{\mathrm{S}}} \tag{I.13}$$

   Note that f here is the continuous-time frequency. The quantity $2f_{\mathrm{B}}/f_{\mathrm{S}}$ is called the oversampling ratio OSR. The IDR is then:

$$\mathrm{IDR}_{\mathrm{oversampled}} = \frac{V_{\mathrm{FS}}^2/8}{Q_{\mathrm{oversampled}}} = \frac{3}{2}2^{2M}\mathrm{OSR} \tag{I.14}$$

and in dB:

$$\mathrm{IDR}_{\mathrm{oversampled}}\ [\mathrm{dB}] \simeq 1.76 + 6.02M + 10\log_{10}(\mathrm{OSR}) \tag{I.15}$$

3. *M-bit ADC in an L-th order, noise-differencing ΣΔ-modulator*
Assume the M-bit ADC is both oversampled and embedded in a ΣΔ-modulator that implements the L-th order noise-differencing:

$$H_{\mathrm{E}}(z) = \left(1 - z^{-1}\right)^{L} \tag{I.16}$$

Assuming the same input bandwidth restrictions as in the oversampled ADC case, the total inband quantization noise of the ΣΔ-modulator is then:

$$Q_{\Sigma\Delta} = \int_{-f_{\mathrm{B}}}^{f_{\mathrm{B}}} \frac{Q_{\mathrm{noise}}}{f_{\mathrm{S}}} |H_{\mathrm{E}}(f)|^2 \,\mathrm{d}f = \int_{-f_{\mathrm{B}}}^{f_{\mathrm{B}}} \frac{\Delta^2/12}{f_{\mathrm{S}}} 2^{2L} \left|\sin\left(\frac{\pi f}{f_{\mathrm{S}}}\right)\right|^{2L} \mathrm{d}f \tag{I.17}$$

Assuming $2f_{\mathrm{B}} \ll f_{\mathrm{S}}$, Equation (I.17) can be approximated by noting that $\sin(x) \simeq x$ for small x and thus:

$$Q_{\Sigma\Delta} \simeq \frac{\Delta^2/12}{f_{\mathrm{S}}} \int_{-f_{\mathrm{B}}}^{f_{\mathrm{B}}} 2^{2L} \left(\frac{\pi f}{f_{\mathrm{S}}}\right)^{2L} \mathrm{d}f = \left(\frac{\Delta^2}{12}\right)\left(\frac{\pi^{2L}}{2L+1}\right)\left(\frac{1}{\mathrm{OSR}}\right)^{2L+1} \tag{I.18}$$

The IDR of the ΣΔ-modulator is:

$$\mathrm{IDR}_{\Sigma\Delta} = \frac{V_{\mathrm{FS}}^2/8}{Q_{\Sigma\Delta}} \simeq \frac{3}{2} 2^{2M} \frac{2L+1}{\pi^{2L}} \mathrm{OSR}^{2L+1} \tag{I.19}$$

and in dB:

$$\begin{aligned} \mathrm{IDR}_{\Sigma\Delta}\ [\mathrm{dB}] \simeq\ & 1.76 + 6.02M + (2L+1)10\log_{10}(\mathrm{OSR}) \\ & + 10\log_{10}(2L+1) - (2L)10\log_{10}(\pi) \end{aligned} \tag{I.20}$$

The benefit of just oversampling can be seen by comparing $\mathrm{IDR}_{\mathrm{standard}}$ with $\mathrm{IDR}_{\mathrm{oversampled}}$: every doubling in OSR delivers a 3-dB improvement in IDR. In the ΣΔ-modulator, this benefit is further augmented by noise shaping. Comparing $\mathrm{IDR}_{\mathrm{standard}}$ with $\mathrm{IDR}_{\Sigma\Delta}$, every doubling in OSR delivers a $(6L+3)$-dB improvement in IDR: for

the 1st-order and 2nd-order noise-differencing ΣΔ-modulators, the improvement is thus 9-dB per OSR octave, and 15-dB per OSR octave, respectively.[13] Hence, even with very coarse quantizers, noise-differencing ΣΔ-modulators can deliver high IDR via a combination of noise-shaping and aggressive oversampling.

I.3.3 Design Issues

In reality, many practical design issues often prevent ΣΔ-modulators from achieving the IDR predicted by Equation (I.20), especially in the case of high-IDR modulators. Many of these concerns are addressed in the following subsections. It should be noted, though, that this discussion is by no means complete: many other issues can limit ΣΔ-modulator performance.

I.3.3.1 Choice of sub-ADC Resolution

Classically, ΣΔ-modulators use single-bit ADCs and, correspondingly, single-bit DACs. The reason is that, given the aforementioned sensitivity of feedback systems to nonidealities in their feedback path, for high-IDR applications the sub-DAC must be highly linear. Notably, single-bit sub-DACs can be made inherently linear (see footnote 26 of Chapter 4).

Single-bit ΣΔ-modulators, though, present their own challenges. For example, single-bit ΣΔ-modulators are known to be only conditionally stable, in particular, such modulators are bounded-input bounded-output (BIBO) stable only over a limited range of input signals [*Adams and Schreier*, 1997]. Hence the input signal value must be properly bounded to prevent triggering instability. While these input-range bounds do not overly compromise IDR in lower-order (i.e., $L \leq 2$) noise-differencing modulators, in higher-order (i.e., $L \geq 3$) modulators the bounds result in severe IDR loss, especially in high-IDR modulators [*Adams and Schreier*, 1997, pp. 153–154]. However, analytically predicting the stability of single-bit ΣΔ-modulators via standard linear methods is complicated by the gain ambiguity of the single-bit

[13] Offsetting this gain with OSR is an initial penalty that grows with L, represented by the final two terms of Equation (I.20).

sub-ADC: the gain of this nonlinear element is not well-defined. Hence stability analysis of such modulators must rely on a combination of analytical predictions, time-domain simulations, and designer experience.[14,15,16]

Stability is easier to analyze with multibit sub-ADCs, where the sub-ADC gain is no longer ambiguous. To guarantee the linearity of the corresponding multibit sub-DAC, though, designers often recourse to a variety of linearizing techniques, such as element trimming (e.g., laser trimming), dynamic element matching (e.g., *Vleugels et al.* [2001]),[17] or digital calibration (e.g., *Carley et al.*

[14]One solution for preventing unbounded states inside the modulator is to limit the integrator outputs, either by imposing nonlinear limiting circuits or by adding detection-and-correction circuits that sense overrange and subsequently reset various circuits to quash it [*Adams*, 1997, pp. 183–185].

[15]Notably, lower-order noise-differencing modulators are also suspect to another type of instability: idle tones. Idle tones are periodic signals output from the ΣΔ-modulator resulting from the DC input value [*Adams et al.*, 1991; *Adams*, 1997, pp. 185–186]. Idle tones arise more readily in lower-order modulators (such as 1st-order and 2nd-order modulators) than higher-order modulators—the intuitive explanation being that the increased "randomness" incurred around the modulator feedback loop of higher-order modulators tends to decohere and attenuate the tones [*Adams et al.*, 1991]—although they may nonetheless still appear in the latter. A common technique for preventing idle tones is to increase the circuit noise to be higher than the quantization noise: the increased randomness upsets and reduces the tones [*Adams*, 1997, p. 185]. The effectiveness of this strategy, though, depends on the application: for example, the human ear can detect tones as deep as 20-dB below a white noise floor [*Adams*, 1997, p. 185] and hence even small tones can compromise performance in audio applications.

[16]Another suspect analysis in single-bit ΣΔ-modulators is the use of the PQN model itself. Recall from footnote 13 of Chapter 3 that, for bounded quantizers, the PQN model holds approximately when 1) the quantizer does not overrange, 2) M_{level} (the number of quantizer levels) is asymptotically large, 3) Δ (the quantizer step size) is asymptotically small, and 4) the joint PDF of the quantizer input at different times is smooth [*Gray*, 1997, p. 48]. However, a single-bit sub-ADC in a ΣΔ-modulator violates many—if not all—of these conditions. In fact, for a 1st-order, noise-differencing ΣΔ-modulator, it can be shown that the autocorrelation of the quantizer error is not white nor even continuous [*Gray*, 1997, pp. 62–64]. Thus, though the PQN model approximation provides valuable insight into the noise-shaping process, and can even provide accurate IDR estimates for lower order, single-bit, noise-differencing ΣΔ-modulators, it must be kept in mind that for single-bit ΣΔ-modulators it is not an exact solution, and does not completely predict modulator behavior (for example, it does not predict the existence of the idle tones discussed in footnote 15 of this appendix).

[17]In the switched-capacitor circuitry often favored by ΣΔ-modulator designers, multibit sub-DACs are typically implemented by switched banks of unit-capacitors: switching a particular subset of the unit-capacitors to one reference or another implements a particular DAC code. Typically each DAC code is mapped to a unique unit-capacitor configuration. Under dynamic element matching, though, this mapping is not fixed but instead changes from sample to sample (typically, although the number of unit-capacitors switched to one reference or another remains constant, the actual set of unit-capacitors so-switched changes). The objective of dynamic element matching is to choose a mapping strategy that decorrelates the mismatch-induced sub-DAC errors from the input signal

[1997]). Nonetheless, multibit ΣΔ-modulators remain popular in applications such as wideband communications where the large input signal bandwidths, combined with process technology limitations, quickly constrain the OSR to lower values: in these applications, proper choice of the sub-ADC resolution M (along with proper choice of the modulator order L as discussed in the next section) can provide a power-efficient means of achieving high IDR [*Vleugels et al.*, 2001].[18]

I.3.3.2 Choice of $H_{\mathrm{E}}(z)$

Another way to increase the IDR is to increase L, the order of the noise-differencing. More aggressive noise-differencing yields greater attenuation of $E(z)$ at low frequencies, boosting the achieved SNR.

Such higher-order noise-differencing modulators can be achieved by continuing the pattern of Figure I.4 and placing additional integrators in the feedforward path. However, if a single-bit sub-ADC is used the resulting noise-differencing modulators often display lower IDR than that predicted by the analytical methods of Section I.3.2. This underperformance is due to the aforementioned stability limitations, which become more severe in higher-order modulators. To compensate for these stability problems, higher-order modulators thus often adopt alternate $H_{\mathrm{E}}(z)$ transfer functions such as Butterworth or inverse-Chebyshev characteristics: while single-bit modulators implementing these alternate $H_{\mathrm{E}}(z)$ are still only conditionally stable, they often have greater stable input ranges, allowing greater IDR. The resulting modulator topologies, though, can become quite complicated, as illustrated in Figure I.6, which shows an architecture used to implement a 5th-order inverse-Chebyshev ΣΔ-modulator [*Adams et al.*, 1991; *Adams*, 1997, Sect. 5.6]. The feedback paths a_k

and ultimately to shift the incurred noise out of the signal band [*Vleugels et al.*, 2001]. Different dynamic element matching techniques differ in their mapping strategies.

[18] An interesting alternative to a multibit sub-ADC and multibit sub-DAC is to maintain the multibit sub-ADC (which increases M) but to only use a single-bit sub-DAC (which allows inherently linear sub-DAC design). Indeed, with proper signal processing of the modulator output, it can be shown that such an architecture can achieve the performance promised by the traditional multibit sub-ADC with corresponding multibit sub-DAC modulator [*Carley et al.*, 1997, pp. 273–275]. However, the aforementioned signal processing requires precise matching between the analog-implemented $H_{\mathrm{E}}(z)$ and a collection of digitally implemented filters: the inaccuracies of such matching often limit the performance of this technique in practice [*Carley et al.*, 1997, p. 275].

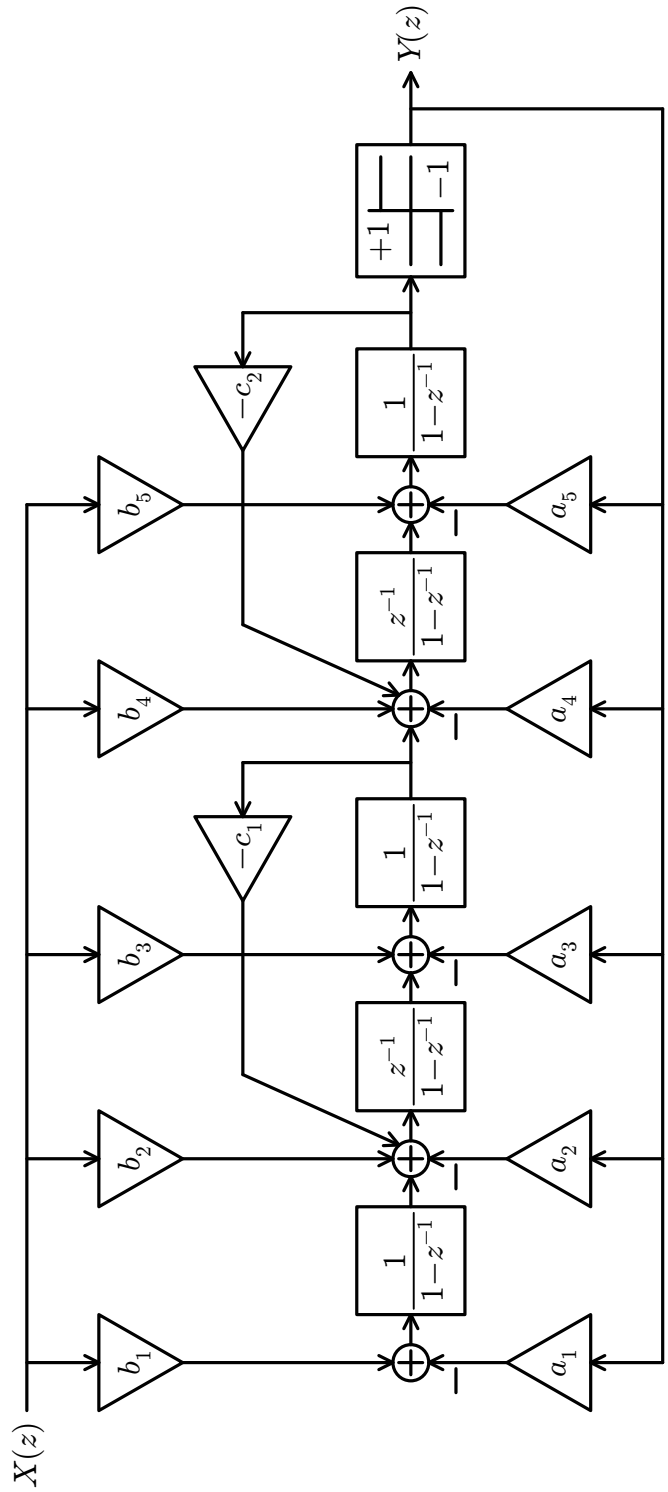

Figure I.6: Example of a general 5th-order ΣΔ-modulator used to implement an inverse-Chebyshev $H_{\mathrm{E}}(z)$ characteristic, as described in *Adams et al.* [1991].

and c_k enable the modulator to implement the inverse-Chebyshev $H_{\mathrm{E}}(z)$; in particular, the local resonator feedback paths c_k enable movement of the zeros of $H_{\mathrm{E}}(z)$ to finite positive frequencies (versus the noise-differencing modulator, wherein all $H_{\mathrm{E}}(z)$ zeros are confined to DC), corresponding to the notches required of an inverse-Chebyshev characteristic [*Adams*, 1997, p. 180]. The feedforward paths b_k grant a certain degree of decoupling of $H_{\mathrm{E}}(z)$ and $H_{\mathrm{X}}(z)$, helping to reduce IDR loss from the $H_{\mathrm{X}}(z)$ characteristic [*Adams*, 1997, pp. 179–180].

