

A High-Speed, Low-Power Analog-to-Digital Converter in Fully Depleted Silicon-on-Insulator Technology

by

Kent H. Lundberg

S.M. Electrical Engineering
Massachusetts Institute of Technology, 1997

S.B. Physics
Massachusetts Institute of Technology, 1992

S.B. Electrical Engineering
Massachusetts Institute of Technology, 1992


Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2002



© Massachusetts Institute of Technology 2002. All rights reserved.

Author

Department of Electrical Engineering and Computer Science
August 2002

Certified by

James K. Roberge
Professor of Electrical Engineering
Thesis Supervisor

Accepted by

Arthur C. Smith
Chairman, Committee on Graduate Students

A High-Speed, Low-Power Analog-to-Digital Converter in Fully Depleted Silicon-on-Insulator Technology

by
Kent H. Lundberg

Submitted to the Department of Electrical Engineering and Computer Science
on August 26, 2002, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis demonstrates a one-volt, high-speed, ultra-low-power, six-bit flash analog-to-digital converter fabricated in a fully depleted silicon-on-insulator CMOS technology.

Silicon-on-insulator CMOS technology provides a number of benefits for low-power low-voltage analog design. The full dielectric isolation of the silicon island, where the transistors are built, allows higher layout packing density and reduces parasitic junction capacitances. Fully depleted silicon-on-insulator (SOI) exhibits improved subthreshold slope, which allows for lower transistor threshold voltages. Significant savings in power consumption can be obtained by leveraging these advantages. However, the floating-body effect can create significant problems in analog circuits, leading to potential circuit malfunction.

A single-ended auto-zeroed comparator topology is optimized to leverage the advantages of fully depleted SOI technology and avoid the floating-body effect. Using this comparator topology and other circuit techniques that operate with a one-volt supply, a six-bit 500-MS/s flash A/D converter is designed with the lowest power-consumption figure of merit in its class. Consuming only 32 mA from a one-volt supply, the quantization energy figure of merit for this design is calculated to be $E_Q = 2$ pJ.

Test chips were fabricated in MIT Lincoln Laboratory's 0.25- μm fully depleted SOI CMOS process. Testing of this design demonstrates the potential of SOI technology for the production of high-speed, low-power analog-to-digital converters.

Thesis Supervisor: James K. Roberge
Title: Professor of Electrical Engineering

Dedication

For Cynthia.

I couldn't have done it without you.

Acknowledgments

I owe a great debt to Professor James Roberge, for being my teacher, supervisor, employer, mentor, role model, and friend. The impact he has had on my development as an engineer is immeasurable. He will forever have my deepest gratitude.

I am also indebted to Dr. Craig Keast, Leader of Group 88 at MIT Lincoln Laboratory, for supporting and funding this work, and for finding help and resources for me when I really needed them.

Special thanks to Tony Soares, Eric Austin, and Anthony Broccio, the Group 88 testing team, for indispensable help driving the equipment, making numerous service calls, and even building ancillary circuits and adapters.

Thank you to Dr. Vyshi Suntharalingam and Dr. Peter Wyatt for many helpful discussions. Vyshi's suggestions and encouragements during the eleventh hour were particularly appreciated.

Thank you to Paul Davis and Hollis Whittington for answering all my layout questions and commiserating in the struggle with Mentor Graphics. I learned more about layout and layout software in an hour of watching over their shoulders than I could have in a thousand hours of reading books and manuals.


To Group Secretary Sue Moriarty, for keeping the paperwork tamed, keeping my badge white, and keeping the candy dish stocked. To George Young and Rich D'Onofrio for keeping the computers running, even when called, at home, on Sunday afternoons.

And to Bob Berger, Jim Burns, Chang-Lee Chen, Chenson Chen, Pascal Gouker, Jeff Knecht, Brian Tyrrell, Donna-Ruth Yost, and the rest of Group 88 for answering a million questions about fabrication details, transistor parameters, and processing variables. Also, a thank you to the MEL fabrication team that actually made my chips, most of whom I know by face but not by name.

I have benefited, both intellectually and professionally, from discussions with, and consulting for, Jennifer Lloyd, Kimo Tam, and Mike Walsh at Analog Devices in Wilmington. I owe them thanks for indirectly supporting this work.

Of course, to my parents.

In Memoriam



Hazel (1993–2002)

Contents

1	Introduction	15
1.1	High-Speed Low-Power A/D Converters	15
1.2	Silicon-on-Insulator Technology	16
1.3	Thesis Contribution and Organization	16
2	Power Consumption of High-Speed A/D Converters	19
2.1	Quantization Energy Figure of Merit	19
2.2	High Speed	22
2.3	Low Power	22
2.4	Absolute Minimum Power	23
2.5	Aperture Jitter	26
2.6	Additional Power Consumption	26
2.6.1	Transistor Matching	27
2.6.2	Device Parasitics	28
2.6.3	Calibration	28
2.7	Future Needs	28
3	Silicon-on-Insulator Technology Overview	29
3.1	Conventional Bulk CMOS Technology	30
3.2	SOI CMOS Technology	31
3.3	Advantages of SOI	32
3.3.1	Full Dielectric Isolation	32
3.3.2	Reduced Capacitances	33
3.3.3	Improved Subthreshold-Slope Factor	33
3.3.4	Short-Channel Effects	35
3.3.5	Harsh Environments	35
3.4	Disadvantages of SOI	36
3.4.1	Floating-Body Effect	36
3.4.2	Kink Effect	37
3.4.3	Self Heating	38
3.4.4	Noise in SOI	38
3.5	Advantages for Analog Circuits	38
3.5.1	Lower Analog Supply Voltage	39
3.5.2	Higher Maximum Gain	39

4	High-Speed Flash Design	41
4.1	Analog Front End	43
4.2	Latch Circuit	43
4.3	Encoding Logic	45
4.4	Output ROM Decoder	47
5	Comparator Design for SOI	49
5.1	Differential Comparators in Bulk and SOI	49
5.2	Single-Ended Auto-Zeroed Comparator	52
5.2.1	Comparator Gain	55
5.2.2	Comparator Speed	58
5.2.3	CMOS Switches	60
5.3	Comparator Specifications	62
5.3.1	Calculated Speed	62
5.3.2	Power Consumption	63
5.4	Comparison to Bulk CMOS Implementation	64
6	Complete Circuit and Simulations	67
6.1	Pseudo-Differential Comparator Topology	67
6.2	Interface Circuitry	68
6.2.1	Clock Phase Production	68
6.2.2	Clock Input Driver	69
6.2.3	50 Ω Output Drivers	70
6.3	Complete Circuit Simulations	72
6.3.1	Functionality	72
6.3.2	Power Consumption	76
6.4	Comparison to Bulk CMOS Implementation	78
7	Experimental Results	79
7.1	Test Structures	83
7.2	High-Speed Testing	83
7.3	Simulations of Observed Problems	87
7.3.1	Switch Resistance	87
7.3.2	Reference Resistors	89
7.4	Other Observations	90
7.5	Quantization Energy	91
8	Conclusions	93
8.1	Summary of Results	93
8.2	Future Directions	94
8.2.1	Second Fabrication	94
8.2.2	Improved Test Structures	94
8.2.3	Improved Design	94
A	Analog-to-Digital Converter Testing	97
A.1	A/D Converter Figures of Merit	99
A.1.1	Static Parameters	99
A.1.2	Resolution and Quantization Noise	102

A.1.3	Frequency-Domain Dynamic Parameters	104
A.1.4	Time-Domain Dynamic Parameters	109
A.2	Dynamic A/D Converter Testing Methods	112
A.2.1	Testing with the Fast Fourier Transform	112
A.2.2	Signal Coherence for FFT Tests	112
A.2.3	Histogram Test for Linearity	115
A.2.4	Sine Wave Curve Fit for ENOB	119
A.2.5	Overall Noise and Aperture Jitter	124
B	Noise Sources in Bulk CMOS	125
B.1	Noise Notation	125
B.2	Flicker Noise	128
B.2.1	The McWhorter Model (Number Fluctuations)	129
B.2.2	The Hooge Model (Mobility Fluctuations)	131
B.2.3	Combined Models	133
B.2.4	Simple Empirical Models	134
B.3	Thermal Noise	135
B.4	Shot Noise	135
B.5	Generation/Recombination Noise	136
B.6	Popcorn Noise	137
B.7	kT/C Noise	139
C	Reported Noise Sources in SOI CMOS	141
C.1	Noise Contribution of the Body Resistance	143
C.2	Shot Noise Induced Excess Low Frequency Noise	146
C.3	Pre-Kink Excess Low Frequency Noise	150
C.4	Generation/Recombination Noise	153
C.5	Comparison of Theories	155
D	Simulation Code	157
D.1	SPICE File asystem.sp	158
D.2	SPICE File xcompare.sp	159
D.3	SPICE File xlatch.sp	160
D.4	SPICE File xbubble.sp	161
D.5	SPICE File xclocks.sp	163
D.6	Program adc.pl	167
D.7	Sample adc.pl Output File	170
E	Agilent 93000 SOC Tester	175
E.1	AC Testing with the Agilent 93000	175
E.2	Agilent Vector Format and Pin Assignments	176
E.3	Agilent Output File Format	178
E.3.1	Agilent JWA to Octave Converter	178
E.3.2	Sample Agilent JWA Output File	179
E.4	Vector File Generator in Perl	180
E.5	Plots of Available Vectors	183
E.5.1	Sine Wave Vector	183
E.5.2	Sawtooth Wave Vector	184

E.5.3	Square Wave Vector	185
E.5.4	White Noise Vector	186
E.5.5	Two-tone Sine Wave Vector	187
E.6	Sample HSPICE Vector File	188
E.7	Sample Agilent Vector File	189

Bibliography	192
---------------------	------------

List of Figures

2.1	Table of nineteen high-speed analog-to-digital converter specifications gathered from the literature, along with commercial parts from Fairchild, Rockwell, and Maxim. . .	21
2.2	Published CMOS flash A/D converters: effective resolution bandwidth versus year of publication. The best-fit line shows a factor of ten increase in ERBW over five years.	22
2.3	All high-speed flash A/D converters: quantization energy versus effective resolution bandwidth.	24
2.4	All high-speed flash A/D converters: quantization energy versus year of publication. The best-fit line shows a factor of ten improvement in quantization energy over nine years.	24
2.5	All high-speed flash A/D converters: effective number of bits versus sampling rate. The line corresponds to the limitation imposed by an rms aperture jitter of 0.4 ps. .	27
3.1	Cross section of a conventional bulk CMOS technology, showing devices separated by field oxide and isolated by an n-well, along with the required contacts to substrate and well.	31
3.2	Cross section of a fully depleted SOI CMOS technology, showing mesa-isolated devices supported by buried oxide layer.	32
3.3	I_D versus V_{DS} curves for a typical NMOS SOI transistor, showing the kink effect at large drain voltages (when V_{DS} is larger than the bandgap of silicon, which is required for impact ionization — this voltage is referred to as the “kink onset voltage”). . . .	37
4.1	Schematic of simplified flash topology.	42
4.2	Ideal analog-to-digital converter characteristic.	43
4.3	Regenerative and doubled half-latch schematic.	44
4.4	Clock phases required for latch operation. The dashed line denotes a complete clock cycle, while the dotted lines show quarter periods.	45
4.5	Schematic of the majority circuit.	45
4.6	Examples of errors in the thermometer code.	47
4.7	Schematic of optimized majority circuit topology with transition point detection. . .	47
4.8	Schematic of an 8-input pseudo-PMOS NOR gate, used in each bit of the output ROM decoder for a 4-bit flash converter.	48
5.1	Schematic of standard continuous-time differential-input comparator topology. . . .	50
5.2	Step response for a differential comparator in SOI technology. The floating-body effect causes significant initial undershoot and long-tail transient as the body charge changes.	51

5.3	Step response for a differential comparator in a comparable bulk CMOS technology. Note absence of any initial error or long-tail transient.	51
5.4	Simplified schematic of single-ended auto-zeroed comparator topology.	52
5.5	Inverter transfer characteristic from input to output. In the auto-zero mode, the inverter self-biases at the intersection point, where the maximum slope of the transfer characteristic provides maximum gain for the next compare phase.	53
5.6	Step response for a single-ended auto-zeroed comparator in SOI technology. Note absence of any initial error or long-tail transient due to offset-cancellation scheme.	54
5.7	Comparator gain versus transistor gate length, from minimum length to two microns.	56
5.8	A close-up of the comparator-gain-versus-transistor-gate-length relationship.	56
5.9	Simplified small-signal equivalent circuit for the single-ended auto-zeroed comparator topology during the auto-zero phase.	58
5.10	Simplified small-signal equivalent circuit for the single-ended auto-zeroed comparator topology during the compare phase.	59
5.11	Complementary CMOS switch, with both NMOS and PMOS transistors to minimize charge dump and clock feedthrough and allow rail-to-rail operation.	60
5.12	Resistance of the switch versus input voltage.	61
5.13	A single comparison accomplished by the single-ended auto-zeroed comparator.	61
5.14	Table of typical SOI transistor and circuit parameters from the designed comparator	62
5.15	A single comparison accomplished by the single-ended auto-zeroed comparator in comparable bulk CMOS technology, showing significantly slower performance in the compare mode.	65
6.1	Pseudo-differential comparator topology.	67
6.2	The three primary clock phases and their purposes. Phase PH12 and phase PH23 are also produced in complementary form as phases $\overline{\text{PH12}}$ and $\overline{\text{PH23}}$	69
6.3	Differential clock input buffer.	70
6.4	Clock phase generator and distribution buffers.	71
6.5	Fully differential current-steering output driver schematic. The 50 Ω resistors shown are the off-chip termination resistors.	71
6.6	Reconstructed output codes of complete circuit with ramp input, from SPICE simulation.	73
6.7	Reconstructed output codes of complete circuit with 2 MHz sine wave input, from SPICE simulation.	73
6.8	Reconstructed output codes of complete circuit with 20 MHz sine wave input, from SPICE simulation.	74
6.9	Reconstructed output codes of complete circuit with 125 MHz sine wave input, from SPICE simulation.	74
6.10	Reconstructed output codes of complete circuit with 250 MHz sine wave input, from SPICE simulation, showing adequate full-power bandwidth.	75
6.11	Integral nonlinearity due to comparator input currents loading the reference resistor string.	77
6.12	Power consumption comparison for flash converter subcircuits in FDSOI and bulk CMOS technologies at 100 MHz (ignoring reference resistors).	78
7.1	Microphotograph of fabricated test chip.	80
7.2	Layout map of test chip.	81
7.3	Bonding diagram and pin-out for packaged test chip.	82