Alternately, cascade topologies can be used to achieve higher-order noise-differencing. Cascaded ΣΔ-modulators are constructed from a collection of noise-differencing modulators cascaded such that the order of the noise-differencing of the cascade is the sum of the orders of the noise-differencing of the constituent modulators. Since the cascaded ΣΔ-modulator uses only feedforward (and not feedback) paths, it largely adopts the stability of its constituents: if lower-order noise-differencing modulators are used, the cascaded ΣΔ-modulator stability can be quite good. The concept is illustrated in Figure I.7, which shows a cascade of a 2nd-order noise-differencing ΣΔ-modulator with a 1st-order noise-differencing ΣΔ-modulator, called a 2-1 cascaded ΣΔ-modulator. Here, the quantization noise of the first modulator is computed and then input to a second modulator. With proper digital signal processing of the two modulator outputs (indicated by $H_1(z)$ and $H_2(z)$ in Figure I.7), it can be shown that only the quantization error of the second sub-ADC remains, and furthermore that this quantization error is noise-shaped by the combined order of both modulators, for a 3rd-order noise-differencing response.[19] In general, cascaded

[19]For example, assume that $a_1 = 1/2$ and $a_2 = 2$, and that $a_3 = 1$. Then the 2nd-order and 1st-order modulators of the 2-1 cascaded ΣΔ-modulator of Figure I.7 implement the canonical transfer functions of Equations (I.10) and (I.9), respectively, that is:

$$Y_1(z) = z^{-2}X(z) + \left(1 - z^{-1}\right)^2 E_1(z)$$

and:

$$Y_2(z) = z^{-1}X_2(z) + \left(1 - z^{-1}\right) E_2(z)$$

Further assume that $X_2(z) = E_1(z)$, as accomplished by setting $b_1 = b_2 = b_3 = 1$. Then, if the digital signal processing is implemented as:

$$H_1(z) = z^{-1}$$

and:

$$H_2(z) = \left(1 - z^{-1}\right)^2$$

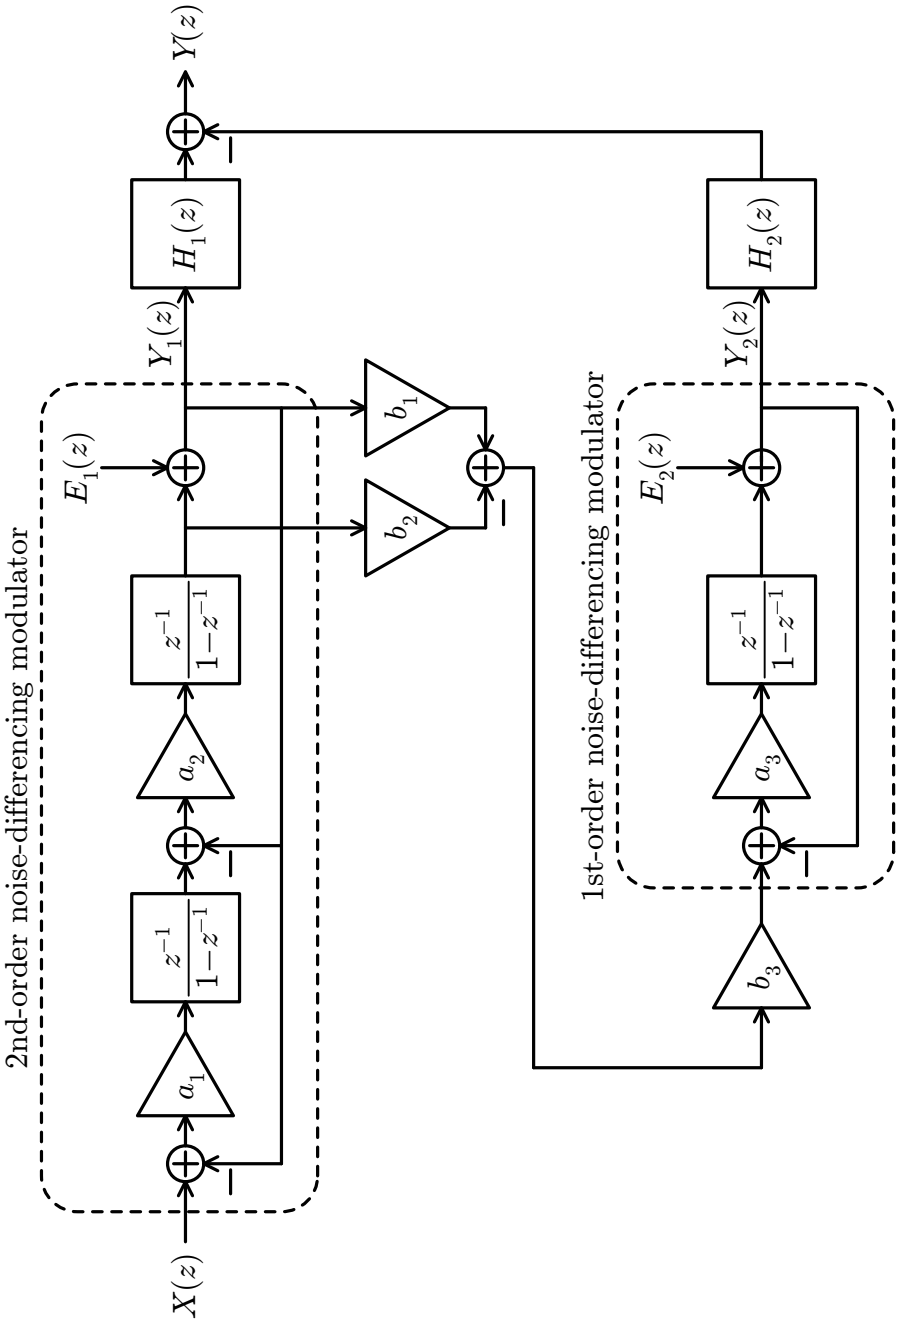

Figure I.7: Example of a general 2-1 cascaded ΣΔ-modulator.

modulators eliminate the quantizer error of the upstream modulators so that only the quantizer error of the last modulator remains, and furthermore is shaped by the combined order of the entire cascade. This principle, then, can be readily applied to 2-2 cascaded ΣΔ-modulators (e.g., *Nam et al.* [2005]), or extended to create 2-2-1 cascaded ΣΔ-modulators (e.g., *Vleugels et al.* [2001]), for example. This effect enables cascaded modulators to achieve high-order noise-differencing while retaining simpler, more stable individual modulators.

---

it can be shown that:

$$Y(z) = z^{-3}X(z) + \left(1 - z^{-1}\right)^3 E_2(z)$$

Hence the cascaded modulator cancels $E_1(z)$ (the quantizer error of the first modulator) and provides 3rd-order noise-differencing of $E_2(z)$ (the quantizer error of the second modulator). In practice, the values of the a_k and b_k coefficients are often altered to accommodate various circuit concerns (e.g., the linear output range of the integrators).

Appendix J

Breakout Test Circuits on SVADC-1

In addition to the SVADC-1 converter proper (described throughout this dissertation) and the NMOS test devices (described in Appendix F), the SVADC-1 chip also includes breakout circuits. These breakout circuits are intended for separate radiation testing of the internal circuitry of the SVADC-1 converter. Each breakout circuit is thus a circuit used within the SVADC-1 converter, isolated and given its own pad ring.[1] Care is taken to ensure that both the schematics and layouts of the breakout circuits mimic their equivalents in the converter proper as closely as possible.

These breakout circuits are seen to the right of the chip micrograph in Figure 5.40. Three pad rings, vertically oriented, are present. From top to bottom, these pad rings encapsulate a bias breakout circuit, an operational amplifier breakout circuit, and a flash ADC breakout circuit. The remainder of this appendix describes each breakout circuit in turn.

All of the breakout circuits were confirmed functional, however, given the strong radiation performance of the SVADC-1 converter proper, they were not rigorously tested for performance, nor radiation tested.

[1] Each pad ring includes complete ESD protection as described in Section 5.5.5. Even though some breakout circuits contain both analog and digital pads, due to limited area each breakout circuit has only one ESD ring; for all breakout circuits, this ring is connected to $V_{\mathrm{DD,A}}$.

J.1 Bias Breakout Circuit

The bias breakout circuit contains a portion of the analog biasing for an operational amplifier. It is intended to assess the impact of, for example, radiation-induced threshold voltage shifts on such circuits.

The schematic for the bias breakout circuit is shown in Figure J.1. Both NMOS and PMOS biasing circuits are included. Note that the bias breakout circuit repeats the biasing circuits of operational amplifier of the track-and-hold amplifier (compare with Figure 5.28): the track-and-hold amplifier—as opposed to the calibrated or uncalibrated residue amplifier—is chosen for breakout since, being upstream, it is highly likely that changes in the track-and-hold amplifier would noticeably affect overall converter performance.

J.2 Operational Amplifier Breakout Circuit

For the same reason, the operational amplifier breakout circuit contains the operational amplifier used in the track-and-hold amplifier. The schematic for the operational amplifier breakout circuit is shown in Figure J.2. Note that the device sizings are the same as in Figure 5.28. In addition to the operational amplifier inputs and outputs and the clock input, a variety of other signals are also brought to pads, including the V_{SET} and V_{CM} circuits of the $V_{\mathrm{DD}}/2$ generator, and multiple signals from the operational amplifier core. To better catalogue the latter, the operational amplifier core is shown in Figure J.3 with the padded signals labeled. Finally, the operational amplifier breakout circuit also adopts the same timing as the track-and-hold amplifier so that, for the common mode feedback, $\phi_{1\mathrm{D}}$ is in phase with the input clock while ϕ_2 is anti-phase.

To test performance, the operational amplifier would most likely be configured in a negative feedback configuration. As only the operational amplifier proper is included in the breakout circuit, the particular configuration would be implemented on the breakout circuits test board. However, such a scheme is highly sensitive to board-level component and layout variations. While the breakout circuit of Figure J.2

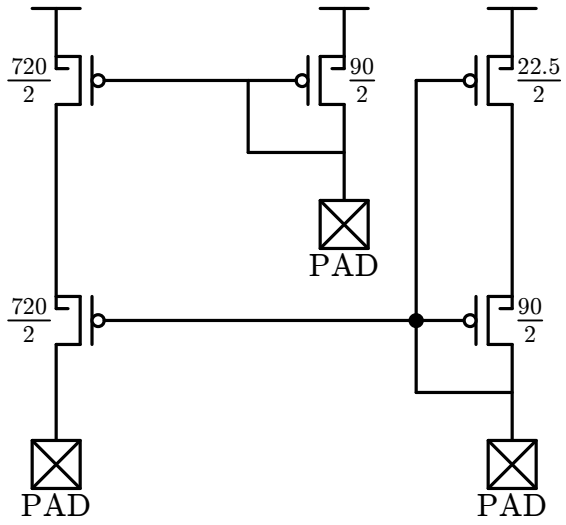

(a) PMOS biasing section.

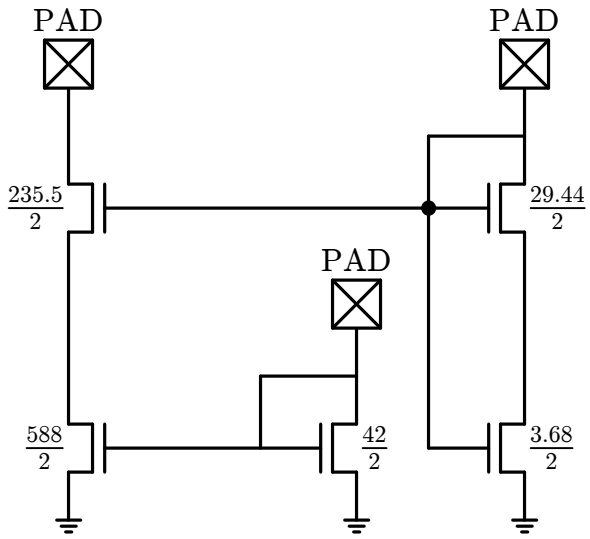

(b) NMOS biasing section.

Figure J.1: Bias breakout circuit. (a) PMOS biasing section and (b) NMOS biasing section. All supplies tied to $V_{\mathrm{DD,A}}$. Biasing circuits taken from the operational amplifier of the track-and-hold amplifier; compare Figure 5.28.

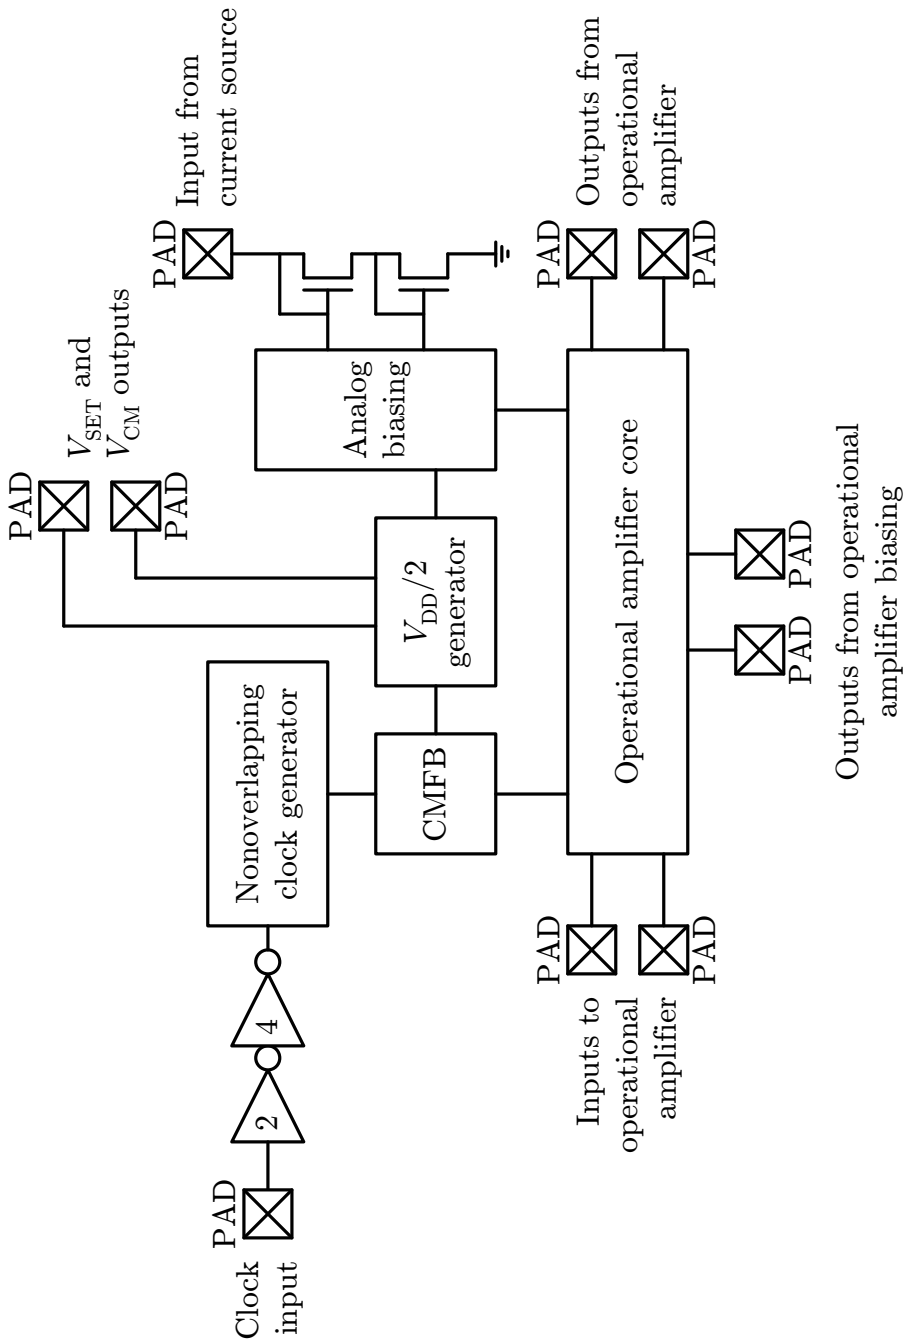

Figure J.2: Operational amplifier breakout circuit. Operational amplifier taken from the track-and-hold amplifier: core is as per Figure 5.28, although see Figure J.3 for a more detailed view of the pad-connected signals. Non-overlapping clock generator as per Figure 5.14, operational amplifier common mode feedback (CMFB) as per Figure 5.31, $V_{DD}/2$ generator as per Figure 5.36. Supplies for clocking circuits tied to $V_{DD,CLK}$, all other supplies tied to $V_{DD,A}$.