7.4	Raw A/D converter output data from sine-wave test, with best fit sine wave for ENOB calculation.	84
7.5	FFT of A/D converter output for sine-wave test, showing poor SNDR.	85
7.6	ENOB performance versus input frequency for two parts.	85
7.7	Histogram for DC input, showing no input noise.	86
7.8	Histogram for DC input, showing considerable input “noise.”	86
7.9	Measured converter behavior with leaky switches.	88
7.10	Simulated converter behavior with leaky switches.	88
7.11	A single comparison accomplished by the single-ended auto-zeroed comparator, with leaky switches.	89
7.12	INL versus input code for several different reference resistor sizes.	90
A.1	Ideal analog-to-digital converter characteristic.	98
A.2	Analog-to-digital converter characteristic, showing offset and gain error.	100
A.3	Analog-to-digital converter characteristic, showing nonlinearity errors and a missing code. The dashed line is the ideal characteristic, and the dotted line is the best fit.	101
A.4	Quantization error voltage for ideal analog-to-digital converter.	103
A.5	Ideal signal-to-noise ratio due to quantization versus resolution.	104
A.6	Quantization noise floor for an ideal 10-bit A/D converter (4096 point FFT).	105
A.7	A/D converter with significant nonlinearity, showing poor SFDR and THD. Note that high frequency harmonics of the input signal are aliased down to frequencies below the Nyquist frequency.	107
A.8	Two-tone IMD test with second-order nonlinearity, showing (from left) $F_2 - F_1$ product, F_1 input, F_2 input, $2F_1$ product, $F_1 + F_2$ product, and $2F_2$ product.	108
A.9	Effects of aperture jitter.	109
A.10	FFT of non-integer number of cycles (4096 points, 127.5 cycles).	113
A.11	FFT of even divisor number of cycles (4096 points, 128 cycles).	113
A.12	FFT of non-divisor prime number of cycles (4096 points, 127 cycles).	114
A.13	Histogram of 1000 points for 6-bit A/D converter (insufficient).	116
A.14	Histogram of 100,000 points for 6-bit A/D converter.	117
A.15	Partial data plot showing sine wave and added noise, along with best-fit sine wave. ENOB calculated from curve fit is 3.01.	123
A.16	Partial FFT showing sine wave and added noise. ENOB calculated from SNDR is 3.02.	123
B.1	Plot of noise amplitude versus frequency (on linear scales) as measured in two dissimilar vacuum tubes by Johnson in 1925, from [77, Figure 7]. The vertical axis has been normalized by the expected shot noise amplitude.	128
B.2	Flicker noise in ten metal films, from [86, Figure 1]. The solid lines correspond to a value of $\alpha_H = 2 \times 10^{-3}$ for Hooge’s empirical parameter.	132
B.3	Typical popcorn noise, showing discrete levels of channel current modulation due to the trapping and release of a single carrier, for three different bias conditions, from [99, Figure 1].	138
C.1	Noise spectral density for an SOI transistor, showing the Lorentzian hump when the drain voltage is above the kink onset voltage, from [101, Figure 2].	142

C.2	SOI noise spectrum, referred to the gate, in H-gate transistors observed by Faccio et al, showing $1/f$ noise, white (thermal) noise and excess noise components, from [50, Figure 2].	143
C.3	Schematic of the noise model by Faccio et al, showing the noisy body resistance and the filtering capacitors due to C_{ox} and C_{box} , from [50, Figure 6].	144
C.4	Layout of special transistor, showing independent contacts to the body allowing resistance measurements, from [50, Figure 7].	145
C.5	Schematic of noise model by Wei et al, showing low pass filter created by the incremental source-to-body junction resistance and the capacitances to the channel, from [51, Figure 2].	147
C.6	Spectral density of excess noise, showing increased noise for drain voltages above and below the kink onset voltage, from [51, Figure 7].	148
C.7	Spectral density of excess noise for a fully depleted device in full depletion and backside accumulation. The excess noise disappears when the device is fully depleted, from [51, Figure 10].	149
C.8	Output conductance plot, showing AC floating body effect. Inset plot shows derivative of AC conductance. Note how the derivative is peaky at the onset of the AC kink, from [105, Figure 2].	151
C.9	Spectral density of SOS noise, clearly showing multiple Lorentzians with multiple corner frequencies due to generation/recombination noise, from [102, Figure 7]. . . .	154
D.1	Single-ended auto-zeroed comparator.	159
D.2	Regenerative latch and doubled half-latch.	160
D.3	Bubble-error correction majority circuit.	161
D.4	The three primary clock phases and their purposes.	163
D.5	Differential clock input buffer.	164
D.6	Clock phase generator and distribution buffers.	164
E.1	Block diagram of test board interface to Agilent 93000 tester.	175
E.2	Sine wave plot, histogram, and FFT from vector file data.	183
E.3	Sawtooth wave plot and FFT from vector file data.	184
E.4	Square wave plot and FFT from vector file data.	185
E.5	White noise plot, histogram, and FFT from vector file data.	186
E.6	Two-tone sine wave plot, histogram, and FFT from vector file data.	187

Chapter 1

Introduction

Analog-to-digital converters are irreplaceable components in the interface circuitry between the real world and the digital domain. Efforts to increase the speed and decrease the power consumption of analog-to-digital (A/D) converters open up new application fields and make possible the improvement of existing systems.

1.1 High-Speed Low-Power A/D Converters

High-speed A/D converters are used in radar signal receivers, digital sampling oscilloscopes, read channels for magnetic storage systems, and local-area-network interfaces. Although these applications benefit from very high speeds, converter resolution of six bits is often sufficient. The power consumption of A/D converters is a concern in portable electronics, including wireless digital phones and networked laptop computers. While lower power consumption and longer battery life are relentless economic drivers, progress in power consumption reduction has been slow.

The need for high-speed, low-resolution A/D converters is usually satisfied by flash converters fabricated in conventional bulk CMOS technology. However, large parasitic capacitances in bulk CMOS technology often lead to large power consumption for this topology.

1.2 Silicon-on-Insulator Technology

Silicon-on-insulator CMOS technology improves on conventional bulk CMOS in a number of important ways. The full dielectric isolation of the silicon island, where the transistors are built, allows higher layout packing density and reduces parasitic junction capacitances. Fully depleted silicon-on-insulator (SOI) exhibits improved subthreshold slope, which allows for lower transistor threshold voltages. Potentially significant savings in power consumption can be obtained by leveraging these advantages. However, the floating-body effect and other SOI drawbacks can create significant problems in analog circuits, leading to circuit malfunction and negating potential benefits.

1.3 Thesis Contribution and Organization

Using a comparator topology optimized to leverage the advantages of fully depleted SOI technology and avoid the floating-body effect, this thesis details the design of a one-volt, 500-MS/s flash A/D converter with the lowest power-consumption figure of merit in its class. Fabrication and testing of this design demonstrates the potential of SOI technology for the production of high-speed, ultra-low-power analog-to-digital converters.

Chapter 2 compares the speed and power consumption of several high-speed, low-resolution A/D converters reported in the literature. Also, potential limits to performance and power consumption are calculated and discussed. Chapter 3 provides an overview of fully depleted silicon-on-insulator technology, including its advantages and disadvantages for general analog circuit design. Chapter 4 discusses the architecture of the flash A/D converter, as well as the design of some of the digital components and error-correction logic. Chapter 5 details the design and optimization of the single-ended auto-zeroed comparator, which leverages the advantages of SOI technology, avoids the floating-body effect, and operates with a one-volt supply.

Using the circuits from the Chapters 4 and 5, a complete six-bit flash converter is assembled and simulated. Chapter 6 discusses the interface circuitry and the simulations of the complete circuit. Chapter 7 presents the results from test chips fabricated in MIT Lincoln Laboratory's 0.25- μm fully depleted SOI CMOS process.

Chapter 8 summarizes the results and suggests future work on high-speed, low-power flash A/D converters in SOI technology.

Chapter 2

Power Consumption of High-Speed A/D Converters

Every year, higher and higher sampling rates as well as lower and lower power dissipations are reported in the literature. The table in Figure 2.1 lists published results for nineteen high-speed, low-resolution A/D converters (mostly six to eight bits), along with a few commercial parts, representing a wide variety of fabrication technologies. In the following sections, these reported specifications are analyzed and compared.

2.1 Quantization Energy Figure of Merit

The analog program committee of the IEEE International Solid-State Circuits Conference suggested a figure of merit for A/D converters that takes into account power dissipation, resolution, and sampling rate [1]. It has units of energy, and represents the energy used per conversion step

$$E = \frac{P}{2^N F_S}$$

where P is the power dissipation, N is the stated number of bits, and F_S is the sampling rate. This quantity is nearly the inverse of a similar figure of merit suggested earlier by Robert Walden [2, 3] that uses the *effective* number of bits (ENOB), B , instead of the stated number of bits

$$F = \frac{2^B F_S}{P}$$

Walden’s figure of merit F correctly includes the performance limitation of signal-to-noise-and-distortion ratio (SNDR), but still produces optimistic results for some A/D converters¹. In most applications, A/D converters are expected to faithfully convert all input-signal frequencies below the Nyquist frequency (one half of the sampling rate F_S), but many A/D converters exhibit severe degradation of SNDR at frequencies well below the Nyquist frequency. For this reason, the literature has recently started using the effective resolution bandwidth (ERBW) instead of the sampling rate in the equation for the figure of merit [4, 5]. This figure of merit is known as the “quantization energy”

$$E_Q = \frac{P}{2^B (2F_{BW})}$$

where P is the power dissipation, B is the high-frequency ENOB (calculated from SNDR), and F_{BW} is either the effective resolution bandwidth or the Nyquist frequency, whichever is less.

¹Hopeless pedants, like this author [6], also complain about its nonsensical units of “inverse joules.”

Pub Year	F_S MS/s	stated bits	ENOB bits	ERBW MHz	Power mW	E_Q pJ	Pub Ref
CMOS Flash Topologies							
1998	200	6	5.0	100	380	59.4	[7]
1998	350	6	5.8	87	225	23.2	[8]
1998	400	6	5.0	100	190	29.7	[9]
1998	400	6	4.7	100	200	38.5	[10]
1999	500	6	5.3	250	400	20.3	[11]
1999	500	6	4.7	50	330	127	[12]
1999	500	6	5.3	180	225	15.9	[13]
2000	600	6	4.7	150	193	24.8	[14]
2000	700	6	5.5	136	187	15.2	[15]
2000	800	6	5.0	200	400	31.3	[16]
2001	1100	6	5.3	450	300	8.5	[4]
2001	1300	6	5.0	650	545	13.1	[17]
2002	1600	6	5.2	600	328	7.4	[5]
Bipolar Processes (including Si, GaAs, InP, etc)							
1995	2000	8	6.5	1000	5300	29.3	[18]
1995	4000	6	5.1	1800	5700	46.2	[3, 19]
1996	8000	3	2.4	4000	3500	82.9	[20]
1999	500	8	7.2	250	950	12.9	[21]
Interleaved Topologies							
1997	8000	8	5.3	2000	13500	85.7	[22]
2002	4000	8	6.5	900	2400	14.7	[23]
Commercial Parts							
1995	1000	8	5.5	500	5500	122	[24]
1998	2000	8	5.9	500	5000	83.7	[25]
2001	1500	8	7.5	750	5300	19.5	[26]

Figure 2.1: Table of nineteen high-speed analog-to-digital converter specifications gathered from the literature, along with commercial parts from Fairchild, Rockwell, and Maxim.

2.2 High Speed

High-speed analog-to-digital converters are used in radar signal receivers, digital sampling oscilloscopes, read channels for magnetic storage systems, and local-area-network interfaces. Figure 2.2 shows the steady increase in effective resolution bandwidth with time. The average speed of high-speed A/D converters has increased by a factor of ten over the past five years.




Figure 2.2: Published CMOS flash A/D converters: effective resolution bandwidth versus year of publication. The best-fit line shows a factor of ten increase in ERBW over five years.

2.3 Low Power

In this era of wireless communications, portable high-speed test equipment, and networked laptop computers, the motivation for low-power electronics is clear. Reducing the power requirements of existing components extends the battery life of portable devices, or allows additional functionality to be incorporated on the same power budget. Alternatively, battery size (a significant, often

dominant, contribution to the size and weight of portable devices) can be reduced by decreasing the power consumption while maintaining feature set and battery life.

The total power consumption of massively-parallel systems, such as phased-array radar systems, is also a concern, if for no other reason than the attendant cooling issues.

Figure 2.3 shows a minimum quantization energy of about 10 picojoules/conversion-step for a wide range of bandwidths, with a sharp rise above one gigahertz. However, Figure 2.4 shows that much of this progress in efficiency has occurred recently, with the lowest reported quantization energy decreasing yearly. Unfortunately, this data shows slow progress, with only a factor of ten improvement over nine years.