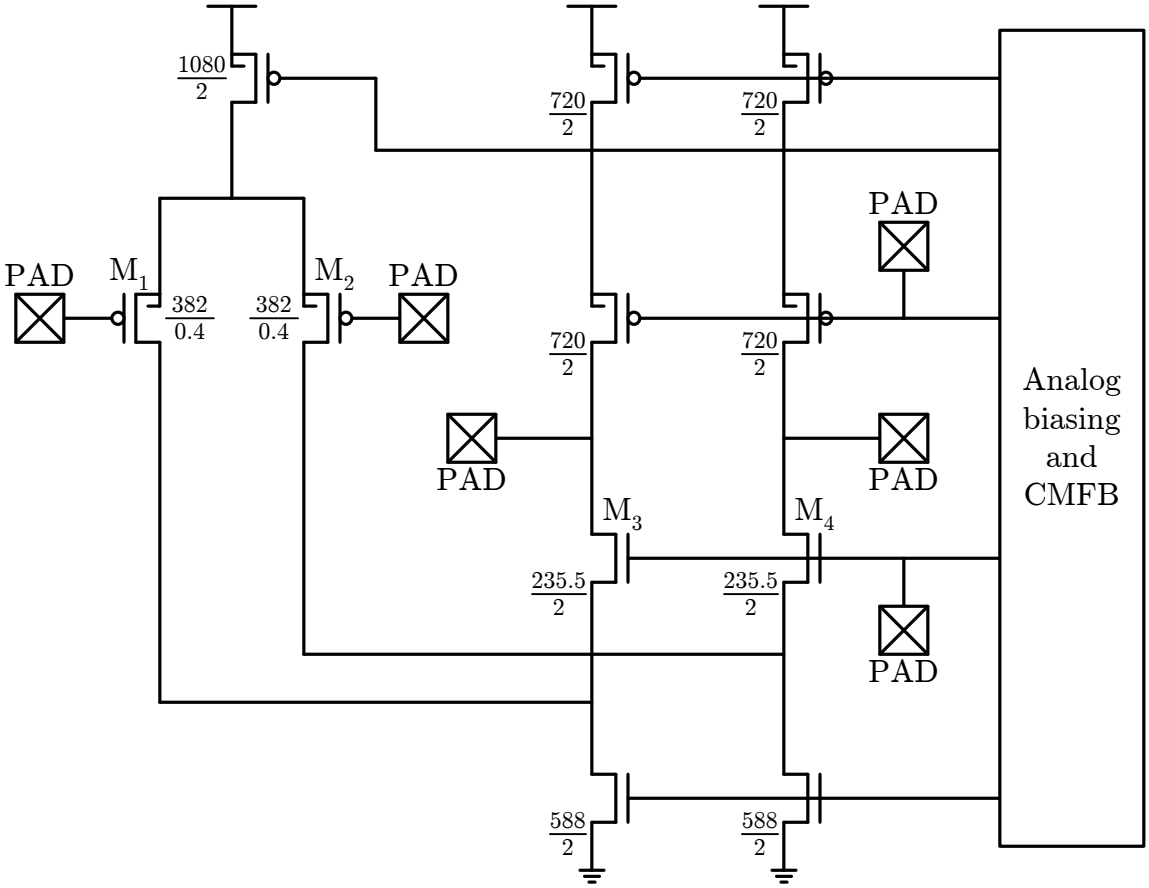

Figure J.3: Operational amplifier breakout circuit core, labeling signals connected to pads. All supplies tied to $V_{\mathrm{DD,A}}$. For use with Figure J.2.

nonetheless remains viable for performance testing under radiation, in future versions it is recommended that, instead of just the operational amplifier itself, the feedback paths (if not a complete switched-capacitor amplifier configuration) be implemented on chip.

J.3 Flash ADC Breakout Circuit

The flash ADC breakout circuit contains a 3.1-bit flash ADC as used in the calibrated stages. The schematic for the flash ADC breakout circuit is shown in Figure J.4. The heart of the flash ADC is the series of 8 comparators, each of which is implemented as described in Section 5.5.2. The output of each comparator is then buffered and outputted to a pad. This output buffer adopts the buffer-flop-buffer structure described in Section 5.5.5.

Regarding timing, each comparator of the flash ADC breakout circuit samples the same analog input on the falling edge of the input clock. The comparison result that appears on the output pad is triggered on the following rising edge of the input clock so that, from an external perspective, the comparison requires half an input clock cycle. Being flopped, the comparison result remains valid until the next rising edge of the input clock.

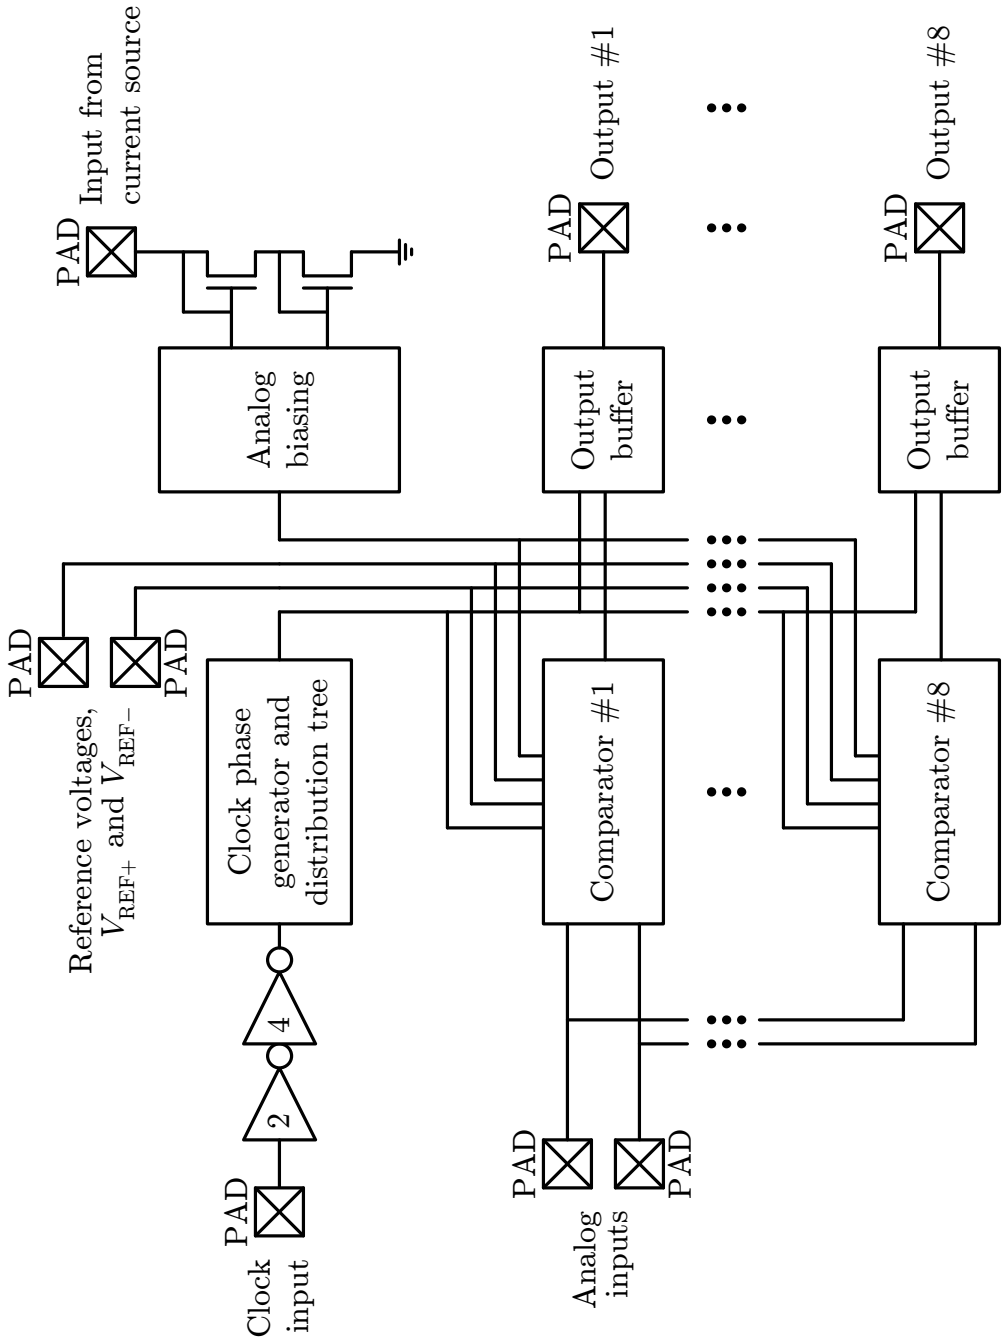

Figure J.4: Flash ADC breakout circuit. Comparators as described in Section 5.5.2 and configured for a 3.1-bit ADC as used in the calibrated stages. Output buffers as per description of Section 5.5.5. Clock phase generator and distribution tree includes non-overlapping clock generator as per Figure 5.14. Supplies for clocking circuits tied to $V_{\mathrm{DD,CLK}}$, supplies for comparators tied to $V_{\mathrm{DD,D}}$, and supplies for the remainder tied to $V_{\mathrm{DD,A}}$.

Appendix K

SVADC-1 Burn-In Testing

This appendix describes the burn-in testing of the SVADC-1. This testing screens SVADC-1 parts for incorporation into the BBR (BroadBand Receiver) instrument of the WIPER (Wave-Induced Precipitation of Electron Radiation) experiment flying aboard the DSX (Deployable Structures eXperiment) satellite. Two batches of SVADC-1s are tested, the first consisting of 23 SVADC-1s, the second of 16 SVADC-1s. Both batches are burn-in tested at a temperature of 125°C for at least 160 hours, and assessed for functionality and performance both before and after burn-in.

The testing of the first batch took place in the summer of 2008, and of the second batch in the autumn of 2008. All testing was conducted at Lockheed Martin Corporation, in particular, at their facilities in Palo Alto, California. The experimenters were Bob Bumala, Clem Tillier, and Ken Holsworth, all with Lockheed Martin Corporation. The testing was performed in compliance with Military Standard 883G, Method 1015.9, Class B [*MIL-STD-883G*, 2006, Method 1015.9]. All tested devices remained functional after burn-in and a ranking of final performance, based on a custom metric developed in support of the WIPER experiment, is presented.

K.1 Purpose

Burn-in testing is a method for ensuring reliability of a population of parts. Specifically, it is

> performed for the purpose of screening or eliminating marginal devices, those with inherent defects or defects resulting from manufacturing aberrations which cause time and stress dependent failures. In the absence of burn-in, these defective devices would be expected to result in infant mortality or early lifetime failures under use conditions. Therefore, it is the intent of this screen to stress microcircuits at or above maximum rated operating conditions or to apply equivalent screening conditions, which will reveal time and stress dependent failure modes with equal or greater sensitivity. [*MIL-STD-883G*, 2006, Method 1015.9]

Such screening is justified since the failure rate[1] for a population of parts often displays a "bathtub curve" characteristic as sketched in Figure K.1. As shown, the failure rate is often higher in early life (labeled as the "Early Failure Period") before dropping and

[1] Intuitively, the concept of failure rate can be illustrated by the following example:

> Consider a population of 1000 units that start operating at time zero. Over time, units fail one by one. Say that at 5000 hours the fourth unit has already failed and another unit fails in the next hour. How would we define a "rate of failure" for the units operating in the hour between 5000 and 5001? Since 996 units were operating at the start of that hour and one failed, a natural estimate of the failure rate (for units at 5000 hr of age) would be 1/996 per hour. [*Tobias and Trindade*, 1986, p. 24]

This intuitive understanding can be formalized, as described in *Tobias and Trindade* [1986, pp. 21–29]. Assume that $F(t)$ gives the probability of a random part drawn from the population failing by time t, where t is the lifetime of the part. (An alternative interpretation of $F(t)$ is that it is the fraction of the population that has failed by time t.) Note that $F(t)$ is thus a cumulative distribution function (CDF) over t (for example, it approaches 1 as t approaches infinity): the corresponding probability distribution function (PDF) is $f(t) = \mathrm{d}F(t)/\mathrm{d}t$. The inverse probability of $F(t)$ is the reliability function $R(t) = 1 - F(t)$. $R(t)$ thus gives the probability that a random part drawn from the population is still operating at (i.e., survives to) time t. (An alternative interpretation of $R(t)$ is that it is the fraction of the population that survives to at least time t.) The failure rate at time t, then, is defined in terms of the conditional probability of failure after surviving up to t. Considering the next Δt interval, this conditional probability is:

$$\mathrm{P}\Big\{[\text{failure in next } \Delta t] \;\Big|\; [\text{survival to } t]\Big\} = \frac{\mathrm{P}\{[\text{failure in next } \Delta t] \text{ and } [\text{survival to } t]\}}{\mathrm{P}\{[\text{survival to } t]\}}$$

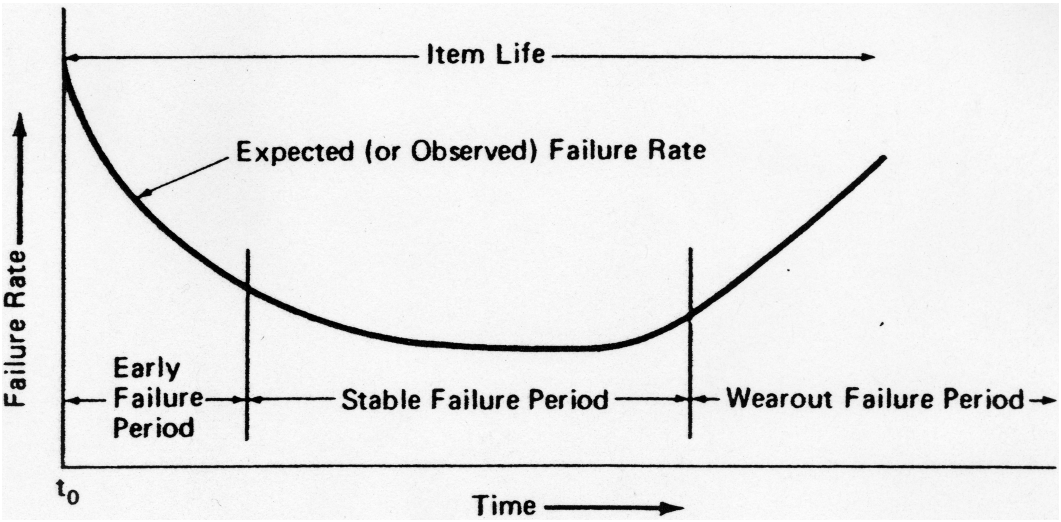

Figure K.1: Illustration of a typical plot of failure rate versus time depicting the characteristic bathtub curve shape. Reproduced in total from *Tobias and Trindade* [1986, p. 28].

leveling out in midlife (labeled as the "Stable Failure Period"). Burn-in testing thus stresses the population to eliminate parts that fail during the early failure period. This stress is typically applied by continuously operating the parts at high temperature, giving rise to the name "burn-in". Parts are characterized both before and after burn-in and the failed parts subsequently removed from the population.

---

where the equality $\mathrm{P}\{A|B\} = \mathrm{P}\{AB\}\,/\mathrm{P}\{B\}$ is simply the definition of conditional probability [*Leon-Garcia*, 1994, p. 210]. Given the established definitions:

$$\mathrm{P}\Big\{[\text{failure in next } \Delta t] \;\Big|\; [\text{survival to } t]\Big\} = \frac{F(t+\Delta t) - F(t)}{R(t)}$$

To convert this conditional probability to an instantaneous rate, divide by the time interval Δt and take the limit as Δt goes to 0:

$$h(t) = \lim_{\Delta t \to 0} \frac{F(t+\Delta t) - F(t)}{R(t)\Delta t} = \frac{f(t)}{R(t)}$$

$h(t)$ is the failure rate. Note that, while it is based on probability, $h(t)$ is not a probability itself; for example, it can take on values greater than 1. Finally, "the reader should be cautioned that not all authors use the same definition when talking about failure rates. Some authors define the failure rate to be $f(t)$, which is the rate of failure of the original time zero population at time t" [*Tobias and Trindade*, 1986, p. 25].

K.2 Device-Under-Test

Burn-in testing of the SVADC-1 involves two batches of devices-under-test (DUTs). The first batch consists of 23 SVADC-1s manufactured in the BiCMOS8iED manufacturing processes, and housed in packages with wider cavities. The second batch consists of 16 SVADC-1s manufactured in the BC8BPLUS manufacturing process, and housed in packages with the same footprint, but smaller cavities. All packages are lidded, with the lids hermetically sealed.

The necessity of the second batch of SVADC-1s was due to unbonding of bondwires in the package cavity of the first batch during vibrational testing. As a result, a package with a smaller cavity was adopted. Unfortunately, by this time the BiCMOS8iED manufacturing process was no longer available. Hence for the second batch the SVADC-1 design was migrated to the BC8BPLUS process, another 0.25-μm BiCMOS process generously provided by National Semiconductor Corporation. Luckily, while BiCMOS8iED and BC8BPLUS differ in their bipolar layers, commercially the CMOS technology of both processes is the same. The second batch SVADC-1 design is thus a direct migration of the first batch SVADC-1 design to the newer process.[2]

K.3 Setup

Burn-in testing requires two experimental setups: one for biasing the DUTs during burn-in, and one for evaluating DUT functionality and performance before and after burn-in.

K.3.1 Burn-In Bias Setup

During the burn-in of both batches, a custom board is used to bias the DUTs. This board was designed and developed by Bob Bumala of Lockheed Martin Corporation with assistance from Charles C. Wang of Stanford University. The schematic for the

[2]The only design change between the two versions of the SVADC-1 design is an adjustment in the RC circuit that triggers the ESD clamps: see footnote 35 of Chapter 6.

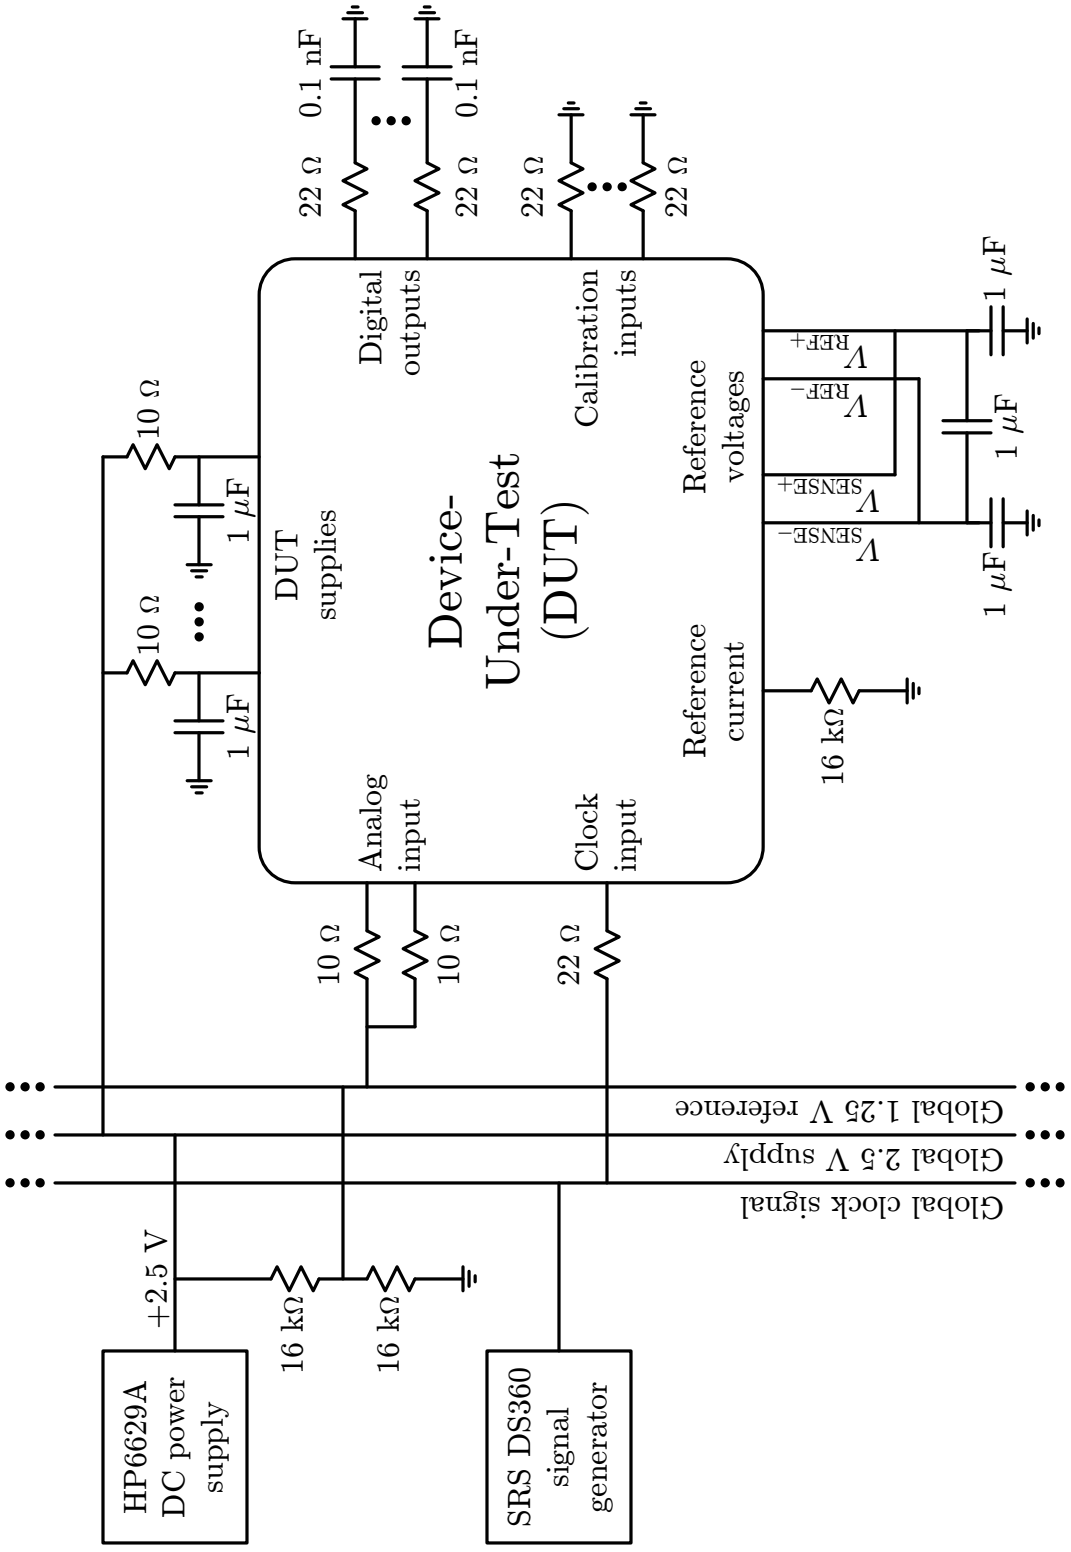

Figure K.2: Schematic of bias board used during burn-in testing of SVADC-1. Board supports up to 28 DUTs simultaneously: circuitry for only a single DUT shown.

board is shown in Figure K.2. The board enables up to 28 DUTs to be simultaneously biased and thus, for space reasons, the circuitry for each DUT is minimized. For example, the DUT reference voltages are provided by the on-chip sources. Also, if possible, circuitry is shared amongst the DUTs, such as in the case of the DUT analog input, power supplies, and clock signal. To maintain isolation between DUTs (important as otherwise a failed DUT can compromise unfailed DUTs), 0.1 W series resistances are imposed on these shared lines so that, should a DUT pin short to a supply or ground, the resistor "pops" and creates an open-circuit.

During burn-in, all DUTs are powered by a single 2.5 V supply provided by an HP6629A power supply, and clocked by a single 100-kHz square wave signal provided by a Stanford Research Systems (SRS) DS360.

K.3.2 Performance Evaluation Setup

The performance of the DUTs is assessed by populating them a DUT-at-a-time into a prototype BBR wideband receiver. In this setting, the DUT is clocked at 100 kHz and uses external voltage and current references. Its input is provided by an SRS DS360 buffered by an operational amplifier buffer circuit configured for a gain of 2. The DUT is thus driven by a fully differential signal whose common mode is set by the buffer.

K.4 Procedure

In burn-in testing, the performance of the DUTs is first measured before burn-in. The DUTs are then placed under bias in a thermal test chamber and heated to undergo the burn-in process. After burn-in, the DUTs are removed from the oven, allowed to cool, and their performance measured again.

K.4.1 Burn-In Procedure

The burn-in procedure is described in Method 1015.9 of the Military Standard [*MIL-STD-883G*, 2006, Method 1015.9]. Briefly, for burn-in, all DUTs are populated in the

burn-in bias board, which is then placed inside the thermal test chamber. The DUTs are biased and the temperature of the thermal test chamber is ramped to 125°C over 30 minutes. Once the thermal test chamber has reached its steady state temperature, the chamber door is opened for a maximum of 10 minutes and the supply and bias voltages of one DUT measured to confirm that the bias board is operational.

The thermal test chamber is then left at 125°C, and the DUTs left operating under bias, for at least 160 hours. Upon completion of burn-in, the chamber door is again opened at temperature and the supply and bias voltages of one DUT measured to confirm bias board operation. Following this confirmation, the thermal test chamber is allowed to cool to room temperature. Note that the DUTs remain biased during cooling. After cooling, the DUTs are removed from the thermal test chamber, their biases are removed, and the DUTs proceed to performance characterization.

K.4.2 Performance Characterization

For the SVADC-1s, DUT performance is assessed by sinusoidal testing. Notably, the Military Standard dictates that, following burn-in, the performance characterization of all DUTs must be completed within 96 hours [*MIL-STD-883G*, 2006, Method 1015.9].[3] Hence each DUT performance characterization must be done relatively quickly.[4]

A list of the datasets collected and the measurements made follows. In these descriptions, the fully differential input signal amplitude is given at the signal generator: recall that there is a gain-of-2 buffer between the signal generator and the DUT.

- *Power dataset*

  The input signal is disconnected, resulting in a ∼0 V DC signal input to the DUT via the buffer. The power of each of the five DUT supplies is measured.

[3] Otherwise, the DUTs must be re-burned-in before additional burn-in performance measurements can be made [*MIL-STD-883G*, 2006, Method 1015.9].

[4] For the burn-in experiment described here, DUTs of both the SVADC-1 and SVLNAE-3 designs were simultaneously burn-in tested, further limiting the characterization time per DUT.

- *Gain-bandwidth dataset*
  With the input signal amplitude set to 0.35 V_{RMS}, the input signal frequency is logarithmically swept from 30 Hz to 50 kHz. The gain at the output of the DUT is then measured as a function of frequency and the 3-dB bandwidth determined.

- *Linearity dataset*
  The input signal amplitude is linearly swept from 17.5 mV_{RMS} to 0.37 V_{RMS} at set frequencies of 100 Hz, 1 kHz, and 10 kHz. The gain and SFDR at the output of the DUT is assessed at each frequency. To determine the gain, a least-squares linear fit is made to the input-amplitude to output-amplitude curve: the slope of the fit is the gain. To determine the SFDR, a 4096-point FFT (for a spectral resolution of ∼24 Hz/bin at 100 kS/s) of the DUT output is taken at each input amplitude and the peak SFDR determined.[5]

- *SINAD and missing codes dataset*
  The input signal frequency is set to 7 kHz and the amplitude to 0.35 V_{RMS}. The SINAD is assessed via the techniques described in footnote 22 of Chapter 3.[6] In addition, the missing output codes are catalogued.

K.5 Results

The measured performance of the first batch of DUTs, before and after burn-in, is summarized in Tables K.1 and K.2, respectively; the same for the second batch is summarized in Tables K.3 and K.4, respectively. Specifically, for each DUT the tables list the total power consumption, the gain and peak SFDR as measured by

[5] In particular, the DUT output is windowed by the 4-term cosine window described in Equations (6.1) and (6.2), with the window length set equal to the FFT length. The SFDR is then determined by first finding the peak of the spectrum: this peak determines the fundamental. From there, a search is made of the highest peak of the spectrum over a bandwidth starting from 10 bins higher than the fundamental peak (∼240 Hz) to $f_S/2$. The ratio of the two peaks is taken as the SFDR.

[6] In particular, 100,000 samples of the DUT output $y'[n]$ are fitted to a signal sinusoid $s[n]$ via a nonlinear least-squares method. This fit is ultimately a four-parameter fit that determines the amplitude, phase, offset, and frequency of $s[n]$. The noise-and-distortion component $d[n]$ is then $d[n] = y'[n] - s[n]$.

the linearity dataset, and the SINAD as measured by the SINAD and missing codes dataset. The tables also include calculations of the aggregate gain μGain and the aggregate SFDR μSFDR. These aggregate quantities collapse the multiple frequency measurements of gain and peak SFDR into a single metric by simple averaging, that is:

$$\mu\text{SFDR} = \frac{1}{3}\left(\left[\text{Peak SFDR}\right]\big|_{100\text{ Hz}} + \left[\text{Peak SFDR}\right]\big|_{1\text{ kHz}} + \left[\text{Peak SFDR}\right]\big|_{10\text{ kHz}}\right) \quad \text{(K.1)}$$

and:

$$\mu\text{Gain} = \frac{1}{3}\left(\text{Gain}\big|_{100\text{ Hz}} + \text{Gain}\big|_{1\text{ kHz}} + \text{Gain}\big|_{10\text{ kHz}}\right) \quad \text{(K.2)}$$

Different weightings of the individual measurements to weigh different frequency bandwidths are possible, but an egalitarian approach is taken here.

K.6 Analysis

Notably, all tested DUTs from both batches remain functional after burn-in testing.

In addition, for part selection for the flown BBR instrument, the DUTs are ranked by a scheme custom-designed for the requirements of the WIPER experiment. For this ranking, both the performance of the DUTs, as well as the change in performance due to burn-in, needs to be considered. Specifically, the ideal DUT is one which not only displays high performance, but also maintains that high performance, with little change, through burn-in. Hence a DUT that achieves high performance, but only after experiencing a large change in performance through burn-in, is not desirable. Nor is a DUT that experiences little change in performance through burn-in, but displays low performance. To capture both requirements, a figure of merit FoM_A is defined for the performance metric A as:

$$\text{FoM}_\text{A} = \text{A}_\text{after} - \left(\text{A}_\text{after} - \text{A}_\text{before}\right)^2 \quad \text{(K.3)}$$

where A_before and A_after are the values of A before and after burn-in, respectively. This figure of merit thus penalizes the after burn-in performance by the square of

SN	Power	Gain				SINAD	Peak SFDR			
		100 Hz	1 kHz	10 kHz	μGain		100 Hz	1 kHz	10 kHz	μSFDR
	[mW]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]
001	39.9	5.87	5.97	5.96	5.93	57.5	77.6	67.4	65.2	70.0
002	40.5	5.80	6.02	6.00	5.94	59.3	77.8	67.2	64.8	70.0
003	42.3	5.81	6.02	6.01	5.95	59.0	77.4	67.5	65.9	70.3
004	42.0	5.91	6.01	6.00	5.97	62.2	77.6	72.1	70.8	73.5
005	42.0	5.75	5.99	5.98	5.91	61.2	76.1	71.3	68.9	72.1
006	41.6	5.83	5.99	5.98	5.93	59.3	76.7	67.5	65.4	69.9
007	42.3	5.87	6.01	6.00	5.96	57.9	75.4	66.5	65.7	69.2
008	42.1	5.70	5.97	5.96	5.88	59.4	76.5	66.6	65.9	69.7
009	41.4	5.77	6.00	5.99	5.92	58.5	76.4	66.4	64.8	69.2
010	40.7	5.89	6.02	6.01	5.97	59.5	77.1	68.6	67.5	71.1
013	41.5	5.68	5.96	5.95	5.86	58.4	77.8	65.9	66.8	70.2
014	44.4	5.82	6.03	6.02	5.96	61.4	76.5	71.3	70.2	72.7
015	40.0	5.82	5.98	5.97	5.92	57.0	76.7	66.8	63.8	69.1
016	40.5	5.72	5.93	5.92	5.86	57.8	74.7	65.2	66.8	68.9
018	42.4	5.71	5.97	5.96	5.88	59.5	75.6	67.6	66.7	70.0
019	40.5	5.84	5.99	5.98	5.94	56.4	75.3	65.4	63.7	68.1
020	42.2	5.70	5.97	5.96	5.88	58.4	75.4	65.5	66.3	69.1
021	43.3	5.75	5.97	5.96	5.89	62.5	76.7	74.9	74.7	75.5
022	40.5	6.02	6.04	6.03	6.03	59.6	76.2	68.6	65.0	69.9
024	41.7	5.70	5.95	5.94	5.86	58.9	76.2	67.0	64.8	69.3
025	42.0	5.92	5.99	5.98	5.96	54.1	72.7	66.0	65.4	68.0
026	42.2	5.81	5.97	5.96	5.91	59.6	72.6	66.9	66.1	68.5
027	42.0	5.75	5.99	5.98	5.91	57.9	72.9	66.3	65.4	68.2
μ	41.7	5.80	5.99	5.98	5.92	58.9	76.0	67.8	66.6	70.1
σ	1.0	0.08	0.03	0.03	0.04	1.8	1.5	2.4	2.5	1.8

Table K.1: Measured DUT performance *before* burn-in for the first batch of SVADC-1s, ordered by DUT serial number SN. Mean (μ) and standard deviation (σ) included for each metric.