2.4 Absolute Minimum Power

The absolute minimum power for an analog-to-digital converter can be calculated from first principles [27, 28]. Assuming that the voltage to be converted is sampled and stored on a single capacitor, C , the minimum thermal noise is

$$V_{\text{ntc}}^2 = \frac{kT}{C}$$

If this capacitor is charged and discharged during every conversion cycle, the maximum power dissipated is

$$P = CV_{FS}^2 F_S$$

where V_{FS} is the full-scale voltage of the A/D converter and F_S is the sampling rate. The quantization noise for this A/D converter is²

$$V_{\text{nq}}^2 = \frac{V_{\text{LSB}}^2}{12} = \frac{V_{FS}^2}{2^{2N} 12}$$

²For a complete derivation of the quantization noise, see Section A.1.2 in Appendix A.

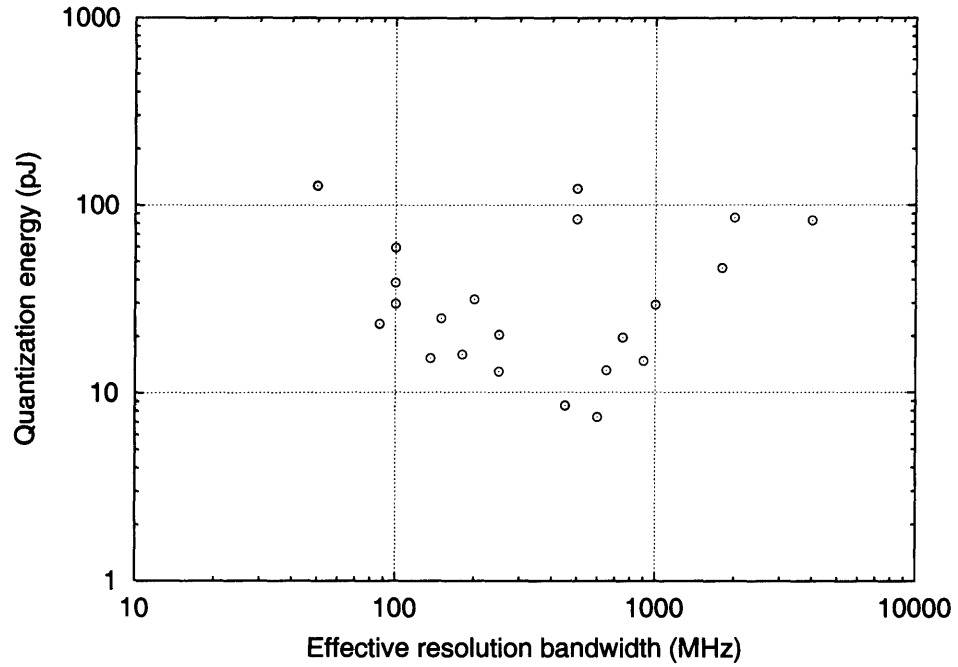


Figure 2.3: All high-speed flash A/D converters: quantization energy versus effective resolution bandwidth.




Figure 2.4: All high-speed flash A/D converters: quantization energy versus year of publication. The best-fit line shows a factor of ten improvement in quantization energy over nine years.

Limiting the thermal-noise power to be less than one quarter of the quantization-noise power, the total signal-to-noise ratio of the A/D converter will be degraded by less than 1 dB (since 1 dB corresponds to a factor of 1.25 in power units)

$$\frac{kT}{C} < \frac{1}{4} \left(\frac{V_{FS}^2}{2^{2N} 12} \right)$$

Solving this inequality for the expression $V_{FS}^2 C$

$$V_{FS}^2 C > 48kT 2^{2N}$$

and substituting this quantity into the above expression for power gives the final result

$$P > 48kT 2^{2N} F_S$$

This expression implies an absolute lower bound on the quantization energy E_Q for an A/D converter of a given resolution N at any speed

$$E_Q = \frac{P}{2^N F_S} > \frac{48kT 2^{2N} F_S}{2^N F_S} = 48kT 2^N$$

for a six-bit A/D converter at room temperature (such as most of the A/D converters listed in Figure 2.1), this limit is

$$E_Q > 1.3 \times 10^{-17} J$$

which is more than five orders of magnitude below the best A/D converter in Figure 2.4.

2.5 Aperture Jitter

A limitation in the performance of state-of-the-art A/D converters has been observed and commented on by Robert Walden [3]. Aperture jitter is the the sample-to-sample variation in the instant of conversion. There is an rms voltage error caused by rms aperture jitter (measured in picoseconds), and it decreases the overall signal-to-noise ratio.

Since the aperture-time variations are random, these voltage errors behave like a random noise source. The signal-to-jitter-noise ratio can be expressed as

$$\text{SJNR} = 20 \log \left(\frac{1}{2\pi t_a F_S} \right)$$

The graph in Figure 2.5 shows the effective number of bits versus sampling rate for the A/D converters listed in Figure 2.1. No A/D converter exhibits an aperture uncertainty better than 0.4 ps. Walden observes that this situation has not improved much over time, and will not without significant research and development.

2.6 Additional Power Consumption

Clearly, thermal noise is not the only limiting factor, and the simply-derived lower bound on power may never be approached. The development in Section 2.4 considers only the power consumption of an ideal sample-and-hold. It completely ignores the power consumption of active circuitry in the A/D converter for analog amplification or digital encoding. The power required to drive parasitic capacitances was also ignored. Considering these real-world requirements, several other limitations on performance can be calculated.




Figure 2.5: All high-speed flash A/D converters: effective number of bits versus sampling rate. The line corresponds to the limitation imposed by an rms aperture jitter of 0.4 ps.

2.6.1 Transistor Matching

Random voltage offsets in the comparators cause a limitation to accuracy that can only be overcome with improved transistor matching. Offset reduction through transistor matching increases power consumption. Transistor matching can be improved by increasing device size (which increases parasitic capacitances), by introducing calibration or error-correcting techniques (which consume additional power), or by introducing auto-zero cycles (which decrease speed). The additional power required to overcome the effects of transistor mismatch in high-speed CMOS systems is two orders of magnitude higher than the limit imposed by thermal noise [29, 30].

Reported data shows that transistor matching improves with decreasing gate-oxide thickness [30], thus for the same transistor area higher accuracy and speed, or lower power consumption, should be achievable in a technology with smaller feature sizes. However, lower supply voltages

adversely affect the signal-to-noise ratio and decrease the potential improvements [31].

2.6.2 Device Parasitics

Scaled technologies with smaller feature sizes often have larger parasitics, since the increased doping increases the values of the depletion capacitances that form around the junctions. These parasitics directly affect performance. In particular, gate-overlap capacitances and drain-to-bulk depletion-layer capacitances have been shown to have increasing influence in scaled processes [32]. It is possible that power consumption will actually need to increase to maintain accuracy and speed in the face of shrinking feature sizes [33].

2.6.3 Calibration

Calibration of A/D converters generally increases power consumption and decreases speed. Comparable digital and analog calibration techniques [34, 35], designed for similar performance improvements, significantly increase power consumption.

2.7 Future Needs

Aperture jitter, transistor matching, and increasing drain-bulk capacitance remain major problems in the development of high-speed, low-power A/D converters. In spite of these difficulties, steady progress in speed and power has been made in recent years, as shown in Figures 2.2 and 2.4. However, it seems clear that the benefits of future scaling will not achieve lower power in the face of increasing parasitic capacitances.

Chapter 3

Silicon-on-Insulator Technology

Overview

Over the past thirty years, the continued shrinking of features in CMOS technologies has provided significant gains in the performance and density of integrated circuits. However, future scaling of conventional bulk CMOS technology offers only diminishing returns. The limits of deep-submicron CMOS are evident as parasitic capacitances and leakage currents increase sharply for smaller transistor geometries. As the benefits of bulk CMOS scaling disappear, silicon-on-insulator technology promises a bright future for low-power, low-voltage, high-speed integrated circuits.

The use of SOI devices is seen as only an evolutionary step in CMOS technology development [36, page 6]. Wafer fabrication can be carried out in standard bulk CMOS processing lines, and as an additional economic advantage, SOI processing even requires fewer steps than an equivalent CMOS technology [37, page 105]. The absence of wells and isolating trenches offers additional flexibility and also makes possible the fabrication of many innovative devices [38, Section 3.8].

SOI technology improves on conventional bulk CMOS with three major advantages: less capacitance, less body effect, and the possibility for lower threshold voltages [36, page 67]. Leveraging

these advantages promises significant savings in power consumption [39]. Less junction capacitance allows lower-power digital circuits and higher-bandwidth analog circuits. Less body effect means less threshold-voltage shift and improved small-signal parameters. Lower threshold voltages allow lower supply voltages and decreased power consumption overall.

For the design of analog circuits, SOI improves on bulk CMOS in a number of important and useful ways, including better subthreshold slope, reduced short-channel effects, and increased small-signal transconductance [40].

3.1 Conventional Bulk CMOS Technology

In conventional bulk CMOS technology [41], such as shown in Figure 3.1, n-channel transistors are built directly into a p-type layer in the top of the silicon wafer. To build p-channel transistors, regions of the silicon are counter-doped with an n-type dopant to create n-wells. Highly doped source and drain regions are implanted to define the active transistor areas. A field oxide is grown between the devices to isolate adjacent transistors. Electrical isolation between transistors is achieved by keeping the diode junctions created by the source and drain implants reversed biased. To maintain this reverse bias, contacts to the substrate and to each well are created and connected to appropriate voltages.

Only one-tenth of one percent of the silicon wafer is useful for electron transport [38, page 1]. The rest of the silicon volume creates undesirable parasitic effects and potential problems. The reverse-biased junctions create significant parasitic capacitances that affect ultimate device speed. Electrical interaction between the transistor and the substrate causes leakage currents, body effect, and noise coupling. The depletion regions around the source and drain interfere with the gate control of the channel, causing a long list of short-channel effects, such as threshold-voltage roll-off, degradation of the subthreshold slope, and drain-induced barrier lowering [42, Chapter 5].

Also of particular concern in CMOS circuits is the destructive phenomenon of latchup, caused by the thyristor-like *nnpn* structure created by the numerous adjacent n-type and p-type implants [42]. If this thyristor is triggered into latchup (by leakage currents, supply glitches, radiation events, or noise), it can conduct large currents through the chip, causing significant damage.




Figure 3.1: Cross section of a conventional bulk CMOS technology, showing devices separated by field oxide and isolated by an n-well, along with the required contacts to substrate and well.

3.2 SOI CMOS Technology

The advantages of SOI transistors derive from their physical structure. In contrast to bulk CMOS technology, the transistors are built on top of a layer of silicon dioxide, called the buried oxide layer, as shown in Figure 3.2. From the figure, the relative simplicity of SOI technology is readily apparent. Adjacent transistors are completely isolated from each other by etching away unused silicon, creating individual islands. There is no need for wells, since the silicon islands can be individually doped for n-channel or p-channel transistors. There is no need for well contacts or substrate contacts, since there are no wells and no common substrate. This full dielectric isolation, created by the silicon island and the buried oxide, create the primary advantages of SOI technology.




Figure 3.2: Cross section of a fully depleted SOI CMOS technology, showing mesa-isolated devices supported by buried oxide layer.

3.3 Advantages of SOI

SOI technology has emerged, not only as a solution to many submicron CMOS problems [40], but also as a promising opportunity for high-speed, low-power applications [43, 44]. The advantages of SOI technology are directly related to its physical structure.

3.3.1 Full Dielectric Isolation

The complete isolation of the transistors on individual silicon islands creates a number of significant advantages. Isolating wells are not required, as in bulk CMOS, which allows a higher layout packing density. With higher density circuits, there is less parasitic capacitance due to interconnect. Also, due to the complete isolation, there are no thyristor-like structures, and the notorious thyristor latchup effect is no longer present.

The buried oxide layer not only isolates the transistors from each other, but also from the wafer substrate. With a layer of oxide between the transistors and the substrate, there is no leakage current from any circuit node to the substrate. Also, due to this isolation, substrate noise coupling is virtually nonexistent, and can be ignored [36, page 68].

3.3.2 Reduced Capacitances

In bulk CMOS, the source-to-bulk and drain-to-bulk junction capacitances are often the dominant parasitic capacitances. As higher and higher substrate dopant concentrations are used in scaled technologies, these junction capacitances increase significantly. In SOI devices, the source and drain diffusions do not abut the substrate. The parasitic source-to-bulk and drain-to-bulk capacitors occur across the buried oxide, which results in capacitances 4 to 7 times smaller than in bulk [44], depending on the thickness of the buried oxide.

3.3.3 Improved Subthreshold-Slope Factor

Current flow in a MOSFET transistor occurs through the formation of a channel between source and drain, which is indirectly controlled by the gate voltage. The degree of control between the gate voltage and the surface potential of the channel is denoted by the subthreshold-slope factor n [44, 45], sometimes called the body factor [43], body-effect factor [45], or body-effect coefficient [46, 47].

The subthreshold-slope factor n appears in many equations of interest [46], including the expression for subthreshold slope (with units of decades-of-current/volt)

$$\frac{1}{I_D} \left(\frac{dI_D}{dV_G} \right) = \frac{q}{nkT}$$

the expression for drain current in strong inversion

$$I_{D,sat} = \frac{\mu_n C_{ox}}{2n} \frac{W}{L} (V_{GS} - V_T)^2$$

and the expression for transconductance in weak inversion

$$g_m = \frac{qI_D}{nkT}$$

Inefficient coupling from the gate to the channel occurs because of the capacitive divider between gate, channel, and substrate (ground). The subthreshold-slope factor depends on the ratio of the capacitance from the gate to the channel, C_{ox} , and the capacitance from the channel to the substrate. In bulk CMOS, the capacitance from the channel to the substrate is created by the depletion region under the channel

$$n = 1 + \frac{\epsilon_s}{C_{\text{ox}}x_d}$$

where x_d is the maximum depletion width in strong inversion. Typical values for n in bulk CMOS range from 1.4 to 1.6 [47].