SN	Power	Gain				SINAD	Peak SFDR			
		100 Hz	1 kHz	10 kHz	μGain		100 Hz	1 kHz	10 kHz	μSFDR
	[mW]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]
001	39.0	5.84	6.01	6.00	5.95	58.1	78.9	68.5	66.1	71.2
002	40.6	5.87	6.06	6.04	5.99	59.0	80.4	67.8	68.9	72.3
003	42.2	5.91	6.06	6.05	6.01	50.5	79.0	68.8	65.4	71.1
004	41.9	5.95	6.05	6.03	6.01	62.6	80.3	72.6	74.7	75.9
005	41.9	5.89	6.03	6.02	5.98	61.3	79.7	71.2	69.0	73.3
006	41.6	5.83	6.03	6.02	5.96	59.3	79.7	67.7	68.4	71.9
007	42.2	5.93	6.05	6.04	6.01	57.5	78.3	67.6	64.7	70.2
008	41.9	5.90	6.01	6.00	5.97	59.6	80.2	67.7	71.3	73.0
009	41.4	5.98	6.05	6.04	6.02	58.6	79.9	66.9	64.4	70.4
010	40.6	5.97	6.06	6.05	6.03	59.9	80.4	68.4	67.8	72.2
013	41.3	5.70	6.00	6.00	5.90	58.4	78.3	67.6	66.0	70.6
014	44.3	5.95	6.07	6.36	6.13	61.8	78.4	71.6	73.5	74.5
015	40.0	5.99	6.02	6.01	6.01	57.2	77.9	67.3	64.2	69.8
016	40.5	5.79	5.97	5.96	5.91	57.3	78.2	69.9	66.8	71.6
018	42.4	5.96	6.02	6.01	6.00	59.4	79.6	68.1	65.5	71.0
019	40.4	5.98	6.04	6.03	6.02	56.6	77.0	67.6	62.6	69.1
020	42.0	5.85	6.01	6.00	5.95	56.4	79.7	69.5	67.6	72.3
021	43.1	5.80	6.01	6.00	5.94	62.7	80.2	74.3	76.8	77.1
022	40.4	5.94	6.08	6.07	6.03	59.9	79.4	69.2	71.0	73.2
024	41.7	5.84	5.99	5.98	5.94	59.3	81.3	70.0	68.1	73.1
025	41.6	5.83	6.04	6.03	5.97	54.6	80.3	69.2	70.3	73.3
026	42.1	5.79	6.01	6.59	6.13	60.5	81.0	68.5	71.4	73.6
027	41.9	5.79	5.05	6.01	5.62	58.4	78.5	68.3	66.3	71.0
μ	41.5	5.88	5.99	6.06	5.98	58.7	79.4	69.1	68.3	72.3
σ	1.1	0.08	0.20	0.14	0.09	2.6	1.1	1.8	3.5	1.9

Table K.2: Measured DUT performance *after* burn-in for the first batch of SVADC-1s, ordered by DUT serial number SN. Mean (μ) and standard deviation (σ) included for each metric.

SN	Power	Gain				SINAD	Peak SFDR			
		100 Hz	1 kHz	10 kHz	μGain		100 Hz	1 kHz	10 kHz	μSFDR
	[mW]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]
001	42.6	5.85	6.03	6.02	5.97	59.9	77.4	75.8	73.5	75.6
002	42.2	5.95	6.03	6.02	6.00	59.6	75.5	73.8	73.9	74.4
003	42.6	5.92	6.04	6.02	5.99	58.5	75.5	75.0	70.8	73.8
004	40.8	5.91	6.04	6.03	5.99	60.7	75.1	76.5	72.1	74.6
005	42.6	5.90	6.04	6.02	5.99	59.8	75.3	76.8	72.8	75.0
006	42.3	5.87	6.04	6.03	5.98	60.2	76.2	78.2	75.7	76.7
007	42.3	5.79	5.89	6.05	5.91	59.6	74.0	73.9	78.7	75.5
008	42.5	5.91	6.04	6.02	5.99	59.7	75.7	74.2	72.1	74.0
009	42.6	5.78	6.04	6.02	5.95	60.1	76.3	76.3	74.5	75.7
010	42.5	5.86	6.02	6.04	5.97	59.9	75.2	74.5	75.2	75.0
012	42.8	5.88	6.04	6.03	5.98	60.7	75.7	74.7	72.4	74.3
013	41.6	5.88	6.04	6.08	6.00	60.2	76.0	76.5	77.3	76.6
014	41.8	5.95	6.04	6.02	6.00	59.9	75.4	75.8	72.6	74.6
015	42.7	5.77	6.04	6.04	5.95	59.6	75.5	74.7	75.5	75.2
016	41.8	5.92	6.04	6.03	6.00	59.7	75.0	75.0	76.6	75.6
μ	42.2	5.88	6.03	6.03	5.98	59.9	75.6	75.4	74.2	75.1
σ	0.5	0.06	0.04	0.02	0.02	0.5	0.7	1.2	2.1	0.8

Table K.3: Measured DUT performance *before* burn-in for the second batch of SVADC-1s, ordered by DUT serial number SN. Mean (μ) and standard deviation (σ) included for each metric.

SN	Power	Gain				SINAD	Peak SFDR			
		100 Hz	1 kHz	10 kHz	μGain		100 Hz	1 kHz	10 kHz	μSFDR
	[mW]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]	[dB]
001	42.7	5.85	6.03	6.02	5.97	59.9	76.5	74.6	75.4	75.5
002	42.2	5.95	6.03	6.02	6.00	59.7	76.8	73.2	69.4	73.1
003	42.3	5.87	6.03	6.02	5.97	60.0	76.4	76.3	71.8	74.9
004	40.9	5.96	6.04	6.03	6.01	60.8	75.3	75.0	71.5	73.9
005	42.6	5.87	6.03	6.01	5.97	59.8	74.2	77.0	70.9	74.0
006	42.2	5.89	6.03	6.02	5.98	57.7	76.5	77.4	75.5	76.5
007	42.4	5.78	6.03	6.02	5.94	59.6	73.5	74.0	75.5	74.3
008	42.5	5.95	6.03	6.02	6.00	59.7	74.1	73.8	74.4	74.1
009	42.6	5.73	6.03	6.02	5.93	60.1	77.9	74.2	78.4	76.8
010	42.5	5.77	6.04	6.02	5.94	59.9	75.4	75.4	74.9	75.2
012	42.8	5.90	6.03	6.02	5.98	60.7	75.9	75.9	71.5	74.4
013	41.5	5.89	6.03	6.02	5.98	60.2	77.4	76.0	74.7	76.0
014	41.6	5.84	6.03	6.02	5.96	55.5	75.0	75.1	72.3	74.2
015	42.7	5.82	6.03	6.02	5.96	59.7	74.7	75.3	71.9	74.0
016	41.8	5.89	6.03	6.02	5.98	59.7	76.0	75.2	72.7	74.6
μ	42.2	5.86	6.03	6.02	5.97	59.5	75.7	75.2	73.4	74.8
σ	0.5	0.07	0.00	0.00	0.02	1.3	1.2	1.1	2.3	1.0

Table K.4: Measured DUT performance *after* burn-in for the second batch of SVADC-1s, ordered by DUT serial number SN. Mean (μ) and standard deviation (σ) included for each metric.

the change in the performance.[7] The after burn-in performance is chosen as this performance is closest to that expected of the DUT during flight.[8]

The DUT rank is then determined by a weighted sum of figures of merit over select performance metrics. In particular, for the BBR instrument, the figures of merit for μSFDR and SINAD are used and combined as per:

$$\text{Rank} = 0.6\,\text{FoM}_{\mu\text{SFDR}} + 0.4\,\text{FoM}_{\text{SINAD}} \tag{K.4}$$

Higher ranks are desirable. The coefficients of Equation (K.4) can be trimmed as needed, but the emphasis on SFDR over SINAD—motivated by the strong relationship between SFDR and plasma wave receiver spectrographic performance (see Section 3.2.3)—is reflected in the choice here.[9] The power and μGain of the DUTs are not considered in computing the rank since, for these metrics, 1) there is little change in performance due to burn-in, and 2) there is little spread in performance

[7] It is notable that the FoM_A as defined here does not yield to normal unit analysis: it subtracts a dB-squared quantity from a dB quantity. An alternative that better preserves units is to penalize the after burn-in performance by the absolute value of the difference. However, to enforce a stronger penalty, especially for larger performance changes, the square is nonetheless used here.

[8] As constructed in Equation (K.3), a higher FoM_A is more desirable: this definition thus applies for metrics A wherein a higher value is more desirable. This characteristic is true of the metrics used to rank the DUTs as given in Equation (K.4). However, FoM_A can be modified for metrics A wherein lower A is desirable (such as in power consumption): in these cases, the after burn-in performance is penalized by adding, instead of subtracting, the square of the performance change.

[9] As is seen in Tables K.5 and K.6, for both batches the average value of $\text{FoM}_{\mu\text{SFDR}}$ is larger than that of $\text{FoM}_{\text{SINAD}}$. Hence the μSFDR is weighted even more heavily than the coefficients attest. A more egalitarian approach is to equally weight $\text{FoM}_{\mu\text{SFDR}}$ and $\text{FoM}_{\text{SINAD}}$ so that, on average, each metric composes 50% of the rank. That is, if:

$$\text{Rank} = a\,\text{FoM}_{\mu\text{SFDR}} + b\,\text{FoM}_{\text{SINAD}}$$

then this equal-weighting approach desires:

$$a\,\text{E}\{\text{FoM}_{\mu\text{SFDR}}\} = b\,\text{E}\{\text{FoM}_{\text{SINAD}}\}$$

where $\text{E}\{\cdot\}$ is the expected value. If the additional criterion of:

$$a + b = 1$$

is imposed, then it can be shown that:

$$a = \frac{\text{E}\{\text{FoM}_{\text{SINAD}}\}}{\text{E}\{\text{FoM}_{\mu\text{SFDR}}\} + \text{E}\{\text{FoM}_{\text{SINAD}}\}}$$

between DUTs in both the before and after burn-in characterizations. Thus, from an FoM perspective, these metrics do little to differentiate DUTs one from another.

Based on these definitions, the ranks for the first and second batch of DUTs are given in Tables K.5 and K.6, respectively. The Tables are presented in order of highest to lowest rank.

---

and:

$$b = \frac{\mathrm{E}\{\mathrm{FoM}_{\mu\mathrm{SFDR}}\}}{\mathrm{E}\{\mathrm{FoM}_{\mu\mathrm{SFDR}}\} + \mathrm{E}\{\mathrm{FoM}_{\mathrm{SINAD}}\}}$$

For the first batch of SVADC-1 DUTs, under the equal-weighting approach $a = 0.46$ and $b = 0.54$, while for the second batch of SVADC-1 DUTs, $a = 0.44$ and $b = 0.56$. Given the weightings of Equation (K.4), though, on average the $\mathrm{FoM}_{\mu\mathrm{SFDR}}$ and $\mathrm{FoM}_{\mathrm{SINAD}}$ compose 64% and 36%, respectively, of the rank for the first batch of DUTs, and 66% and 34%, respectively, of the rank for the second batch of DUTs, for roughly a 2-to-1 weighting in both cases.

SN	Before burn-in		After burn-in		Figure of merit		Rank
	μSFDR	SINAD	μSFDR	SINAD	$\text{FoM}_{\mu\text{SFDR}}$	$\text{FoM}_{\text{SINAD}}$	
	[dB]	[dB]	[dB]	[dB]			
021	75.5	62.5	77.1	62.7	74.4	62.7	69.7
005	72.1	61.2	73.3	61.3	71.8	61.3	67.6
014	72.7	61.4	74.5	61.8	71.3	61.6	67.4
004	73.5	62.2	75.9	62.6	70.3	62.4	67.2
010	71.1	59.5	72.2	59.9	70.9	59.8	66.4
018	70.0	59.5	71.0	59.4	69.9	59.4	65.7
013	70.2	58.4	70.6	58.4	70.4	58.4	65.6
001	70.0	57.5	71.2	58.1	69.9	57.8	65.0
009	69.2	58.5	70.4	58.6	69.0	58.6	64.8
007	69.2	57.9	70.2	57.5	69.2	57.3	64.5
015	69.1	57.0	69.8	57.2	69.3	57.2	64.4
006	69.9	59.3	71.9	59.3	67.6	59.3	64.3
019	68.1	56.4	69.1	56.6	68.2	56.6	63.6
002	70.0	59.3	72.3	59.0	66.6	58.9	63.6
022	69.9	59.6	73.2	59.9	62.6	59.8	61.4
016	68.9	57.8	71.6	57.3	64.2	57.2	61.4
027	68.2	57.9	71.0	58.4	63.1	58.2	61.2
008	69.7	59.4	73.0	59.6	61.7	59.5	60.8
024	69.3	58.9	73.1	59.3	58.7	59.1	58.9
020	69.1	58.4	72.3	56.4	62.1	52.7	58.3
026	68.5	59.6	73.6	60.5	47.5	59.6	52.4
025	68.0	54.1	73.3	54.6	45.5	54.3	49.0
003	70.3	59.0	71.1	50.5	70.4	-21.9	33.5
μ	70.1	58.9	72.3	58.7	65.9	55.2	61.6
σ	1.8	1.8	1.9	2.6	7.1	16.6	7.6

Table K.5: Ranking of the first batch of SVADC-1s, ordered by highest to lowest rank. Mean (μ) and standard deviation (σ) included for each metric.

SN	Before burn-in		After burn-in		Figure of merit		Rank
	μSFDR	SINAD	μSFDR	SINAD	$\text{FoM}_{\mu\text{SFDR}}$	$\text{FoM}_{\text{SINAD}}$	
	[dB]	[dB]	[dB]	[dB]			
013	76.6	60.2	76.0	60.2	75.7	60.2	69.5
009	75.7	60.1	76.8	60.1	75.5	60.1	69.4
001	75.6	59.9	75.5	59.9	75.5	59.9	69.3
010	75.0	59.9	75.2	59.9	75.2	59.9	69.0
012	74.3	60.7	74.4	60.7	74.4	60.7	68.9
004	74.6	60.7	73.9	60.8	73.5	60.8	68.4
008	74.0	59.7	74.1	59.7	74.1	59.7	68.3
016	75.6	59.7	74.6	59.7	73.8	59.7	68.1
005	75.0	59.8	74.0	59.8	73.0	59.8	67.7
007	75.5	59.6	74.3	59.6	72.8	59.6	67.5
015	75.2	59.6	74.0	59.7	72.5	59.7	67.4
003	73.8	58.5	74.9	60.0	73.7	57.8	67.3
002	74.4	59.6	73.1	59.7	71.6	59.7	66.8
006	76.7	60.2	76.5	57.7	76.4	51.6	66.5
014	74.6	59.9	74.2	55.5	73.9	36.0	58.8
μ	75.1	59.9	74.8	59.5	74.1	57.7	67.5
σ	0.8	0.5	1.0	1.3	1.3	6.2	2.5

Table K.6: Ranking of the second batch of SVADC-1s, ordered by highest to lowest rank. Mean (μ) and standard deviation (σ) included for each metric.

Bibliography

Abel, Bob and Thorne, Richard M. "Electron scattering loss in Earth's inner magnetosphere 1. Dominant physical processes". *Journal of Geophysical Research*, vol. 103, no. A2, pp. 2385–2396, February 1998.

Abo, Andrew Masami. *Design for Reliability of Low-voltage, Switched-capacitor Circuits*. Ph.D. thesis, University of California, Berkeley, Berkeley, CA, Spring 1999.

Actel Corporation. "RTAX-S/SL RadTolerant FPGAs Datasheet". version 5.4, retrieved 21 May 2009.

Adams, R. W.; Ferguson, P. F., Jr.; Ganesan, A.; Vincelette, A.; Volpe, A.; and Libert, R. "Theory and Practical Implementation of a Fifth-Order Sigma-Delta A/D Converter". *Journal of the Audio Engineering Society*, vol. 39, no. 7/8, pp. 515–528, July/August 1991.

Adams, Robert W. "The Design of High-Order Single-Bit $\Delta\Sigma$ ADCs". In *Delta-Sigma Data Converters: Theory, Design, and Simulation*, Steven R. Norsworthy; Richard Schreier; and Gabor C. Temes, eds., chap. 5, pp. 165–192. New York: IEEE Press, 1997. ISBN 0-7803-1045-4.

Adams, Robert W. and Schreier, Richard. "Stability Theory for $\Delta\Sigma$ Modulators". In *Delta-Sigma Data Converters: Theory, Design, and Simulation*, Steven R. Norsworthy; Richard Schreier; and Gabor C. Temes, eds., chap. 4, pp. 141–164. New York: IEEE Press, 1997. ISBN 0-7803-1045-4.

Allenspach, M.; Brews, J. R.; Galloway, K. F.; Johnson, G. H.; Schrimpf, R. D.; Pease, R. L.; Titus, J. L.; and Wheatley, C. F. "SEGR: A Unique Failure Mode for Power MOSFETs in Spacecraft". *Microelectronics and Reliability*, vol. 36, no. 11-12, pp. 1871–1874, November-December 1996.

Anelli, G.; Campbell, M.; Delmastro, M.; Faccio, F.; Floria, S.; Giraldo, A.; Heijne, E.; Jarron, P.; Kloukinas, K.; Marchioro, A.; Moreira, P.; and Snoeys, W. "Radiation Tolerant VLSI Circuits in Standard Deep Submicron CMOS Technologies for the LHC Experiments: Practical Design Aspects". *IEEE Transactions on Nuclear Science*, vol. 46, no. 6, pp. 1690–1696, December 1999.

Anelli, Giovanni Maria. *Conception et Caracterisation de Circuits Integres Resistants aux Radiations pour les Detecteurs de Particules du LHC en Technologies CMOS Submicroniques Profondes.* Ph.D. thesis, Institut National Polytechnique de Grenoble, 11 December 2000.

Aparicio, Roberto and Hajimiri, Ali. "Capacity Limits and Matching Properties of Integrated Capacitors". *IEEE Journal of Solid-State Circuits*, vol. 37, no. 3, pp. 384–393, March 2002.

Attix, Frank Herbert. *Introduction to Radiological Physics and Radiation Dosimetry.* New York: John Wiley & Sons, 1986.

Bell, T. F. and Helliwell, R. A. "The Stanford University VLF Wave Injection Experiment on the ISEE-A Spacecraft". *IEEE Transactions on Geoscience Electronics*, vol. 16, no. 3, pp. 248–252, July 1978.

Blair, J. "Histogram Measurement of ADC Nonlinearities Using Sine Waves". *IEEE Transactions on Instrumentation and Measurement*, vol. 43, no. 3, pp. 373–383, June 1994.

Bracewell, Ronald N. *The Fourier Transform and Its Applications.* San Francisco: McGraw-Hill, 1986. ISBN 0-07-007015-6.

Brandt, Brian; Ferguson, Paul F.; and Rebeschini, Mike. "Analog Circuit Design for $\Delta\Sigma$ ADCs". In *Delta-Sigma Data Converters: Theory, Design, and Simulation*, Steven R. Norsworthy; Richard Schreier; and Gabor C. Temes, eds., chap. 11, pp. 333–379. New York: IEEE Press, 1997. ISBN 0-7803-1045-4.

Brandt, Brian P. and Wooley, Bruce A. "A Low-Power, Area-Efficient Digital Filter for Decimation and Interpolation". *IEEE Journal of Solid-State Circuits*, vol. 29, no. 6, pp. 679–687, June 1994.

Buchner, Stephen; Campbell, Arthur B.; Sternberg, Andrew; Massengill, Lloyd; McMorrow, Dale; and Dyer, Clive. "Validity of Using a Fixed Analog Input for Evaluating the SEU Sensitivity of a Flash Analog-to-Digital Converter". *IEEE Transactions on Nuclear Science*, vol. 52, no. 1, pp. 462–467, February 2005.

Buchner, Stephen; McMorrow, Dale; Poivey, Christian; Howard, James, Jr.; Boulghassoul, Younes; Massengill, Lloyd W.; Pease, Ron; and Savage, Mark. "Comparison of Single-Event Transients Induced in an Operational Amplifier (LM124) by Pulsed Laser Light and a Broad Beam of Heavy Ions". *IEEE Transactions on Nuclear Science*, vol. 51, no. 5, pp. 2776–2781, October. 2004.

Calin, T.; Nicolaidis, M.; and Velazco, R. "Upset Hardened Memory Design for Submicron CMOS Technology". *IEEE Transactions on Nuclear Science*, vol. 43, no. 6, pp. 2874–2878, December 1996.

Carley, Richard L.; Schreier, Richard; and Temes, Gabor C. "Delta-Sigma ADCs with Multibit Internal Converters". In *Delta-Sigma Data Converters: Theory, Design, and Simulation*, Steven R. Norsworthy; Richard Schreier; and Gabor C. Temes, eds., chap. 8, pp. 244–281. New York: IEEE Press, 1997. ISBN 0-7803-1045-4.

Ceschia, M.; Paccagnella, A.; Cester, A.; Scarpa, A.; and Ghidini, G. "Radiation Induced Leakage Current and Stress Induced Leakage Current in Ultra-Thin Gate Oxides". *IEEE Transactions on Nuclear Science*, vol. 45, no. 6, pp. 2375–2382, December 1998.

Chen, Hsin-Shu; Song, Bang-Sup; and Bacrania, Kantilal. "A 14-b 20-MSamples/s CMOS Pipelined ADC". *IEEE Journal of Solid-State Circuits*, vol. 36, no. 6, pp. 997–1001, June 2001.

Chiu, Yun; Gray, Paul R.; and Nikolić, Borivoje. "A 14-b 12-MS/s CMOS Pipeline ADC With Over 100-dB SFDR". *IEEE Journal of Solid-State Circuits*, vol. 39, no. 12, pp. 2139–2151, December 2004.

Cho, Thomas Byunghak and Gray, Paul R. "A 10 b, 20 Msample/s, 35 mW Pipeline A/D Converter". *IEEE Journal of Solid-State Circuits*, vol. 30, no. 3, pp. 166–172, March 1995.

Chuang, Shang-Yuan (Sean) and Sculley, Terry L. "A Digitally Self-Calibrating 14-bit 10-MHz CMOS Pipelined A/D Converter". *IEEE Journal of Solid-State Circuits*, vol. 37, no. 6, pp. 674–683, June 2002.

Chun, Jung-Hoon and Murmann, Boris. "Analysis and Measurement of Signal Distortion due to ESD Protection Circuits". *IEEE Journal of Solid-State Circuits*, vol. 41, no. 10, pp. 2354–2358, October 2006.

Cline, David W. and Gray, Paul R. "A Power Optimized 13-b 5 Msamples/s Pipelined Analog-to-Digital Converter in 1.2 μm CMOS". *IEEE Journal of Solid-State Circuits*, vol. 31, no. 3, pp. 294–303, March 1996.

Cline, David William. *Noise, Speed, and Power Trade-offs in Pipelined Analog to Digital Converters*. Ph.D. thesis, University of California, Berkeley, Berkeley, CA, 1995.

Coilcraft, Incorporated. "P7LP-604L Datasheet: Low Pass LC Filter Modules", 30 October 2008.

Cover, Thomas M. and Thomas, Joy A. *Elements of Information Theory*. New York: John Wiley & Sons, 1991. ISBN 0-471-06259-6.

Cutler, C. C. "Transmission Systems Employing Quantization". U. S. Patent 2,927,962, issued 08 March 1960, filed 26 April 1954.

de Jager, F. "Deltamodulation, A Method of P.C.M. Transmission Using the 1-Unit Code". *Philips Research Reports*, vol. 7, no. 6, pp. 442–466, December 1952.

Devarajan, Siddharth; Singer, Larry; Kelly, Dan; Decker, Steven; Kamath, Abhishek; and Wilkins, Paul. "A 16b 125MS/s 385mW 78.7dB SNR CMOS Pipeline ADC". In *Digest of Technical Papers, 2009 IEEE International Solid-State Circuits Conference (ISSCC 2009)*, pp. 86–87. San Francisco, 8–12 February 2009.

Dodd, Paul E. and Massengill, Lloyd W. "Basic Mechanisms and Modeling of Single-Event Upset in Digital Microelectronics". *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 583–602, June 2003.

Dressendorfer, Paul V. "Radiation Effects on MOS Devices and Circuits". In *Ionizing Radiation Effects in MOS Devices and Circuits*, T. P. Ma and Paul V. Dressendorfer, eds., chap. 4, pp. 256–332. New York: John Wiley & Sons, 1989a. ISBN 0-471-84893-X.

Dressendorfer, Paul V. "Radiation-Hardening Technology". In *Ionizing Radiation Effects in MOS Devices and Circuits*, T. P. Ma and Paul V. Dressendorfer, eds., chap. 2, pp. 333–400. New York: John Wiley & Sons, 1989b. ISBN 0-471-84893-X.

Edwards, Christopher F.; Redman-White, William; Bracey, Mark; Tenbroek, Bernard M.; Lee, Michael S. L.; and Uren, Michael J. "A Multibit $\Sigma\Delta$ Modulator in Floating-Body SOS/SOI CMOS for Extreme Radiation Environments". *IEEE Journal of Solid-State Circuits*, vol. 34, no. 7, pp. 937–948, July 1999.

EIA/JESD57. "Test Procedures for the Measurement of Single-Event Effects in Semiconductor Devices from Heavy Ion Irradiation". EIA/JEDEC Standard, December 1996.

Ergun, R. E.; Carlson, C. W.; Mozer, F. S.; Delory, G. T.; Temerin, M.; McFadden, J. P.; Pankow, D.; Abiad, R.; Harvey, P.; Wilkes, R.; Primbsch, H.; Elphic, R.; Strangeway, R.; Pfaff, R.; and Cattell, C. A. "The FAST Satellite Fields Instrument". *Space Science Reviews*, vol. 98, no. 1-2, pp. 67–91, August 2001.

Evans, Robley D. *The Atomic Nucleus*. New York: McGraw-Hill, 1955.

Galton, Ian. "Digital Cancellation of D/A Converter Noise in Pipelined A/D Converters". *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 47, no. 3, pp. 185–196, March 2000.

George, Jeffery S.; Koga, Rocky; and McMahan, Margaret A. "Neutron Soft Errors in Xilinx FPGAs at Lawrence Berkeley National Laboratory". In *2008 IEEE Radiation Effects Data Workshop*, pp. 118–123. Tucson, AZ, 14–18 July 2008.

Giraldo, A.; Paccagnella, A.; and Minzoni, A. "Aspect ratio calculation in n-channel MOSFETs with a gate-enclosed layout". *Solid State Electronics*, vol. 44, no. 6, pp. 981–989, Jun. 2000.

Goes, João; Vital, João C.; and Franca, José E. "Systematic Design for Optimization of High-Speed Self-Calibrated Pipelined A/D Converters". *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 45, no. 12, pp. 1513–1526, December 1998.

Gołkowski, M.; Inan, U. S.; Gibby, A. R.; and Cohen, M. B. "Magnetospheric amplification and emission triggering by ELF/VLF waves injected by the 3.6 MW HAARP ionospheric heater". *Journal of Geophysical Research (Space Physics)*, vol. 113, pp. 10201–+, October 2008.

Gray, Robert M. "Quantization Noise in $\Sigma\Delta$ A/D Converters". In *Delta-Sigma Data Converters: Theory, Design, and Simulation*, Steven R. Norsworthy; Richard Schreier; and Gabor C. Temes, eds., chap. 2, pp. 44–74. New York: IEEE Press, 1997. ISBN 0-7803-1045-4.

Gregorian, Roubik and Temes, Gabor C. *Analog MOS Integrated Circuits for Signal Processing*. New York: John Wiley & Sons, 1986. ISBN 0-471-09797-7.

Grignoux, Patrice and Geiger, Randall L. "Modeling of MOS Transistors with Nonrectangular-Gate Geometries". *IEEE Transactions on Electron Devices*, vol. 29, no. 8, pp. 1261–1269, August 1982.

Gurnett, D. A.; Persoon, A. M.; Randall, R. F.; Odem, D. L.; Remington, S. L.; Averkamp, T. F.; Debower, M. M.; Hospodarsky, G. B.; Huff, R. L.; Kirchner, D. L.; Mitchell, M. A.; Pham, B. T.; Phillips, J. R.; Schintler, W. J.; Sheyko, P.; and Tomash, D. R. "The Polar Plasma Wave Instrument". *Space Science Reviews*, vol. 71, no. 1-4, pp. 597–622, February 1995.

Gurnett, D. A.; Scarf, F. L.; Fredricks, R. W.; and Smith, E. J. "The ISEE-1 and ISEE-2 Plasma Wave Investigation". *IEEE Transactions on Geoscience Electronics*, vol. 16, no. 3, pp. 225–230, July 1978.

Gurnett, Donald A. "Principles of Space Plasma Wave Instrument Design". In *Measurement Techniques in Space Plasmas: Fields*, Robert F. Pfaff; Joseph E. Borovsky; and David T. Young, eds., pp. 121–136. Washington, DC: American Geophysical Union, 1998. ISBN 0-87590-086-0.

Halliday, David; Resnick, Robert; and Krane, Kenneth S. *Physics, Extended Version*, vol. 2. New York: John Wiley & Sons, Inc., 1992. ISBN 0-471-54804-9.

Händel, Peter; Skoglund, Mikael; and Pettersson, Mikael. "A Calibration Scheme for Imperfect Quantizers". *IEEE Transactions on Instrumentation and Measurement*, vol. 49, no. 5, pp. 1063–1068, October 2000.

Harris, Fredric J. "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform". *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, January 1978.

Harvey, P. R.; Curtis, D. W.; Heetderks, H. D.; Pankow, D.; Rauch-Leiba, J. M.; Wittenbrock, S. K.; and McFadden, J. P. "The FAST Spacecraft Instrument Data Processing Unit". *Space Science Reviews*, vol. 98, no. 1-2, pp. 113–149, August 2001.

Hatanaka, Shingo and Taniguchi, Kenji. "A Ratio-Independent Algorithmic Pipeline Analog-to-Digital Converter". In *The 8th IEEE International Conference on Electronics, Circuits and Systems (ICECS 2001) Conference Proceedings*, vol. 1, pp. 345–348. Malta, 2–5 September 2001. ISBN 0-7803-7057-0.

Helliwell, Robert A. *Whistlers and Related Ionospheric Phenomena*. Stanford, CA: Stanford University Press, 1965.

Higgins, J. R. "Five Short Stories About the Cardinal Series". *Bulletin of the American Mathematical Society*, vol. 12, no. 1, pp. 45–89, 1985.

Holmes-Siedle, Andrew and Adams, Len. *Handbook of Radiation Effects*. 2nd edn. New York: Oxford University Press, 2002. ISBN 0-19-850733-X.

Hughes, H. "Historical Perspective". In *Ionizing Radiation Effects in MOS Devices and Circuits*, T. P. Ma and Paul V. Dressendorfer, eds., chap. 2, pp. 47–86. New York: John Wiley & Sons, 1989. ISBN 0-471-84893-X.

Hughes, H. L. and Benedetto, J. M. "Radiation Effects and Hardening of MOS Technology: Devices and Circuits". *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 500–521, June 2003.

i-Lotus Corporation. "M12M Timing Oncore^{TM} Receiver Brochure". Specifications 1.0, retrieved 01 May 2009.

IEEE Std 1241-2000. "IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters". Waveform Measurement and Analysis Technical Committee of the IEEE Instrumentation and Measurement Society, 7 December 2000.

IEEE Std 488.1-2003. "IEEE Standard for Higher Performance Protocol for the Standard Digital Interface for Programmable Instrumentation". Technical Committee on Automated Test Systems and Instrumentation (TC-8) of the IEEE Instrumentation and Measurement Society, 12 December 2003.

Inan, U. S.; Bell, T. F.; Bortnik, J.; and Albert, J. M. "Controlled precipitation of radiation belt electrons". *Journal of Geophysical Research (Space Physics)*, vol. 108, no. A5, pp. 1186–+, May 2003.

Inan, U. S.; Helliwell, R. A.; and Kurth, W. S. "Terrestrial Versus Jovian VLF Chorus; A Comparative Study". *Journal of Geophysical Research*, vol. 88, no. A8, pp. 6171–6180, August 1983.

Inan, U. S.; Platino, M.; Bell, T. F.; Gurnett, D. A.; and Pickett, J. S. "Cluster measurements of rapidly moving sources of ELF/VLF chorus". *Journal of Geophysical Research (Space Physics)*, vol. 109, pp. 5214–+, May 2004.

Inan, Umran S. and Inan, Aziz S. *Electromagnetic Waves*. New Jersey: Prentice Hall, 2000. ISBN 0-201-36179-5.