In a fully depleted SOI device, the capacitance from the channel to the substrate is reduced by the buried oxide thickness and is significantly smaller, resulting in excellent coupling between the gate voltage and surface potential [40]. The space-charge region below the gate extends through the silicon island, across the thick buried oxide, and into the handle wafer. Thus, the subthreshold-slope factor is greatly improved, and is now equal to

$$n = 1 + \frac{C_{\text{film}}C_{\text{box}}}{C_{\text{ox}}(C_{\text{film}} + C_{\text{box}})}$$

where $C_{\text{film}} = \epsilon_s/t_s$ is the capacitance across the island, and $C_{\text{box}} = \epsilon_{\text{ox}}/t_{\text{box}}$ is the capacitance across the buried oxide. Typical values for n in SOI CMOS range from 1.05 to 1.1 [47].

This thirty to fifty percent improvement in subthreshold-slope factor has a number of positive effects. Compared to a bulk transistor with similar parameters, this improvement predicts SOI

devices with higher saturation drain current [47]. Measurements show that g_m/I_D is always larger for fully depleted SOI transistors than for comparable bulk devices [45].

Indirectly, improved subthreshold slope allows the design of SOI devices with lower threshold voltages for identical subthreshold leakage-current values as their bulk counterparts [45]. Lower threshold voltages allow lower supply voltages, and thus lower total system power.

3.3.4 Short-Channel Effects

In small-geometry bulk CMOS transistors, the depletion regions around the source and drain junctions interfere with the active channel region underneath the gate. This causes a long list of short-channel effects: threshold voltage roll-off, degradation of the subthreshold slope, drain-induced barrier lowering, and others. In SOI technology, the extension of the source and drain depletion regions is restricted by the thickness of the silicon island, limiting interaction with the channel and thus limiting short-channel effects.

3.3.5 Harsh Environments

SOI circuits have long been used in harsh environments, such as high-radiation environments and high-temperature applications. Radiation hardness in SOI CMOS is improved over bulk CMOS, since carriers created in the bulk wafer by energetic particles cannot flow into the active circuits (due to the buried oxide layer), and because of the reduced volume of silicon used by the transistor (less volume for induced defects). High temperature performance is improved due to the absence of thermally-activated latchup, reduced leakage currents, and a smaller variation of threshold voltage with temperature [37, Chapter 7].

3.4 Disadvantages of SOI

Unfortunately, SOI technology does include drawbacks, including the floating body effect, the kink effect, self heating, and excess noise.

3.4.1 Floating-Body Effect

In bulk CMOS technology, excess charge carriers in the body can escape to the substrate or to a well contact. In SOI technology, there is no conducting substrate, and contacts to the body are often left out (to improve device packing density). In fully depleted SOI technology, body contacts are often not effective because of the high resistance of the depleted body. Thus, unwanted carriers in the body (holes in an NMOS device) can charge the body and change the effective threshold voltage of the transistor through the backgate effect. This phenomenon is known as the floating-body effect.

Charge in the body can come from a number of sources, including leakage currents from the drain and source (depending on operating point) and impact ionization in the channel (when V_{DS} is high enough). Thus, the threshold voltage of the transistor is dependent on operating point history. This hysteresis effect can have grave consequences for many applications.

Also, although SOI circuits are free of the standard thyristor latchup effect, there exists a parasitic bipolar transistor. For example, in a NMOS device, the $n+$ source and drain diffusions and the p -doped channel make a NPN transistor with a floating base terminal. If the body-to-source voltage ($V_{BS} = V_{BE}$) is high enough, this transistor can conduct unwanted excess current from the collector (drain) to the emitter (source). Fortunately for low-voltage applications, this parasitic bipolar transistor does not conduct when the drain-to-source voltage is less than two volts [44].

3.4.2 Kink Effect

The kink effect is a special case of the floating-body effect and is characterized by a sharp upturn in the I_D versus V_{DS} curve of the transistor. See the typical family of curves shown in Figure 3.3. In an NMOS device when the drain voltage is high enough (larger than the bandgap voltage, about 1.1 V in silicon), energetic carriers (electrons) can create electron-hole pairs through impact ionization. The new electrons flow normally to the drain, while the free holes charge up the body (in the absence of a body contact). This body charge reduces the effective threshold voltage (analogous to a backgate bias), which further increases the drain current. Also, if this impact ionization hole current is large enough, it can provide enough base current to the parasitic bipolar NPN transistor and create a second kink [37, page 159]. The kink effect can be avoided by keeping the drain-to-source voltage less than the bandgap voltage.

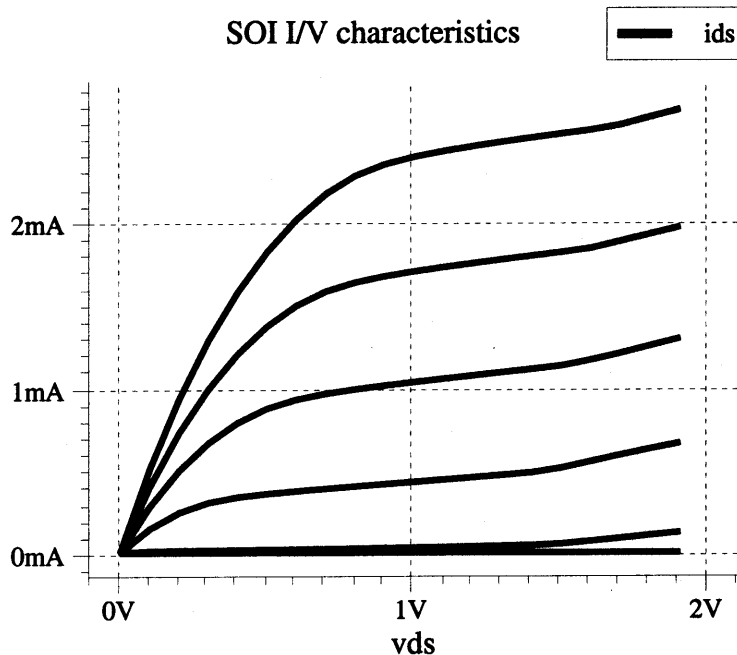


Figure 3.3: I_D versus V_{DS} curves for a typical NMOS SOI transistor, showing the kink effect at large drain voltages (when V_{DS} is larger than the bandgap of silicon, which is required for impact ionization — this voltage is referred to as the “kink onset voltage”).

3.4.3 Self Heating

In conventional bulk CMOS wafers, neighboring devices are assumed to be at the same temperature because of the excellent thermal conductivity of the substrate silicon. However, in addition to being an excellent electrical insulator, the silicon dioxide in the buried oxide layer of SOI is also an excellent thermal insulator. In fact, the thermal conductivity of silicon dioxide is one hundred times smaller than the thermal conductivity of bulk silicon.

This thermal insulation gives rise to the self heating of SOI transistors. As power is dissipated in the channel of an SOI device, the generated heat does not conduct to the substrate, and can cause a significant local temperature increase [48]. This temperature increase can affect the operating point of the transistor. The heating and cooling of a single device can exhibit a time constant of one microsecond to ten seconds [49].

3.4.4 Noise in SOI

SOI transistors have all the noise sources of standard bulk CMOS transistors, including thermal noise and flicker noise. In addition, they also exhibit an additional “excess” noise source [50, 51, 52]. Appendices B and C examine theories of these noise sources.

3.5 Advantages for Analog Circuits

Smaller capacitances, improved subthreshold-slope factor, and reduced short-channel effects make SOI attractive for the design of analog circuits. In fact, a wide variety of analog circuits have been produced in SOI, although mostly in partially depleted technologies.¹ However, fully depleted SOI has several additional advantages for analog circuit design.

¹For examples of analog circuits in partially depleted SOI, see references [53, Chapter 6] and [54, Chapter 10].

3.5.1 Lower Analog Supply Voltage

The lowest acceptable analog supply voltage V_{DD} can be calculated by considering a simple analog switch, consisting of an NMOS transistor and a PMOS transistor connected in parallel [55]. The switch must conduct for any terminal voltage between ground and V_{DD} in the on state. Assuming the switch control voltages are ground and V_{DD} , the supply voltage must be

$$V_{DD} \geq \frac{2V_T}{2-n}$$

If the supply voltage is less than this limit, there will be a range of voltages, near one-half V_{DD} , for which neither transistor conducts. In bulk CMOS, where the threshold voltage is $V_T = 0.7$ volts and the subthreshold slope factor is $n = 1.5$, V_{DD} is limited to a minimum voltage of 3 volts. In fully depleted SOI where the threshold voltages can be reduced to $V_T = 0.4$ volts (for similar $V_{GS} = 0$ leakage) and the subthreshold slope factor is $n = 1.05$, V_{DD} can be as low as 0.85 volts [46].

3.5.2 Higher Maximum Gain

Optimum device biasing in analog circuits is often found by considering the g_m/I_D ratio in saturation. This ratio can be viewed as a measure of device quality, in terms of the transconductance (gain) achieved for a given current (power dissipation). This ratio reaches a maximum in weak inversion, and is

$$\frac{g_m}{I_D} = \frac{q}{nkT}$$

combined with the Early Voltage, the intrinsic voltage gain of a current-source-loaded common-source amplifier is

$$A = \frac{g_m}{g_d} = g_m \frac{V_A}{I_D} = \frac{qV_A}{nkT}$$

Since SOI has comparable Early Voltages to bulk CMOS, operational amplifiers with similar or improved low-frequency gain can be designed that significantly outperform bulk implementations [43]. Leveraging improved transconductances and lower capacitances, fully depleted SOI op amps can be designed that have unity-gain frequencies ten times higher than comparable bulk CMOS designs.

Chapter 4

High-Speed Flash Design

Flash analog-to-digital converters are by far the fastest and conceptually simplest of all A/D converter topologies [56, 57]. Most of the high-speed converters listed in the table of Figure 2.1 are flash converters.

In a typical flash converter, there is one comparator corresponding to each transition in the digital code. The three-bit converter in Figure 4.1 has eight possible digital output codes, thus seven transitions and seven comparators. The resistor string produces a complete set of reference voltages, one for each comparator, corresponding to the analog voltages of the the transition codes. These analog transition voltages are shown along the horizontal axis of Figure 4.2.

The input voltage is applied to the non-inverting input of each comparator. During conversion, each comparator compares the input voltage to its corresponding reference voltage, and outputs a logic 1 if the input is higher than its reference voltage or a logic 0 if the input is lower. Thus the outputs of the comparators produce a “thermometer code” of logic 1s, whose height corresponds to the voltage of the analog input.

Simple digital logic converts the thermometer code into a binary-coded digital word. The outputs of the comparators are first latched, so they do not change while being converted into binary.

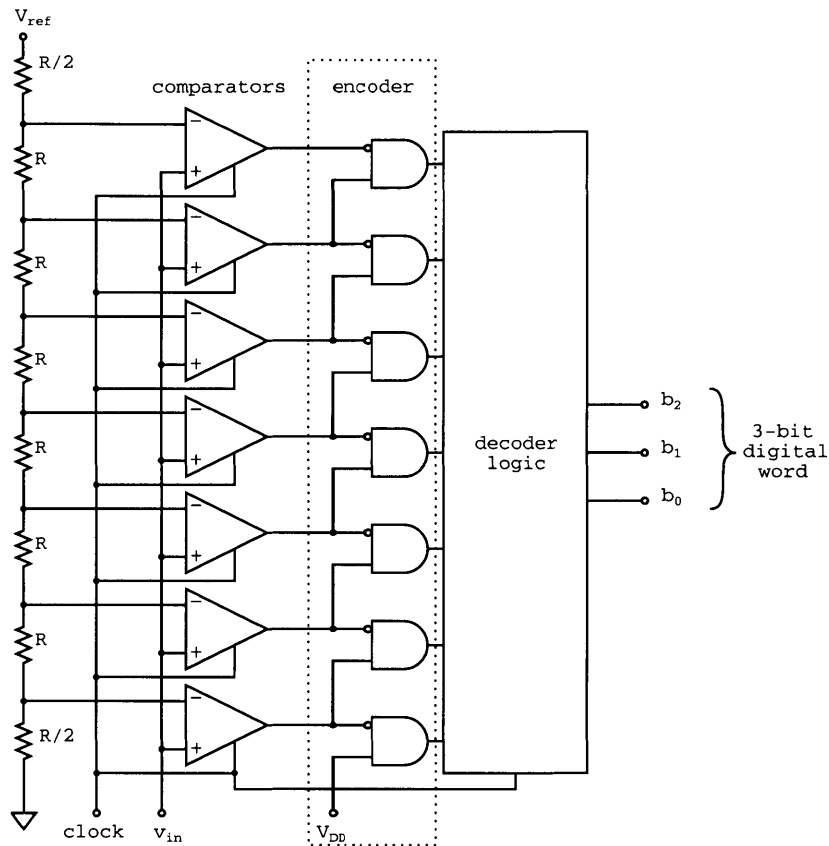


Figure 4.1: Schematic of simplified flash topology.

In the simplest case, shown in Figure 4.1, logic gates look for the single break in thermometer code (where it transitions from logic 1 to logic 0), and a simple decoder converts this one-of-many input into the appropriate digital output.

Flash converters require a large number of parts (at least $2^N - 1$ resistors, comparators, and gates), thus they suffer from large chip area, large power dissipation, and large input capacitance. Flash converters can easily reach speeds faster than 1 GS/s, but due to space and power requirements, are generally limited to 6 to 8 bits of resolution. However, there is no faster converter for high-speed applications.

This thesis demonstrates that in SOI technology, high-speed flash converters can also be ultra-low-power flash converters.




Figure 4.2: Ideal analog-to-digital converter characteristic.

4.1 Analog Front End

As shown in Figure 4.1, the analog front end of the flash converter consists of the reference resistor string and the comparators. The comparators compare the input signal to the $2^N - 1$ reference voltages created by the resistor string. The resistors in the resistor string must be small enough to supply stable reference voltages, even when loaded by the input bias currents flowing into the comparators.

The comparators must be able to resolve to the full accuracy of the converter. Thus the comparator gain must be at least 2^N . This requirement of comparator gain can be relaxed if the latch does not require full-swing inputs. Comparator design is discussed in Chapter 5.