Ingino, Joseph M. and Wooley, Bruce A. "A Continuously Calibrated 12-b, 10-MS/s, 3.3-V A/D Converter". *IEEE Journal of Solid-State Circuits*, vol. 33, no. 12, pp. 1920–1931, December 1998.

Inose, H.; Yasuda, Y.; and Murakami, J. "A Telemetering System by Code Modulation — Δ-Σ Modulation". *IRE Transactions on Space Electronics and Telemetry*, vol. SET-8, no. 3, pp. 204–209, September 1962.

Irom, Farokh; Farmanesh, Farhad; and Kouba, Coy K. "Single-Event Upset and Scaling Trends in New Generation of the Commercial SOI PowerPC Microprocessors". *IEEE Transactions on Nuclear Science*, vol. 53, no. 6, pp. 3563–3568, December 2006.

Jantzi, Stephen; Schreier, Richard; and Snelgrove, Martin. "The Design of Bandpass $\Delta\Sigma$ ADCs". In *Delta-Sigma Data Converters: Theory, Design, and Simulation*, Steven R. Norsworthy; Richard Schreier; and Gabor C. Temes, eds., chap. 9, pp. 282–308. New York: IEEE Press, 1997. ISBN 0-7803-1045-4.

Jeon, Young-Deuk; Lee, Seung-Chul; Kwi-Dong, Kim; Kwon, Jong-Kee; and Kim, Jongdae. "A 4.7mW 0.32mm^2 10b 30MS/s Pipelined ADC Without a Front-End S/H in 90nm CMOS". In *Digest of Technical Papers, 2007 IEEE International Solid-State Circuits Conference (ISSCC 2007)*, pp. 456–457, 615. San Francisco, 11–15 February 2007.

Jerri, Abdul J. "The Shannon Sampling Theorem—Its Various Extensions and Applications: A Tutorial Review". *Proceedings of the IEEE*, vol. 65, no. 11, pp. 1565–1596, November 1977.

Jerri, Abdul J. "Correction to "The Shannon Sampling Theorem—Its Various Extensions and Applications: A Tutorial Review"". *Proceedings of the IEEE*, vol. 67, no. 4, pp. 695–695, April 1979.

Johnson, Michael B.; McMahan, Margaret A.; Galloway, Michelle; Leitner, Daniela; Morel, James R.; Gimpel, Thomas L.; Ninemire, Brien F.; Siero, Reba; and Thatcher, Raymond K. ""Super" Cocktails for Heavy Ion Testing". In *2007 IEEE Radiation Effects Data Workshop*, vol. 0, pp. 34–37. Waikiki, HI, 23–27 July 2007.

Johnston, A. H. "The Influence of VLSI Technology Evolution on Radiation-Induced Latchup in Space Systems". *IEEE Transactions on Nuclear Science*, vol. 43, no. 2, pp. 505–521, April 1996.

JPL, Jet Propulsion Laboratory. "NASA/JPL A/D Converter Selection Guide". Tech. rep., Electronic Parts Engineering, Office 514, 14 January 2005.

Karanicolas, Andrew N.; Lee, Hae-Seung; and Bacrania, Kantilal L. "A 15-b 1-Msample/s Digitally Self-Calibrated Pipeline ADC". *IEEE Journal of Solid-State Circuits*, vol. 28, no. 12, pp. 1207–1215, December 1993.

Ker, Ming-Dou. "Whole-Chip ESD Protection Design with Efficient VDD-to-VSS ESD Clamp Circuits for Submicron CMOS VLSI". *IEEE Transactions on Electron Devices*, vol. 46, no. 1, pp. 173–183, January 1999.

Kerns, S. E.; Shafer, B. D.; Rockett, L. R., Jr.; Pridmore, J. S.; Berndt, D. F.; van Vonno, N.; and Barber, F. E. "The Design of Radiation-Hardened ICs for Space: A Compendium of Approaches". *Proceedings of the IEEE*, vol. 76, no. 11, pp. 1470–1509, November 1988.

Kerris, K. G. "Source Considerations and Testing Techniques". In *Ionizing Radiation Effects in MOS Devices and Circuits*, T. P. Ma and Paul V. Dressendorfer, eds., chap. 8, pp. 443–484. New York: John Wiley & Sons, 1989. ISBN 0-471-84893-X.

Kester, Walt and Bryant, James. "ADC Architectures". In *Data Conversion Handbook*, Walt Kester, ed., Analog Devices Series, chap. 3-2. Amsterdam: Newnes, 2005. ISBN 0-7506-7841-0.

King, Everett E.; Lacoe, Ronald C.; and Wang-Ratkovic, Janet. "The Role of the Spacer Oxide in Determining Worst-Case Hot-Carrier Stress Conditions for NMOS LDD Devices". In *2000 IEEE International Reliability Physics Symposium Proceedings*, pp. 83–92, 2000.

Koga, R. "Single-Event Effect Ground Test Issues". *IEEE Transactions on Nuclear Science*, vol. 43, no. 2, pp. 661–670, April 1996.

Kulchycki, Scott D.; Trofin, Roxana; Vleugels, Katelijn; and Wooley, Bruce A. "A 77-dB Dynamic Range, 7.5-MHz Hybrid Continuous-Time/Discrete-Time Cascaded $\Sigma\Delta$ Modulator". *IEEE Journal of Solid-State Circuits*, vol. 43, no. 4, pp. 796–804, April 2008.

Kulchycki, Scott Douglas. *Continuous-Time Σ-Δ Modulation for High-Resolution, Broadband, A/D Conversion.* Ph.D. thesis, Stanford University, Stanford, CA, March 2007.

Kwak, Sung-Ung; Song, Bang-Sup; and Bacrania, Kantilal. "A 15-b, 5-Msample/s Low-Spurious CMOS ADC". *IEEE Journal of Solid-State Circuits*, vol. 32, no. 12, pp. 1866–1875, December 1997.

Lacoe, Ronald. "CMOS Scaling, Design Principles and Hardening-by-Design Methodologies". 2003 IEEE Nuclear and Space Radiation Effects Conference (NSREC) Short Course, Monterey, CA, 21 July 2003.

Lauben, D. S.; Inan, U. S.; Bell, T. F.; and Gurnett, D. A. "Source characteristics of ELF/VLF chorus". *Journal of Geophysical Research (Space Physics)*, vol. 107, pp. 1429–+, December 2002.

Law, Waisiu; Guo, Jianjun; Peach, Charles T.; Helms, Ward J.; and Allstot, David J. "A Monotonic Digital Calibration Technique for Pipelined Data Converters". In *Proceedings of The 2003 IEEE International Symposium on Circuits and Systems (ISCAS 2003)*, vol. 1, pp. I–873–I–876. Bangkok, Thailand, 25–28 May 2003.

LBNL, Lawrence Berkeley National Laboratory. "The 88-Inch Cyclotron, Downloads Page". http://cyclotron.lbl.gov/Download.html (accessed 21 April 2009), 2009a.

LBNL, Lawrence Berkeley National Laboratory. "The 88-Inch Cyclotron, Heavy Ions Page". http://cyclotron.lbl.gov/subpage2.html (accessed 21 April 2009), 2009b.

Lee, Byung Geun and Yan, Shouli. "A New Ratio-Independent A/D Conversion Technique for High-Resolution Pipeline A/D Converters". In *IEEE International Symposium on Circuits and Systems (ISCAS 2005) Conference Proceedings*, vol. 3, pp. 1960–1963. Kobe, Japan, 23–26 May 2005.

Lee, Chin-Chen. "A New Switched-Capacitor Realization for Cyclic Analog-to-Digital Converter". In *IEEE 1983 International Symposium on Circuits and Systems (ISCAS 1983)*, vol. 3, pp. 1261–1265. Newport Beach, CA, 2–4 May 1983.

Lee, Kelin Hsiao-Heng. *Fault-Tolerant FFT Processor Design for Space Applications*. Ph.D. thesis, Stanford University, Stanford, CA, in preparation.

Lee, Seung-Hoon and Song, Bang-Sup. "Digital-Domain Calibration of Multistep Analog-to-Digital Converters". *IEEE Journal of Solid-State Circuits*, vol. 27, no. 12, pp. 1679–1688, December 1992.

Lee, Thomas H. *The Design of CMOS Radio-Frequency Integrated Circuits*. New York: Cambridge University Press, 1998. ISBN 0-521-63922-0.

Leon-Garcia, Alberto. *Probability and Random Processes for Electrical Engineering*. 2nd edn. Menlo Park, CA: Addison-Wesley, 1994. ISBN 0-201-50037-X.

Lewis, Dean; Pouget, Vincent; Beaudoin, Felix; Perdu, Phillipe; Lapuyade, Herve; Fouillat, Pascal; and Touboul, André. "Backside Laser Testing of ICs for SET Sensitivity Evaluation". *IEEE Transactions on Nuclear Science*, vol. 48, no. 6, pp. 2193–2201, December 2001.

Lewis, Stephen H. "Optimizing the Stage Resolution in Pipelined, Multistage, Analog-to-Digital Converters for Video-Rate Applications". *IEEE Transactions*

on Circuits and Systems II: Analog and Digital Signal Processing, vol. 39, no. 8, pp. 516–523, August 1992.

Lewis, Stephen H.; Fetterman, H. Scott; Gross, George F., Jr.; Ramachandran, R.; and Viswanathan, T. R. "A 10-b 20-Msample/s Analog-to-Digital Converter". *IEEE Journal of Solid-State Circuits*, vol. 27, no. 3, pp. 351–358, March 1992.

Lewis, Stephen H. and Gray, Paul R. "A Pipelined 5-Msample/s 9-bit Analog-to-Digital Converter". *IEEE Journal of Solid-State Circuits*, vol. SC-22, no. 6, pp. 954–961, December 1987.

Li, Jipeng and Moon, Un-Ku. "Background Calibration Techniques for Multistage Pipelined ADCs With Digital Redundancy". *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 50, no. 9, pp. 531–538, September 2003.

Li, Ping Wai; Chin, Michael J.; Gray, Paul R.; and Castello, Rinaldo. "A Ratio-Independent Algorithmic Analog-to-Digital Conversion Technique". *IEEE Journal of Solid-State Circuits*, vol. SC-19, no. 6, pp. 828–836, December 1984.

Limotyrakis, Sotirios. *Power-Efficient Broadband A/D Conversion*. Ph.D. thesis, Stanford University, Stanford, CA, September 2004.

Limotyrakis, Sotirios; Kulchycki, Scott D.; Su, David K.; and Wooley, Bruce A. "A 150-MS/s 8-b 71-mW CMOS Time-Interleaved ADC". *IEEE Journal of Solid-State Circuits*, vol. 40, no. 5, pp. 1057–1067, May 2005.

Lin, Yuh-Min; Kim, Beomsup; and Gray, Paul R. "A 13-b 2.5-MHz Self-Calibrated Pipelined A/D Converter in 3-μm CMOS". *IEEE Journal of Solid-State Circuits*, vol. 26, no. 4, pp. 628–636, April 1991.

Linear Technology Corporation. "LT1763 Series Datasheet: 500 mA, Low Noise, LDO Micropower Regulators", 1999.

Ma, T. P. and Dressendorfer, Paul V., eds. *Ionizing Radiation Effects in MOS Devices and Circuits*. New York: John Wiley & Sons, 1989. ISBN 0-471-84893-X.

Manghisoni, Massimo; Ratti, Lodovico; Re, Valerio; and Speziali, Valeria. "Radiation Hardness Perspectives for the Design of Analog Detector Readout Circuits in the 0.18-μm CMOS Generation". *IEEE Transactions on Nuclear Science*, vol. 49, no. 6, pp. 2902–2909, December 2002.

Martin, Donald R. and Secor, Donald J. *High Speed Analog-to-Digital Converters in Communication Systems: Terminology, Architecture, Theory and Performance*. TRW Defense and Space Systems Group, 1981.

Massengill, L. W.; Choi, B. K.; Fleetwood, D. M.; Schrimpf, R. D.; Galloway, K. F.; Shaneyfelt, M. R.; Meisenheimer, T. L.; Dodd, P. E.; Schwank, J. R.; Lee, Y. M.; Johnson, R. S.; and Lucovsky, G. "Heavy-Ion-Induced Breakdown in Ultra-Thin Gate Oxides and High-k Dielectrics". *IEEE Transactions on Nuclear Science*, vol. 48, no. 6, pp. 1904–1912, December 2001.

Mavis, David G. and Eaton, Paul H. "Soft Error Rate Mitigation Techniques for Modern Microcircuits". In *40th Annual Reliability Physics Symposium Proceedings, 2002*, pp. 216–225, 2002.

Mayes, Michael K. and Chin, Sing W. "A 200 mW, 1 Msample/s, 16-b Pipelined A/D Converter with On-Chip 32-b Microcontroller". *IEEE Journal of Solid-State Circuits*, vol. 31, no. 12, pp. 1862–1872, December 1996.

Maynard, Nelson C. "Electric Field Measurements in Moderate to High Density Space Plasmas with Passive Double Probes". In *Measurement Techniques for Space Plasmas: Fields*, Robert F. Pfaff; Joseph E. Borovsky; and David T. Young, eds., pp. 13–27. Washington, DC: American Geophysical Union, 1998. ISBN 0-87590-086-0.

McMahan, M. A. "Radiation Effects Testing at the 88-Inch Cyclotron". In *Fifth European Conference on Radiation and Its Effects on Components and Systems (RADECS), 1999*, pp. 142–147. Abbaye de Fontevraud, France, 13–17 September 1999.

McMahan, M. A.; Leitner, D.; Gimpel, T.; Morel, J.; Ninemire, B.; Siero, R.; Silver, C.; and Thatcher, R. "A 16 MeV/nucleon Cocktail for Heavy Ion Testing". In *2004 IEEE Radiation Effects Data Workshop*, pp. 156–159. Atlanta, GA, 22 July 2004.

McMahan, M. A.; Wozniak, G. J.; Lyneis, C. M.; Bowman, D. R.; Charity, R. J.; Liu, Z. H.; Moretto, L. G.; Kehoe, W. L.; Mignerey, A. C.; and Namboodiri, M. N. "Using a Cyclotron Plus ECR Source for Detector Evaluation and Calibration". *Nuclear Instruments and Methods in Physics Research A*, vol. A253, pp. 1–9, December 1986.

McMorrow, Dale; Buchner, Stephen; Lotshaw, William T.; Melinger, Joseph S.; Maher, Mike; and Savage, Mark W. "Demonstration of Single-Event Effects Induced by Through-Wafer Two-Photon Absorption". *IEEE Transactions on Nuclear Science*, vol. 51, no. 6, pp. 3553–3557, December 2004.

McWhorter, P. J.; Winokur, P. S.; and Pastorek, R. A. "Donor/Acceptor Nature of Radiation-Induced Interface Traps". *IEEE Transactions on Nuclear Science*, vol. 35, no. 6, pp. 1154–1159, December 1988.

Mehr, Iuri and Singer, Larry. "A 55-mW, 10-bit, 40-Msample/s Nyquist-rate CMOS ADC". *IEEE Journal of Solid-State Circuits*, vol. 35, no. 3, pp. 318–325, March 2000.

MIL-STD-883G. "Test Method Standard: Microcircuits". U. S. Department of Defense, 28 February 2006.

Moon, Un-Ku and Song, Bang-Sup. "Background Digital Calibration Techniques for Pipelined ADC's". *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 44, no. 2, pp. 102–109, February 1997.

Moss, Steven C.; LaLumondiere, Stephen D.; Scarpulla, John R.; MacWilliams, Kenneth P.; Crain, William R.; and Koga, Rocky. "Correlation of Picosecond Laser-Induced Latchup and Energetic Particle-Induced Latchup in CMOS Test

Structures". *IEEE Transactions on Nuclear Science*, vol. 42, no. 6, pp. 1948–1956, December 1995.

Mossawir, Benjamin; Linscott, Ivan R.; Inan, Umran S.; Roeder, James L.; Osborn, Jon V.; Witczak, Steven C.; King, Everett E.; and LaLumondiere, Stephen D. "A TID and SEE Radiation-Hardened, Wideband, Low-Noise Amplifier". *IEEE Transactions on Nuclear Science*, vol. 53, no. 6, pp. 3439–3448, December 2006.

Mossawir, Benjamin J. *High-Fidelity, Broadband Amplification and Filtering for Spaceborne Instrumentation.* Ph.D. thesis, Stanford University, Stanford, CA, in preparation.