4.2 Latch Circuit

The latches take the output from the comparators, convert it to a valid logic 1 or logic 0, and hold it until the next comparison is ready. The first part of the latch shown in Figure 4.3 is a simple coupled-inverter regenerative latch. In the precharge state, the tail-current transistor MLT is off,

and the switches are closed allowing the gates of the inverters to be charged. In the latch mode, the input switches are opened, and the tail current is turned on. As the tail current is turned on, the initial condition stored on the gates is regenerated into a valid logic level. Simulations show that this latch topology renders reliable decisions with as little as 100 mV of differential input voltage.

The second half of the latch is a doubled half-latch [58]. When the clock input PH is high, the latch is its transparent mode, and both inverters are enabled. When the clock input is low, both inverters are disabled, and the latch is in its hold mode. The PMOS transistors can still pull up their outputs, but the NMOS transistors cannot pull down when the clock is low. Thus no signal can propagate from the input to the output in the hold mode.




Figure 4.3: Regenerative and doubled half-latch schematic.

Proper operation of the latch requires the clock signals shown in Figure 4.4. In the clock phases shown, inverter gates are precharged during the first half of the clock cycle, and the input is latched during the second half. The half-latches store the new output value during the last quarter of the clock cycle.




Figure 4.4: Clock phases required for latch operation. The dashed line denotes a complete clock cycle, while the dotted lines show quarter periods.

4.3 Encoding Logic

Ideally, the outputs of the latches generate a thermometer code, with all logic 1s below the input level and all logic 0s above. If the input is changing rapidly, slight timing differences in the latches or errors from the comparators can create extraneous zeros or ones in the thermometer code, even though this cannot happen for DC inputs. If the decoder in Figure 4.1 were implemented with OR gates, or a simple ROM topology, these “bubbles” and “sparkles” would create disastrous errors. A bad thermometer code may have several one-to-zero transition points.

Simple bubble errors can be corrected by the majority circuit [59] shown in Figure 4.5. The output of the majority circuit is high if two or three of its inputs are high. Similarly, output is low if two or three of its inputs are low. In effect, each bit of the thermometer code and its two nearest neighbors vote on the true value of each comparator output.




Figure 4.5: Schematic of the majority circuit.

A zero in the thermometer code surrounded by ones is correctly changed to a one. A one surrounded by zeros is similarly corrected. Thus, this circuit corrects all single comparator errors, regardless of the error. Given this behavior, either of the following two patterns will be identified as the true break point in the thermometer code:

$$\begin{array}{cc} 0 & 0 \\ 0 & \text{or } 1 \\ 1 & 0 \\ 1 & 1 \end{array} \left(\begin{array}{cc} 0 & 0 \\ \text{since } 1 & \text{would be corrected to } 0 \\ 0 & 1 \\ 1 & 1 \end{array} \right).$$

The first pattern is the expected pattern at the top of the thermometer code. The second pattern occurs when bubbles are near the top of the thermometer code, or where two comparators have interchanged effective reference voltages (due to large and unfortunate offset voltages). The majority circuit corrects the latter pattern to the correct breakpoint pattern.

This circuit also makes the converter insensitive to stuck comparators. A comparator with a stuck-high output, which would disable all output codes below it in the encoder in Figure 4.1, only causes a single missing code. With the majority circuit, every third comparator could be faulty, and the converter would still give meaningful results, losing less than a bit of resolution. Figure 4.6 shows an example of 4-bit thermometer code with three errors, and the results of the majority circuit's error correction.

The majority circuit used in this design is shown in Figure 4.7. The gates are implemented with standard CMOS logic gates. This implementation takes advantage of the repeating structure of the majority sums. Using a carry bit in the majority circuit saves a NAND gate over a direct implementation. Also shown in Figure 4.7 is the simple logic that detects the zero-to-one transition point in the corrected thermometer code.

0	}		
1		} ignored sparkle	
0			
0			
0			
1	}	} correct breakpoint	
0			
1			
1			
1			
1			
1			
0	}	} ignored bubble	
1			
1			

Figure 4.6: Examples of errors in the thermometer code.

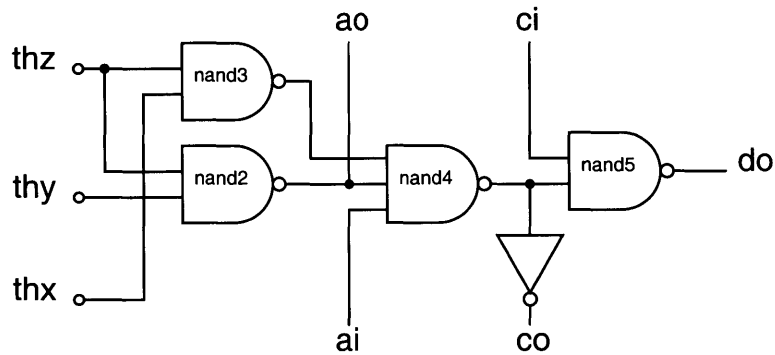


Figure 4.7: Schematic of optimized majority circuit topology with transition point detection.

4.4 Output ROM Decoder

The zero-to-one transition point in the corrected thermometer code is used to address a ROM. The output of the ROM is the final binary word. Each output bit from the ROM is generated by a 32-input pseudo-PMOS-style NOR gate. A simplified 8-input NOR gate, such as would be used in a 4-bit flash converter, is shown in Figure 4.8.




Figure 4.8: Schematic of an 8-input pseudo-PMOS NOR gate, used in each bit of the output ROM decoder for a 4-bit flash converter.

Chapter 5

Comparator Design for SOI

The comparator design is the most important piece of a flash analog-to-digital converter. The accuracy and performance of the internal comparators directly affect the resolution and speed of many A/D converter topologies. In the flash converter architecture, with its $2^N - 1$ comparators, the comparators are also a significant contribution to the total power consumption and die area of the topology.

5.1 Differential Comparators in Bulk and SOI

Many flash converter designs [4, 11, 16] use differential comparators, such as that shown in Figure 5.1. A simple differential comparator of this type can be easily designed to meet the specifications required for the flash A/D converter. The gain and bandwidth can be easily determined from transistor parameters.

However, this topology is problematic in SOI technology. The floating-body effect can cause considerable effective threshold-voltage shift in the input transistor pair. If a large differential voltage is applied to the comparator, one input transistor will be forced into cutoff, while the other input transistor will be in strong inversion. If this difference in bias point is held for a long time,




Figure 5.1: Schematic of standard continuous-time differential-input comparator topology.

the body voltages of the two input transistors can charge to vastly different values, resulting in considerable offset voltage between them.

Figure 5.2 shows the floating-body effect in action. The non-inverting input to the comparator is held at one volt, while the inverting input is stepped from zero volts to nearly one volt. After the initially large differential voltage, the body voltages of the two transistors are very different, and there is a significant initial undershoot in the output voltage due to the effective offset. Eventually, the floating body discharges, producing a long-tail transient, but as can be seen, this recovery takes approximately fifty microseconds.¹

In a flash converter, this error and long-tail transient could create significant errors. The output-voltage error shown in Figure 5.2 is more than enough to cause the latches to store the wrong digital state, creating large strings of incorrect bits in the thermometer code.

This long-tail transient does not occur in standard bulk CMOS technology, since there is no floating-body effect. Figure 5.3 shows the result for the same comparator, simulated in a comparable bulk technology. After the input voltage step, the output quickly changes without any initial undershoot or long-tail transient.

¹The exact time constant is strongly dependent on the impact ionization and carrier lifetime model parameters. These results were simulated with the MIT Lincoln Laboratory 0.25 μm BIM3SOI models (9/1999).

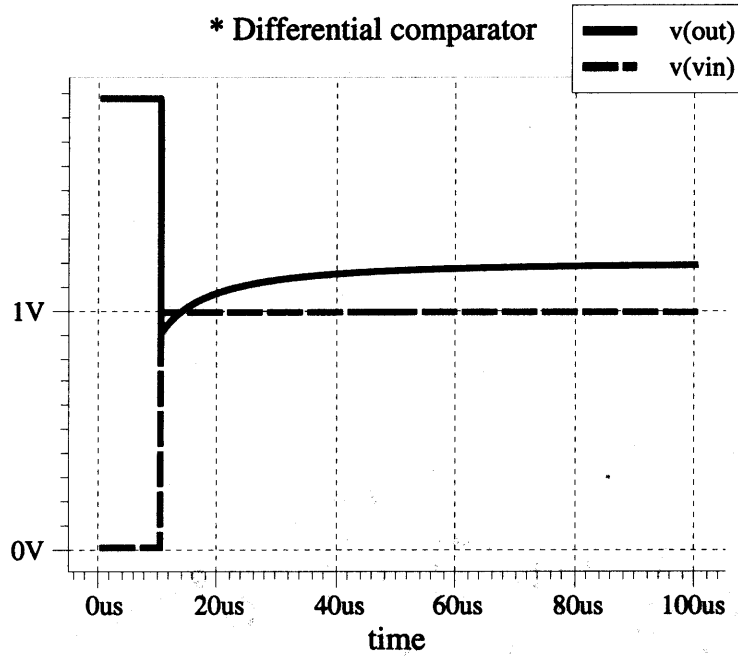


Figure 5.2: Step response for a differential comparator in SOI technology. The floating-body effect causes significant initial undershoot and long-tail transient as the body charge changes.

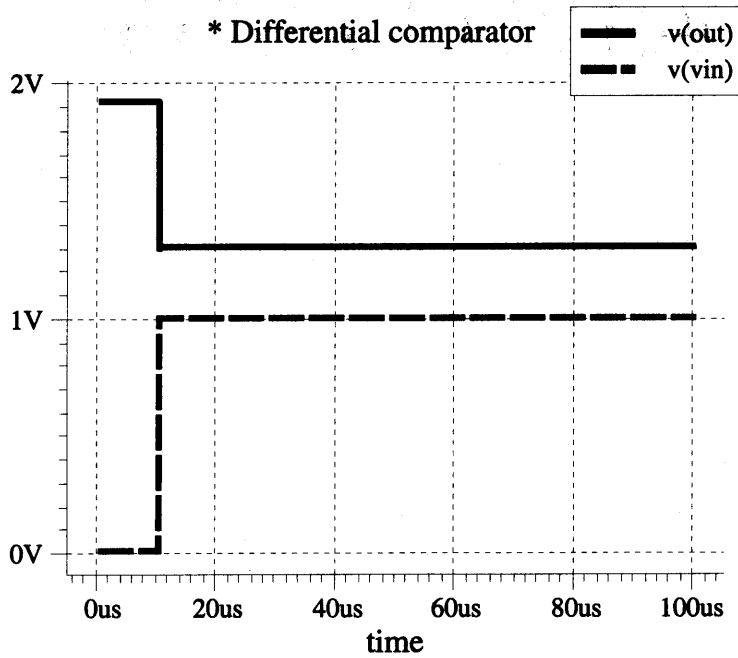


Figure 5.3: Step response for a differential comparator in a comparable bulk CMOS technology. Note absence of any initial error or long-tail transient.

5.2 Single-Ended Auto-Zeroed Comparator

Offset-cancellation schemes are commonly used to reduce the effects of transistor mismatch and low-frequency noise [60]. Usually, transistor mismatch arises from slight imperfections in the fabrication process, such as uncontrolled variations in the photolithography exposure, etch steps, or doping implant concentrations. In SOI technology, the floating-body effect can be treated as transistor mismatch, and corrected using these standard offset cancellation techniques. Figure 5.4 shows a simple single-ended auto-zeroed comparator [61] that meets the necessary specifications.




Figure 5.4: Simplified schematic of single-ended auto-zeroed comparator topology.

The circuit operates in two phases. In the first phase, the auto-zero mode, switch S_1 is opened and switches S_2 and S_3 are closed. The unity-feedback loop around the inverter causes it to self-bias in the middle of its linear region. The transistor dimensions of the inverter are chosen so that the self-bias voltage is approximately half of the supply voltage, $V_B = V_{DD}/2$, and at the point of maximum gain, as shown in Figure 5.5. With switches S_2 and S_3 closed, the capacitor charges to a voltage $V_C = V_M - V_B$, where V_M is the input voltage at the inverting input, and V_B is the self-bias voltage of the inverter.

In the second phase, the compare mode, switches S_2 and S_3 are opened and switch S_1 is closed. With switch S_3 open, the right side of the capacitor is floating, so the capacitor voltage cannot

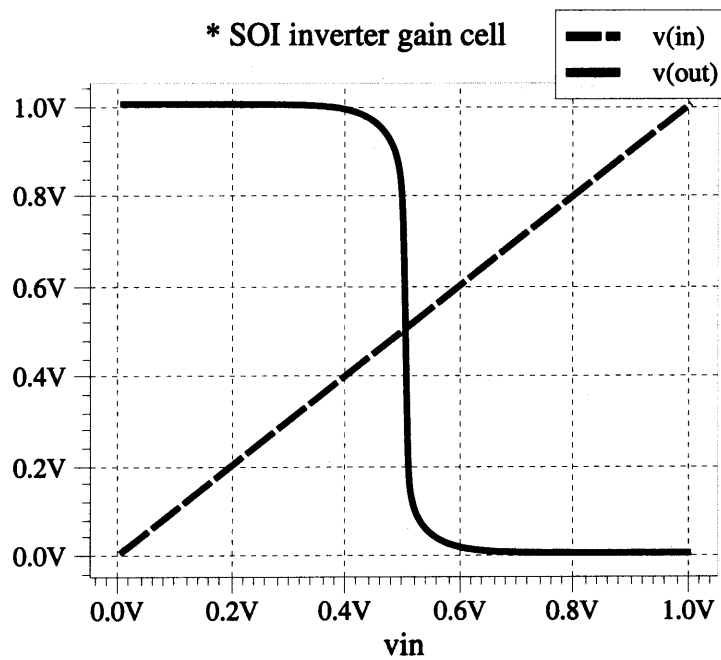


Figure 5.5: Inverter transfer characteristic from input to output. In the auto-zero mode, the inverter self-biases at the intersection point, where the maximum slope of the transfer characteristic provides maximum gain for the next compare phase.

change and thus remains V_C . With switch S_1 closed, the input voltage to the inverter becomes

$$V_P - V_C = V_P - V_M + V_B$$

where V_M is the input voltage at the non-inverting input. This input voltage is amplified by the gain of the inverter. Thus, the output goes low if $V_P - V_M > 0$ and goes high if $V_P - V_M < 0$.

By effectively using the same transistors for both comparator inputs, the floating-body effect can be ignored. Using the same test from Figures 5.2 and 5.3, the absence of any initial undershoot in Figure 5.6 shows that this offset cancellation scheme overcomes any change in threshold voltage due to floating body effect.

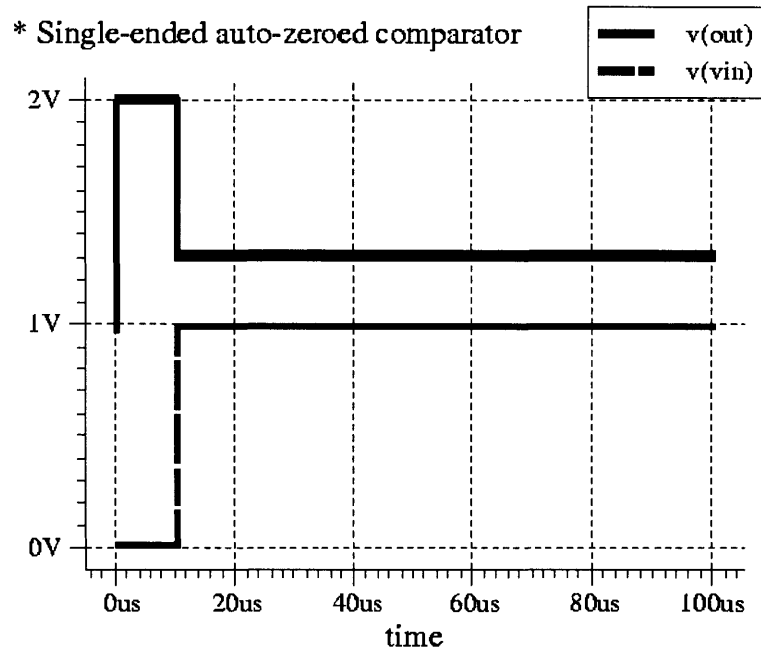


Figure 5.6: Step response for a single-ended auto-zeroed comparator in SOI technology. Note absence of any initial error or long-tail transient due to offset-cancellation scheme.

This simple comparator topology also operates with very low supply voltages. By using a supply voltage of $V_{DD} = 1$ V, the floating-body effect is further suppressed by greatly reducing the impact

ionization and its associated kink effect. Low supply voltage also enables low power consumption, which is also important in this design.

5.2.1 Comparator Gain

The gain of the single-ended auto-zeroed comparator can be calculated from the gain of the inverters and the attenuation due to the input capacitance. The gain of the inverter is simply

$$A_i = \frac{g_{mn} + g_{mp}}{g_{dn} + g_{dp}}$$

where g_{mn} and g_{mp} are the transconductances of the NMOS and PMOS transistors, respectively, and g_{dn} and g_{dp} are the drain conductances of the NMOS and PMOS transistors, respectively. This gain can be between fifty and one hundred.

A capacitive attenuator is formed by the sampling capacitor, C , and the total capacitance at the input of the inverter. This total gate capacitance is formed by the gate-to-source capacitances c_{gs} and the Miller-multiplied gate-to-drain capacitances c_{gd} . The total input capacitance is

$$C_T = c_{gsn} + c_{gsp} + (A_i + 1)(c_{gdn} + c_{gdp})$$

Thus the comparator gain from input to output is

$$A = \frac{A_i C}{C + C_T}$$

Great care is put into the design of the transistors in the inverter. The relative transistor widths are chosen so that the self-bias voltage is approximately half of the supply voltage. The absolute transistor widths were chosen to keep the drain-to-bulk capacitance very small. Widths of




Figure 5.7: Comparator gain versus transistor gate length, from minimum length to two microns.




Figure 5.8: A close-up of the comparator-gain-versus-transistor-gate-length relationship.

four microns and two microns for the PMOS and NMOS transistors, respectively, made the total drain-to-bulk capacitance $c_{db} = 1$ fF.

Due to short-channel effects at the minimum gate length, the gain of the inverter is improved by making the devices longer than minimum length. However, the attenuation due to the input capacitance increases with total gate area, so a compromise is made between inverter gain and input capacitance. As the simulation results show in Figure 5.7, for the chosen gate widths, the comparator gain reaches a maximum for a gate length around $0.5 \mu\text{m}$, but falls off sharply for shorter or longer gates. Figure 5.8 shows detail of the relationship for deep-submicron gate lengths.

5.2.2 Comparator Speed

To assess the speed of the comparator, it is necessary to examine the time constants during the auto-zero phase and during the compare phase.

Auto-Zero Speed




Figure 5.9: Simplified small-signal equivalent circuit for the single-ended auto-zeroed comparator topology during the auto-zero phase.

The small-signal equivalent circuit topology for the auto-zero mode is shown in Figure 5.9, when switch S_1 is open and switches S_2 and S_3 are closed. The negative feedback around the inverter causes it to reach steady-state with a time constant that is approximately

$$\tau_{az} = \frac{(R_{SW} + r_o)(C + C_T)}{A_i + 1}$$

Due to the feedback around the inverter, the effective open-circuit time constant due to drain-to-bulk capacitance c_{db} is negligible. For typical device parameters, the secondary time constant due to the charging time of the sampling capacitor C through switch S_2 ,

$$\tau = R_{SW}C$$

is much smaller, assuming a low-impedance source. The accuracy achieved during the auto-zero

phase determines the total accuracy of the comparator. Any residual unsettled voltage at the end of the auto-zero phase contributes directly to offset error in the comparison phase. To achieve 1/2 LSB accuracy for a 6-bit flash converter (approximately one percent), the auto-zero phase must last at least five time constants.

Compare Speed




Figure 5.10: Simplified small-signal equivalent circuit for the single-ended auto-zeroed comparator topology during the compare phase.

The small-signal equivalent circuit topology for the compare mode is shown in Figure 5.10, when switch S_1 is closed and switches S_2 and S_3 are open. This topology has two independent poles, one pole due to the charging time of the input through the switch, with time constant

$$\tau_{c1} = \frac{R_{SW} C C_T}{C + C_T}$$

and one pole due to the low pass filter at the output

$$\tau_{c2} = r_o C_{db}$$

Depending on the relative values of the components, either of these time constants could dominate the speed of the circuit in the compare mode. Specific numbers for this design are calculated in Section 5.3.1.

5.2.3 CMOS Switches

A significant contribution to comparator error is due to nonidealities in the switches, primarily clock feedthrough and charge dump. When switch S_3 opens, clock feedthrough from its overlap capacitance and charge dump from the channel can affect the voltage stored on the capacitor. Conversely, since switches S_1 and S_2 are operated on opposite phases of the clock, their charge dump and clock feedthrough effectively cancel out.

Complementary switches, with both NMOS and PMOS transistors, as shown in Figure 5.11, were necessary in the design to allow rail-to-rail (0 to 1 volt) input voltage range. Some cancellation of charge dump and clock feedthrough can be achieved by using complementary switches. Unfortunately, this cancellation is not perfect, and thus there are competing constraints on the sizes of the switches. Larger switches have lower resistance thus allow for faster circuit operation, while smaller switches feature less charge dump and thus better accuracy.




Figure 5.11: Complementary CMOS switch, with both NMOS and PMOS transistors to minimize charge dump and clock feedthrough and allow rail-to-rail operation.

The resistance of the switch was designed to be less than 1500Ω for all input signals using minimum-length, eight-micron-wide transistors. Charge dump and clock feedthrough were determined to be small enough by simulation. Switch resistance as a function of the input voltage is shown in Figure 5.12. Simulated off resistance was several megohms.

A simulation of the output-voltage waveform of the complete comparator is shown in Fig-




Figure 5.12: Resistance of the switch versus input voltage.




Figure 5.13: A single comparison accomplished by the single-ended auto-zeroed comparator.

ure 5.13. When the clock goes high, turning on switches S_2 and S_3 , the comparator goes into auto-zero mode, and the output voltage approaches the self-bias point. When the clock waveform goes low, there is some clock feedthrough, pulling the output up and down, but the inverter remains safely near the middle of its linear range. When the complementary clock signal (not shown) goes high, turning on switch S_1 , the comparator makes an unequivocal decision.

5.3 Comparator Specifications

Using typical transistor parameters for MIT Lincoln Laboratory's 0.25- μm fully depleted SOI CMOS process, comparator speed and power are estimated.

5.3.1 Calculated Speed

circuit parameter	symbol	value
sampling capacitor	C	50 fF
inverter gain	A_i	70
total gate-to-source capacitance	$c_{gsn} + c_{gsp}$	10 fF
total gate-to-drain capacitance	$c_{gdn} + c_{gdp}$	1.2 fF
output resistance r_o	$(g_{dn} + g_{dp})^{-1}$	100 k Ω
drain-to-bulk capacitance	$c_{dbn} + c_{dbp}$	1 fF

Figure 5.14: Table of typical SOI transistor and circuit parameters from the designed comparator

The switching transients in Figure 5.13 demonstrate that the designed comparator can operate at clock speeds near 500 MHz. Simple hand calculations using the parameters from the table in Figure 5.14 and the equations from Section 5.2.2 support this result. For the designed inverter, the input capacitance is

$$C_T = c_{gsn} + c_{gsp} + (A_i + 1)(c_{gdn} + c_{gdp}) = 94 \text{ fF}$$

The auto-zero time is dominated by the charging of the sampling capacitor through the switch.

For approximate one percent accuracy, this time is

$$t_{az} = 5\tau_{az} = \frac{5(R_{SW} + r_o)(C + C_T)}{A_i + 1} = 1000 \text{ ps}$$

This time is the full auto-zero period for a 500 MHz sampling rate. In comparison, the time constant for the charging of the sampling capacitor through the switch is only 75 picoseconds.

The compare time is the combined charging times of the input capacitance through the switch and the output capacitance by the inverter. Accuracy is not required in this case, only a decisive output.

$$t_{c1} = 3\tau_{c1} = \frac{3R_{SW}CC_T}{C + C_T} = 3(1500 \Omega)(33 \text{ fF}) = 150 \text{ ps}$$

$$t_{c2} = 3\tau_{c2} = 3r_{ocdb} = 3(100 \text{ k}\Omega)(1 \text{ fF}) = 300 \text{ ps}$$

The decisive output shown in Figure 5.13 takes about 500 ps, which agrees with these numbers. These calculations demonstrate sufficient comparator speed.

5.3.2 Power Consumption

Upper and lower bounds on the power consumption of the complete comparator can be estimated. Unlike a standard digital inverter, the comparator spends significant time in the linear region (at least the entire auto-zero time), dissipating static power. For vanishingly small differences between the input voltages, when the comparator is unable to make a decision, the inverter may stay in the linear region, dissipating power during the compare time as well. Assuming equal time in the auto-zero and compare modes, the power consumption of the comparators is bounded by

$$0.5V_{DD}I_{on} < P < V_{DD}I_{on}$$

where I_{on} is the current that flows through the inverter in its linear region. For the designed comparator, this power dissipation only amounts to a few microwatts per comparator. Thus the power consumed by the comparators pales in comparison to the power required to drive the gates of the switch transistors, which is, worst case

$$P = CV^2F = 3(20 \text{ fF})(1 \text{ V})^2(500 \text{ MHz}) = 30 \mu\text{W}$$

Thus, thus comparators are a small fraction of the total power. Significant power will be dissipated in the clock driver circuitry.

5.4 Comparison to Bulk CMOS Implementation

While the single-ended auto-zeroed comparator topology has been used in bulk CMOS technology [9, 12], the large parasitic capacitances in bulk CMOS severely limit its performance. The designed comparator was also simulated in a comparable bulk CMOS technology.² The much larger drain-to-bulk capacitance in bulk CMOS technology greatly lengthens the compare time, as shown in Figure 5.15. The compare time constant appears to be two or three times larger than for the SOI implementation.

In addition, the power consumption for the bulk design is much larger. Unlike the SOI design, the bulk implementation was nonfunctional for a one-volt supply. With a two volt supply (as shown in Figure 5.15) the design consumed ten times more power in the auto-zero phase, but was still not fast enough in the compare mode.

²Taiwan Semiconductor Manufacturing Company's 2.5-V 0.25- μm logic salicide process.

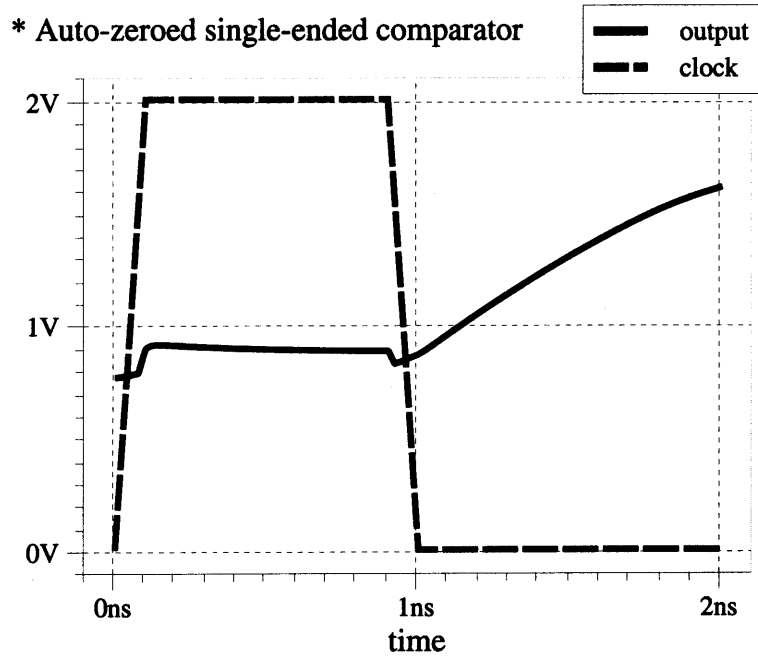


Figure 5.15: A single comparison accomplished by the single-ended auto-zeroed comparator in comparable bulk CMOS technology, showing significantly slower performance in the compare mode.