Murmann, Boris and Boser, Bernhard E. "A 12-bit 75-MS/s Pipelined ADC Using Open-Loop Residue Amplification". *IEEE Journal of Solid-State Circuits*, vol. 38, no. 12, pp. 2040–2050, December 2003.

Nam, KiYoung; Lee, Sang-Min; Su, David K.; and Wooley, Bruce A. "A Low-Voltage Low-Power Sigma-Delta Modulator for Broadband Analog-to-Digital Conversion". *IEEE Journal of Solid-State Circuits*, vol. 40, no. 9, pp. 1855–1864, September 2005.

National Semiconductor Corporation. "Microcircuit Data Sheet: MNLM6172AM-X-RH Dual High Speed, Low Power, Voltage Feedback Amplifiers". Rev. 0A0, 13 August 2001.

Neary, G. J. "Irradiation of simple biological specimens by charged particles". In *The uses of cyclotrons in chemistry, metallurgy and biology: proceedings of a conference held at St. Catherine's College*, pp. 194–203. St. Catherine's College, Oxford, 22–23 September 1969.

NIST 811, NIST Special Publication 811. "Guide for the Use of the International System of Units (SI)". Thompson, Ambler and Taylor, Barry N., National Institute of Standards and Technology (NIST), U. S. Department of Commerce, March 2008.

Normand, Eugene. "Single-Event Effects in Avionics". *IEEE Transactions on Nuclear Science*, vol. 43, no. 2, pp. 461–474, April 1996.

Norsworthy, Steven R. and Crochiere, Ronald E. "Decimation and Interpolation for $\Delta\Sigma$ Conversion". In *Delta-Sigma Data Converters: Theory, Design, and Simulation*, Steven R. Norsworthy; Richard Schreier; and Gabor C. Temes, eds., chap. 13, pp. 406–446. New York: IEEE Press, 1997. ISBN 0-7803-1045-4.

Norsworthy, Steven R.; Schreier, Richard; and Temes, Gabor C., eds. *Delta-Sigma Data Converters: Theory, Design, and Simulation*. New York: IEEE Press, 1997. ISBN 0-7803-1045-4.

Nowlin, R. N.; McEndree, S. R.; Wilson, A. L.; and Alexander, D. R. "A New Total-Dose-Induced Parasitic Effect in Enclosed-Geometry Transistors". *IEEE Transactions on Nuclear Science*, vol. 52, no. 6, pp. 2495–2502, December 2005.

Nuttall, Albert H. "Some Windows with Very Good Sidelobe Behavior". *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, no. 1, pp. 84–91, February 1981.

Nyquist, H. "Certain Topics in Telegraph Transmission Theory". *Transactions of the American Institute of Electrical Engineers (A.I.E.E.)*, vol. 47, no. 2, pp. 617–644, April 1928.

Oldham, T. R. and McLean, F. B. "Total Ionizing Dose Effects in MOS Oxides and Devices". *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 483–499, June 2003.

Oldham, Timothy R. *Ionizing Radiation Effects in MOS Oxides*. New Jersey: World Scientific, 1999.

Oppenheim, Alan V.; Schafer, Ronald W.; and Buck, John R. *Discrete-Time Signal Processing*. 2nd edn. New Jersey: Prentice Hall, 1999. ISBN 0-13-754920-2.

Oppenheim, Alan V.; Willsky, Alan S.; and Nawab, S. Hamid. *Signals & Systems*. 2nd edn. New Jersey: Prentice Hall, 1997. ISBN 0-13-814757-4.

Opris, Ion E.; Lewicki, Laurence D.; and Wong, Bill C. "A Single-Ended 12-bit 20 Msample/s Self-Calibrating Pipeline A/D Converter". *IEEE Journal of Solid-State Circuits*, vol. 33, no. 12, pp. 1898–1903, December 1998.

Ortiz-Conde, A.; García Sánchez, F. J.; Liou, J. J.; Cerdeira, A.; Estrada, M.; and Yue, Y. "A review of recent MOSFET threshold voltage extraction methods". *Microelectronics Reliability*, vol. 42, no. 4–5, pp. 583–596, 2002.

Osborn, J. V.; Mayer, D. C.; Lacoe, R. C.; Moss, S. C.; and LaLumondiere, S. D. "Single Event Latchup Characteristics of Three Commercial CMOS Processes". In *Proceedings of the 7th NASA Symposium on VLSI Design*, pp. 4.2.1–4.2.14. Albuquerque, NM, 1998. Quoted in Ronald Lacoe, "CMOS Scaling, Design Principles and Hardening-by-Design Methodologies". 2003 IEEE Nuclear and Space Radiation Effects Conference (NSREC) Short Course, Monterey, CA, 21 July 2003.

Pan, Hui and Abidi, Asad A. "Spectral Spurs due to Quantization in Nyquist ADCs". *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 51, no. 8, pp. 1422–1439, August 2004.

Papoulis, Athanasios. *Probability, Random Variables, and Stochastic Processes*. 3rd edn. San Francisco: McGraw-Hill, 1991. ISBN 0-07-048477-5.

Pease, R. L. "Total Ionizing Dose Effects in Bipolar Devices and Circuits". *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 539–551, June 2003.

Petersen, E. L.; Pickel, J. C.; Adams, J. H., Jr.; and Smith, E. C. "Rate Prediction for Single Event Effects – a Critique". *IEEE Transactions on Nuclear Science*, vol. 39, no. 6, pp. 1577–1599, December 1992.

Pierret, Robert F. *Field Effect Devices*, vol. 4 of *Modular Series on Solid State Devices*. 2nd edn. Menlo Park, CA: Addison-Wesley, 1990. ISBN 0-201-12298-7.

Platino, M.; Inan, U. S.; Bell, T. F.; Pickett, J.; Kennedy, E. J.; Trotignon, J. G.; Rauch, J. L.; and Canu, P. "Cluster observations of ELF/VLF signals generated

by modulated heating of the lower ionosphere with the HAARP HF transmitter". *Annales Geophysicae*, vol. 22, pp. 2643–2653, July 2004.

Plummer, James D.; Deal, Michael D.; and Griffin, Peter B. *Silicon VLSI Technology: Fundamentals, Practice and Modeling*. New Jersey: Prentice Hall, 2000. ISBN 0-13-085037-3.

Pouget, Vincent; Lewis, Dean; and Fouillat, Pascal. "Time-Resolved Scanning of Integrated Circuits With a Pulsed Laser: Application to Transient Fault Injection in an ADC". *IEEE Transactions on Instrumentation and Measurement*, vol. 53, no. 4, pp. 1227–1231, August 2004.

Preece, W. H. "Earth Currents". *Nature*, vol. 49, p. 554, 12 April 1894.

Rabii, Shahriar. *Design of Low-Voltage Low-Power Sigma-Delta Modulators*. Ph.D. thesis, Stanford University, Stanford, CA, May 1998.

Rasmussen, Robert D. "Spacecraft Electronics Design for Radiation Tolerance". *Proceedings of the IEEE*, vol. 76, no. 11, pp. 1527–1537, November 1988.

Razavi, Behzad. *Design of Analog CMOS Integrated Circuits*. San Francisco: McGraw-Hill, 2001. ISBN 0-07-238032-2.

Redman-White, William; Dunn, Roy; Lucas, Rex; and Smithers, Peter. "A Radiation-Hard AGC Stabilized SOS Crystal Oscillator". *IEEE Journal of Solid-State Circuits*, vol. 25, no. 1, pp. 282–288, February 1990.

Rivetti, Angelo; Anelli, Giovanni; Anghinolfi, Francis; Mazza, Giovanni; and Rotondo, Francesco. "A Low-Power 10-bit ADC in a 0.25-μm CMOS: Design Considerations and Test Results". *IEEE Transactions on Nuclear Science*, vol. 48, no. 4, pp. 1225–1228, August 2001.

Rombouts, P. and Weyten, L. "Comments on "Interstage Gain-Proration Technique for Digital-Domain Multistep ADC Calibration"". *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 46, no. 8, pp. 1114–1116, August 1999.

Schwank, J. R.; Ferlet-Cavrois, V.; Shaneyfelt, M. R.; Paillet, P.; and Dodd, P. E. "Radiation Effects in SOI Technologies". *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 522–538, June 2003.

Schwank, J. R.; Winokur, P. S.; McWhorter, P. J.; Sexton, F. W.; Dressendorfer, P. V.; and Turpin, D. C. "Physical Mechanisms Contributing to Device "Rebound"". *IEEE Transactions on Nuclear Science*, vol. 31, no. 6, pp. 1434–1438, December 1984.

Sedra, Adel S. and Smith, Kenneth C. *Microelectronic Circuits*. New York: Oxford University Press, 1998. ISBN 0-19-511663-1.

Seifert, Norbert; Zhu, Xiaowei; and Massengill, Lloyd W. "Impact of Scaling on Soft-Error Rates in Commercial Microprocessors". *IEEE Transactions on Nuclear Science*, vol. 49, no. 6, pp. 3100–3106, December 2002.

Sexton, F. W.; Fleetwood, D. M.; Shaneyfelt, M. R.; Dodd, P. E.; and Hash, G. L. "Single Event Gate Rupture in Thin Gate Oxides". *IEEE Transactions on Nuclear Science*, vol. 44, no. 6, pp. 2345–2352, December 1997.

Sexton, Fred W. and Schwank, James R. "Correlation of Radiation Effects in Transistors and Integrated Circuits". *IEEE Transactions on Nuclear Science*, vol. 32, no. 6, pp. 3975–3981, December 1985.

Shaneyfelt, M. R.; Dodd, P. E.; Draper, B. L.; and Flores, R. S. "Challenges in Hardening Technologies Using Shallow-Trench Isolation". *IEEE Transactions on Nuclear Science*, vol. 45, no. 6, pp. 2584–2592, December 1998.

Shaneyfelt, Marty R.; Maher, Michael C.; Camilletti, Robert C.; Schwank, James R.; Pease, Ronald L.; Russell, Brian A.; and Dodd, Paul E. "Elimination of Enhanced Low-Dose-Rate Sensitivity in Linear Bipolar Devices Using Silicon-Carbide Passivation". *IEEE Transactions on Nuclear Science*, vol. 53, no. 4, pp. 2027–2032, August 2006.

Shannon, Claude E. "Communication in the Presence of Noise". *Proceedings of the Institute of Radio Engineers (I.R.E.)*, vol. 37, no. 1, pp. 10–21, January 1949.

Shawhan, Stanley D. "Magnetospheric Plasma Waves". In *Solar System Plasma Physics*, Louis J. Lanzerotti; Charles F. Kennel; and E. N. Parker, eds., vol. 3, chap. 3.1.6, pp. 211–270. New York: North-Holland, 1979. ISBN 0-444-85267-0.

Singer, Larry; Ho, Stacy; Timko, Mike; and Kelly, Dan. "A 12b 65 MSample/s CMOS ADC with 82dB SFDR at 120MHz". In *2000 IEEE International Solid-State Circuits Conference (ISSCC 2000) Digest of Technical Papers*, pp. 38–39. San Francisco, 7–9 February 2000.

Siragusa, E. J. and Galton, I. "Gain error correction technique for pipelined analogue-to-digital converters". *Electronics Letters*, vol. 36, no. 7, pp. 617–618, 30 March 2000.

Slepian, D. "On Bandwidth". *Proceedings of the IEEE*, vol. 64, no. 3, pp. 292–300, March 1976.

Smith, Blanchard D., Jr. "An Unusual Electronic Analog-Digital Conversion Method". *IRE Transactions on Instrumentation*, pp. 155–160, June 1956.

Smith, Julius O. *Mathematics of the Discrete Fourier Transform (DFT) with Audio Applications*. 2nd edn. W3K Publishing, May 2008. ISBN 978-0-9745607-4-8.

Snoeys, Walter; Anelli, Giovanni; Campbell, Michael; Cantatore, Eugenio; Faccio, Federico; Heijne, Erik H. M.; Jarron, Pierre; Kloukinas, Kostas C.; Marchioro, Alessandro; Moreira, Paulo; Toifl, Thomas; and Wyllie, Kenneth. "Integrated Circuits for Particle Physics Experiments". *IEEE Journal of Solid-State Circuits*, vol. 35, no. 12, pp. 2018–2030, December 2000.

Song, Bang-Sup; Tompsett, Michael F.; and Lakshmikumar, Kadaba R. "A 12-bit 1-Msample/s Capacitor Error-Averaging Pipelined A/D Converter". *IEEE Journal of Solid-State Circuits*, vol. 23, no. 6, pp. 1324–1333, December 1988.

Spiegel, Murray R. *Schaum's Outline of Theory and Problems of Advanced Mathematics for Engineers and Scientists*. Schaum's Outline Series. San Francisco: McGraw-Hill, 1971. ISBN 07-060216-6.

Srour, Joseph R. and McGarrity, James M. "Radiation Effects on Microelectronics in Space". *Proceedings of the IEEE*, vol. 76, no. 11, pp. 1443–1469, November 1988.

Stassinopoulos, E. G. and Raymond, J. P. "The Space Radiation Environment for Electronics". *Proceedings of the IEEE*, vol. 76, no. 11, pp. 1423–1442, November 1988.

Sze, S. M. *Physics of Semiconductor Devices*. 2nd edn. New York: John Wiley and Sons, 1981. ISBN 0-471-05661-8.

Tascione, Thomas F. *Introduction to the Space Environment*. 2nd edn. Malabar, FL: Krieger, 1994. ISBN 0-89464-044-5.

Teague, Michael J. and Stassinopoulos, E. G. "A Model of the Starfish Flux in the Inner Radiation Zone". NASA/GSFC Report X-601-72-487, NASA, Greenbelt, MA, December 1972.

Thompson, Dwight Ulysses. *A Study of Floating-Point Analog-to-Digital Converters*. Ph.D. thesis, Stanford University, Stanford, CA, March 2001.

Tobias, Paul A. and Trindade, David C. *Applied Reliability*. New York: Van Nostrand Reinhold Company Inc., 1986. ISBN 0-442-28310-5.

Troutman, Ronald R. *Latchup in CMOS Technology: The Problem and Its Cure*. Boston, MA: Kluwer Academic Publishers, 1986. ISBN 0-89838-215-7.

Turflinger, T. L.; Davey, M. V.; and Mappes, B. M. "Single Event Effects in Analog-to-Digital Converters: Device Performance and System Impact". *IEEE Transactions on Nuclear Science*, vol. 41, no. 6, pp. 2187–2194, December 1994.

Turflinger, Thomas L. and Davey, Martin V. "Understanding Single Event Phenomena in Complex Analog and Digital Integrated Circuits". *IEEE Transactions on Nuclear Science*, vol. 37, no. 6, pp. 1832–1838, December 1990.

Vanden Bossche, M.; Schoukens, J.; and Renneboog, J. "Dynamic Testing and Diagnostics of A/D Converters". *IEEE Transactions on Circuits and Systems*, vol. 33, no. 8, pp. 775–785, August 1986.

Vleugels, Katelijn; Rabii, Shahriar; and Wooley, Bruce A. "A 2.5-V Sigma-Delta Modulator for Broadband Communications Applications". *IEEE Journal of Solid-State Circuits*, vol. 36, no. 12, pp. 1887–1899, December 2001.

Voldman, Steven H. *Latchup*. Chichester, West Sussex, England: John Wiley & Sons, 2007. ISBN 978-0470-01642-8.

Walt, Martin. *Introduction to Geomagnetically Trapped Radiation*. New York: Cambridge University Press, 1994. ISBN 0-521-43143-3.

Widrow, Bernard and Kollár, István. *Quantization Noise: Roundoff Error in Digital Communication, Signal Processing, Control, and Communications*. New York: Cambridge University Press, 2008. ISBN 978-0-521-88671-0.

Winokur, Peter S. "Radiation-Induced Interface Traps". In *Ionizing Radiation Effects in MOS Devices and Circuits*, T. P. Ma and Paul V. Dressendorfer, eds., chap. 4, pp. 193–255. New York: John Wiley & Sons, 1989. ISBN 0-471-84893-X.

Witczak, Steven C.; Lacoe, Ronald C.; Osborn, Jon V.; Hutson, John M.; and Moss, Steven C. "Dose-Rate Sensitivity of Modern nMOSFETs". *IEEE Transactions on Nuclear Science*, vol. 52, no. 6, pp. 2602–2608, December 2005.

Ziegler, James F.; Biersack, Jochen P.; and Ziegler, Matthias D. *SRIM: The Stopping and Range of Ions in Matter*. Chester, Maryland: SRIM Co., 2008. ISBN 0-9654207-1-X.