Chapter 6

Complete Circuit and Simulations

A complete six-bit flash converter was designed using the circuits from the previous chapters. With the main circuit pieces and the comparators designed, all that is left is the assembly and the interface circuits.

6.1 Pseudo-Differential Comparator Topology

Since the auto-zeroed comparator designed in Chapter 5 is single ended, but the input signal and the latch described in Section 4.2 are differential, two comparators are used in parallel to get an effectively differential topology, as shown in Figure 6.1.

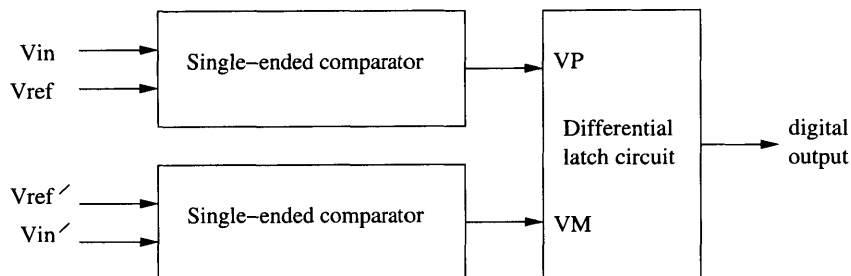


Figure 6.1: Pseudo-differential comparator topology.

When the flash converter is powered from a one volt supply, the input voltages are V_{in} and

$$V'_{in} = 1 - V_{in}$$

and the reference voltages are V_{ref} and

$$V'_{ref} = 1 - V_{ref}.$$

Two reference resistor strings are used to produce the positive and negative reference voltages. This topology allows differential operation, but maintains the advantage of using the same transistors to compare the individual inputs with their respective references.

6.2 Interface Circuitry

The interface circuitry brings the clock signals on chip, produces the necessary clock phases, and drives 50Ω termination resistors off chip.

6.2.1 Clock Phase Production

The CMOS switches in the comparators and latches need three clock phases (and two complementary clock signals). The three primary clock phases are shown in Figure 6.2. The main clock period (shown by the dashed line) is divided into quarters (shown by the dotted lines). The clock signals are named after which quarter or quarters of the clock period they are high.

The operation of the circuit is driven by these clock phases. The clock period begins with clock phase PH12 going high, which causes the comparators to go into auto-zero mode. Then clock phase PH23 goes high, and the latch begins precharging, connecting the floating inverter gates to the output of the comparator. When clock phase PH12 goes low, the comparator goes into compare

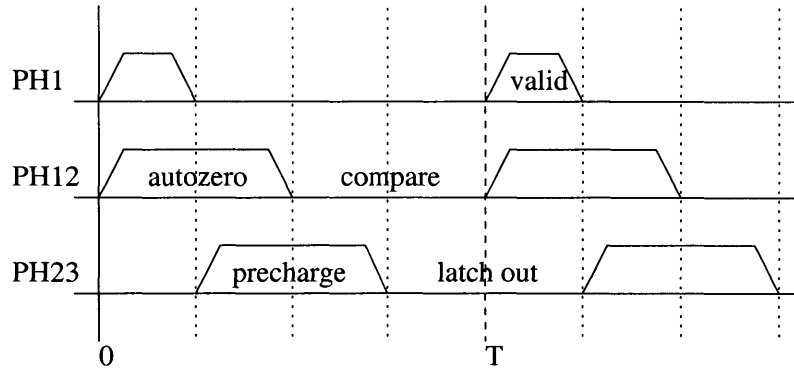


Figure 6.2: The three primary clock phases and their purposes. Phase PH12 and phase PH23 are also produced in complementary form as phases $\overline{\text{PH12}}$ and $\overline{\text{PH23}}$.

mode, and renders a decision, dragging the input gates of the latch along for the ride. When clock phase PH23 goes low, the latch goes into latch mode, amplifying and capturing the comparator decision. Finally, the next clock cycle starts, and the comparator goes back into auto-zero mode, while the latch holds the previous output. At the same time, clock phase PH1 goes high for a short time, enabling the double half-latch and carrying the thermometer code from the regenerative latch to the encoder and decoder circuitry.

Not shown in Figure 6.2, the complements of signals PH12 and PH23, phases $\overline{\text{PH12}}$ and $\overline{\text{PH23}}$, are also produced by the clock circuitry to drive the PMOS transistors in the switches of the comparators and latches.

6.2.2 Clock Input Driver

The clock waveforms are brought on to the chip as two differential sine-wave signals, in quadrature, at the sampling rate. The actual phase relationship between the input sine waves allows fine tuning of the clock phase relationships shown in Figure 6.2.

The clock inputs are buffered by the circuit shown in Figure 6.3. This circuit provides enough gain to turn a small-amplitude sine wave into the proper square waves required to drive the clock

generation logic.




Figure 6.3: Differential clock input buffer.

The clock generation logic produces the five necessary clock signals from the buffered sine waves, as shown in Figure 6.4. The timing of clocks phase PH12 and phase PH23 are determined directly by the input sine waves. Clock phase PH1 is simply the logical AND of clock phases PH12 and $\overline{\text{PH23}}$. The final inverters in the clock generation are distributed throughout the layout as sixty-three small inverter circuits, connected in parallel by metal buses, reducing clock skew.

6.2.3 50 Ω Output Drivers

Driving high-speed output signals off chip is a difficult task. The simple scheme shown in Figure 6.5 uses a fully differential current-steering design, that drives 50 Ω termination resistors and reduces transient peak currents in the supply and ground. The reference voltage V_B and the reference current I_{REF} are provided off chip, and the supply voltage V_{DD2} is independent from the main circuit supply voltage V_{DD} .




Figure 6.4: Clock phase generator and distribution buffers.




Figure 6.5: Fully differential current-steering output driver schematic. The $50\ \Omega$ resistors shown are the off-chip termination resistors.

6.3 Complete Circuit Simulations

The entire circuit was assembled and simulated in SPICE. Details of the SPICE files are shown in Appendix D. The complexity of the circuit only allowed for short simulations, but functionality, performance, and power consumption were demonstrated.

6.3.1 Functionality

The circuit was simulated at several sampling rates to demonstrate speed and functionality. The following figures show the simulation results for a sampling rate of 500 MS/s. Figure 6.6 shows the full A/D converter response to an input ramp, and Figures 6.7, 6.8, and 6.9 show the converter response to input sine waves of frequencies 2 MHz, 20 MHz, and 125 MHz, respectively. The converter begins making small, but noticeable, errors when sampling an input sine wave at the Nyquist frequency of 250 MHz, as shown in Figure 6.10. Despite these errors, this design clearly exhibits a full-power bandwidth above 250 MHz.




Figure 6.6: Reconstructed output codes of complete circuit with ramp input, from SPICE simulation.




Figure 6.7: Reconstructed output codes of complete circuit with 2 MHz sine wave input, from SPICE simulation.

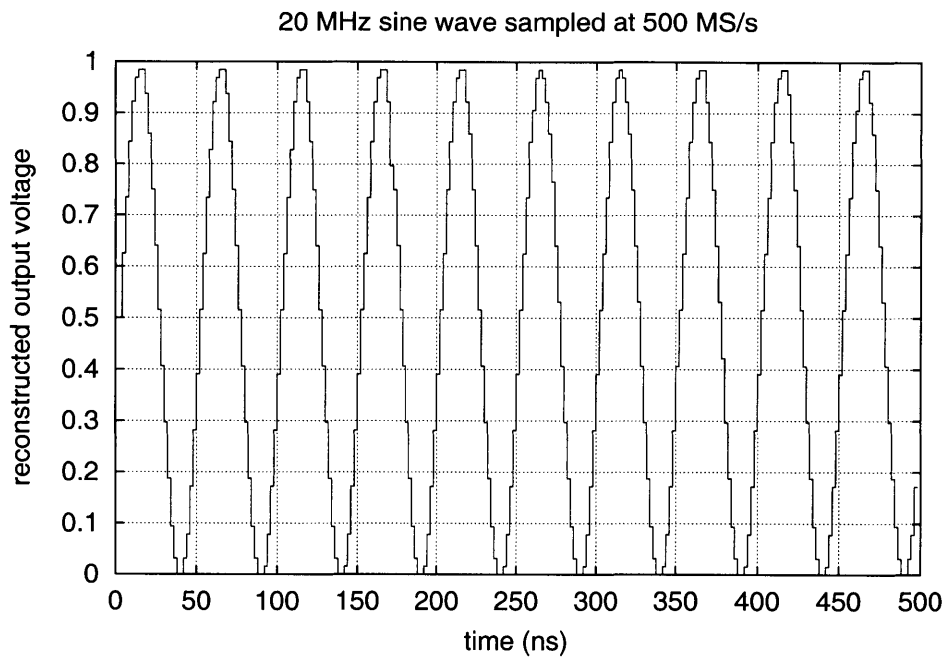


Figure 6.8: Reconstructed output codes of complete circuit with 20 MHz sine wave input, from SPICE simulation.




Figure 6.9: Reconstructed output codes of complete circuit with 125 MHz sine wave input, from SPICE simulation.




Figure 6.10: Reconstructed output codes of complete circuit with 250 MHz sine wave input, from SPICE simulation, showing adequate full-power bandwidth.

6.3.2 Power Consumption

As suggested in Section 5.3.2, most of the power is consumed by the clock driver circuitry. In simulation, the digital logic (latches, majority circuit, output ROM) consumed approximately 1 mW, and the combined power consumption of all the comparators was less than 1 mW. The clock circuitry, however, consumed 7 mW at $F_S = 100$ MHz and 14 mW at $F_S = 500$ MHz.

Significant power was also dissipated in the reference resistor strings. Since the converter design was differential, there were two resistor strings, one for each reference voltage shown in Figure 6.1. In order to keep the reference voltages accurate when loaded by the input current of the comparators, very small resistances were used. Even using $2\ \Omega$ resistors, there was significant integral nonlinearity (INL) due to “bowing” in the reference voltages, as seen in Figure 6.11. The power dissipated was

$$P = \frac{V^2}{R} = \frac{(1\ \text{V})^2}{128\ \Omega} = 8\ \text{mW}$$

With two resistor strings, the total power dissipated is 16 mW!

The total power consumed by the converter, including clock drivers, digital circuitry, the comparators, and reference resistors is 32 mW. The figure of merit is the quantization energy

$$E_Q = \frac{P}{2^B(2F_{BW})}$$

where P is the power dissipation, B is the high-frequency effective number of bits, and F_{BW} is the effective resolution bandwidth.

Using the calculated power consumption, along with assumptions for the effective number of bits (assumed to be $B = 5$) and the effective resolution bandwidth (assumed to be $F_{BW} = F_S/2 =$




Figure 6.11: Integral nonlinearity due to comparator input currents loading the reference resistor string.

250 MHz) the target quantization energy is found

$$E_Q = \frac{32 \text{ mW}}{2^5 (500 \text{ MHz})} = 2 \text{ pJ}$$

which is state of the art compared to the other converters listed in the table in Figure 2.1.

The power consumed by the output drivers was not included in this calculation, as is common practice [17]. The output drivers were powered from a separate, high-voltage supply, and were designed to drive 50Ω loads. No effort was made to reduce their power dissipation. There were seven output drivers, one for each bit and one for clock monitoring, each pushing 20 mA into 50Ω , powered from a 2 volt supply. The total power consumption of the output drivers was 280 mW.

6.4 Comparison to Bulk CMOS Implementation

The same circuit, simulated in a comparable bulk technology,¹ was only functional at a supply voltage of 2 volts and a low sampling rate of 100 MHz. The design was nonfunctional at faster sampling rates or lower supply voltages. Under these conditions, the bulk design drew 58 mA (without the reference resistors) from the 2 volt supply, consuming 116 mW.

At a sampling rate of 100 MHz, the SOI design consumed only 9 mA from a 1 volt supply (also ignoring the reference current). For a breakdown of the power consumption, see the table in Figure 6.12. The lower threshold voltages and the lower parasitic capacitances of SOI technology allow power savings of an order of magnitude. This comparison clearly shows the low power capability of SOI technology.

Circuit	1V FDSOI	2V bulk
Comparators	1 mW	7 mW
Digital Logic	1 mW	19 mW
Clock Drivers	9 mW	90 mW
A/D Total	11 mW	116 mW

Figure 6.12: Power consumption comparison for flash converter subcircuits in FDSOI and bulk CMOS technologies at 100 MHz (ignoring reference resistors).

¹Again, TSMC's 2.5-V 0.25- μm logic salicide process.

Chapter 7

Experimental Results

A layout of the designed chip was completed, and the chip was fabricated in MIT Lincoln Laboratory's 0.25- μm fully depleted silicon-on-insulator CMOS process. The total chip dimensions were 2mm by 2mm, but the actual flash converter circuit takes up a very small area (less than ten percent of the total chip real estate). Most of the chip area is taken up by bypass capacitors and test structures that allow the individual testing of all of the major building blocks of the design. The test structures were used for screening and parameter measurement. A die photo is shown in Figure 7.1, with a map of the important structures in Figure 7.2.

The parts were packaged in a forty-pin LCC (leadless chip carrier) for high-speed testing. See the bonding diagram and package pin-out in Figure 7.3.




Figure 7.1: Microphotograph of fabricated test chip.

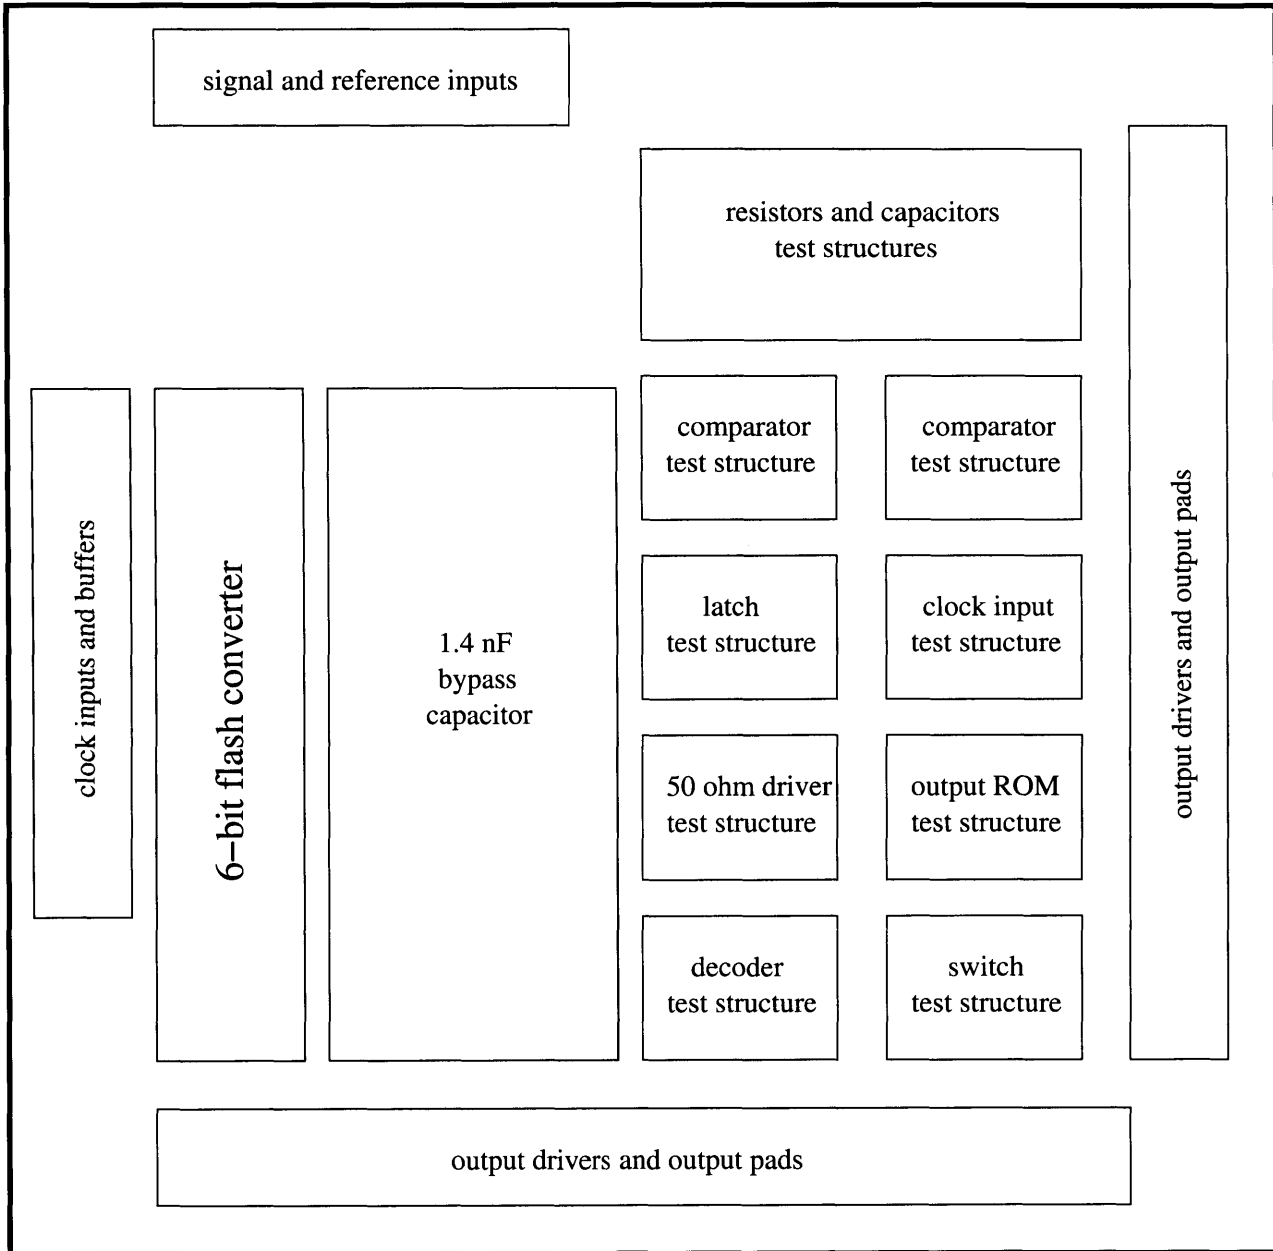


Figure 7.2: Layout map of test chip.




Figure 7.3: Bonding diagram and pin-out for packaged test chip.

7.1 Test Structures

The on-chip test structures were evaluated to screen potentially bad die, and to predict circuit performance. The sampling capacitor was fabricated out of a polysilicon-metal-metal sandwich structure and was very close to the designed value of 50 fF. The large bypass capacitors were also close to their expected value and showed no leakage. The functionality of the comparators, logic circuits, and output drivers was successfully verified at a wafer probe station. Unfortunately, high-speed testing and characterization of these test structures was not possible.

A few concerns were noted during the examinations of the test structures. The reference resistors were fabricated out of silicided polysilicon sheets. The measured resistance value was eight ohms, which is much larger than the intended two-ohm resistance.

Testing of the CMOS switch test circuits revealed that the off resistance of the switches was much lower than expected. With the common-mode voltage near ground, the off resistance was as small as $R_{SW} = 25 \text{ k}\Omega$ instead of the expected megohms. This behavior was due to leaky NMOS transistors, and ultimately limited the measured performance of the flash converter.

7.2 High-Speed Testing

Ten chips were packaged and tested¹ at a clock speed of 100 MHz. Other than a few with completely dead outputs, most of the chips were functional.

Some raw data from one of the sine-wave tests is shown in Figure 7.4, along with the least-squared-error curve fit. This figure seems to show significant noise in the converter, with frequent errors of many LSBs and several large glitches. The Fourier transform of the data shown in in Figure 7.5 shows many large harmonic spurs, contributing to poor signal-to-noise-and-distortion

¹Some details of the test setup are discussed in Appendix E.




Figure 7.4: Raw A/D converter output data from sine-wave test, with best fit sine wave for ENOB calculation.

ratio.

The calculated effective number of bits was approximately 2 bits. The effective number of bits versus input frequency for two typical parts is shown in Figure 7.6.

If the errors and glitches in Figure 7.4 were from fundamental noise sources, they should have also been seen for a constant input voltage. However, taking long histograms with constant inputs didn't show any noise for most input voltages, as shown in Figure 7.7. However, for some input voltages, the output bits oscillate, producing a histograms like that shown in Figure 7.8. This effect is strongly input level dependent, thus the errors and glitches in the output data are not due to noise.

The best performance was actually achieved from chip 2, which had a stuck-high LSB output. This chip was the first chip tested, and was badly abused during the initial test setup and hardware debugging. The chip was put into the socket wrong, and power was initially connected backwards




Figure 7.5: FFT of A/D converter output for sine-wave test, showing poor SNDR.




Figure 7.6: ENOB performance versus input frequency for two parts.




Figure 7.7: Histogram for DC input, showing no input noise.




Figure 7.8: Histogram for DC input, showing considerable input "noise."

by the author. Despite these handicaps, the low-frequency effective number of bits was 2.5 bits, still well below the expected five or six bits from the design and simulations.

7.3 Simulations of Observed Problems

The circuit was simulated again in SPICE using some of the measured parameters to identify and explain some of the observed problematic behavior.

7.3.1 Switch Resistance

From the test structures, the off resistance of the CMOS switches was found to be as small as $R_{SW} = 25 \text{ k}\Omega$. This fabrication problem had a significant impact on the observed behavior of the flash converter. When the fabricated parts are driven with a slow ramp input, large-amplitude limit cycles appear in the outputs, corrupting the output data as shown in Figure 7.9. This exact behavior was simulated. Simulating the entire circuit for a slow ramp input with leaky switches produced the output behavior shown in Figure 7.10. Note the remarkable match between the observed and simulated behavior. This behavior is the cause of the large harmonic spurs in the fast Fourier transform of Figure 7.5.

This behavior can be explained by examining the comparator auto-zero-mode and compare-mode transients in Figure 7.11. The high-going output trace during the compare mode (the second half of the clock cycle) is the designed comparator behavior, the same as in Figure 5.13. During the compare phase, the comparator produces a decisive output. The lower output trace, the one that stays near 0.5 volt, is the comparator output with leaky switches. The output transient is not decisive, and appears similar to a comparator in its metastable state. This output voltage does not successfully trigger the latches into the correct state, and will thus cause significant and random errors in the thermometer code.




Figure 7.9: Measured converter behavior with leaky switches.




Figure 7.10: Simulated converter behavior with leaky switches.




Figure 7.11: A single comparison accomplished by the single-ended auto-zeroed comparator, with leaky switches.

7.3.2 Reference Resistors

Another problem with the tested converters was significant nonlinearities that also may have led to the large harmonic spurs observed in Figure 7.5. Some of these spurs are due to the limit cycles in the output bits. Another potential source of nonlinearity in the converter is the “bowing” of the reference resistor string due to large input currents into the comparators. Although the limit-cycle problem shown in Figure 7.9 prevented a direct measurement of integral nonlinearity (INL), the converter was simulated for a number of reference resistor sizes, shown in Figure 7.12. The fabricated eight-ohm resistors create a peak INL of over 2 LSBs. The designed two-ohm resistors exhibit a peak INL of just 1 LSB. In order to make the INL significantly smaller, the resistor size needs to be decreased significantly (causing much increased power consumption).




Figure 7.12: INL versus input code for several different reference resistor sizes.

7.4 Other Observations

Further testing of the test structures revealed that the off resistance of the CMOS switches improved for negative wafer voltages, $V_W < -4$ V. Changing the wafer voltage changes the effective threshold voltages of the transistors through the back-gate effect. Unfortunately, negative wafer voltages, which make the NMOS transistors less leaky, also make the PMOS transistors more leaky, preventing any systematic improvement. The test chips were not functional for $V_W < -3$ V and no improvement was observed for smaller wafer bias.

The test chips were not functional for clock speeds greater than $F_S > 200$ MHz. It is likely that the PMOS transistors in the clock drivers were not able to overcome the leakage of the NMOS transistors and drive the clock lines at full speed. The best results were achieved with a sampling rate of $F_S = 100$ MHz.

The chips tested also failed for supply voltages $V_{DD} > 1.3$ V. This voltage is just above the

bandgap voltage, which points to the kink effect as culprit in this failure.

7.5 Quantization Energy

To calculate the quantization energy, the operating power consumption was measured. For the designed supply voltage of $V_{DD} = 1$ V, the current draw on the supply was 30 mA at a clock frequency of 100 MHz. The chips tested did not work at higher supply voltages.

Thus, the power consumed by the circuit is $P = 30$ mW. Although this figure is remarkable closed to the total simulated current found in Section 6.3.2, this comparison hides a number of small inconsistencies. The total power dissipated in the reference resistor strings was supposed to be 16 mW, but from the measured resistance of the test structure, the fabricated resistor strings dissipated only 4 mW. Thus, the simulated circuit power consumption (subtracting out the resistors) was 16 mW, while the measured circuit power consumption was 26 mW. It is assumed that the additional power was consumed by leaky NMOS transistors and the clock driver circuitry, driving additional unmodeled parasitic capacitances.

The quantization energy for the converter is calculated as

$$E_Q = \frac{P}{2^B(2F_{BW})} = \frac{30 \text{ mW}}{2^2(100 \text{ MHz})} = 75 \text{ pJ}$$

This result is significantly higher than expected.

Chapter 8

Conclusions

This thesis presented an ultra-low-power flash A/D converter design that leverages the advantages of fully depleted SOI technology for high-speed, low-power analog circuit design.

8.1 Summary of Results

Using a comparator topology optimized for fully depleted SOI technology, the design of a one-volt 500 MS/s flash A/D converter with a simulated state-of-the-art quantization energy of $E_Q = 2$ pJ was completed. Simulations of the design show that it meets the required specifications of low power and high speed. This figure of merit is the lowest reported figure of merit in the entire class of high-speed, low-resolution A/D converters.

The fabricated parts were functional and exhibited low power consumption, but the leaky switch resistances caused very poor signal-to-noise-and-distortion ratio. Despite the poor performance of the tested parts, the potential for ultra-low-power analog-to-digital converters fabricated in fully depleted SOI technology was demonstrated.

8.2 Future Directions

There are a number of improvements that could be explored in the next generation of ultra-low-power flash A/D converters in SOI technology.

8.2.1 Second Fabrication

While the low power consumption of this design was successfully demonstrated, a second fabrication run with improved switches would be necessary to demonstrate the potential accuracy of the technology. Also, the resistors in the reference voltage string should be resized to improve the integral nonlinearity.

8.2.2 Improved Test Structures

Overall, testing of the fabricated parts would have been improved by better test structures and individual circuits. Some of the test structures were designed to interface with a specific piece of test equipment, which was ultimately not used. High-speed testing and characterization of the comparators, logic circuits, and latch as individual circuits were therefore not possible.

8.2.3 Improved Design

The comparator design was carefully optimized to leverage the advantages of SOI technology, however improved SOI circuit techniques could be explored in other circuit components.

Considerable effort was put into the design of the comparator, while comparatively little effort was put into the design of the latch. Some simulations showed clock feedthrough on the output of the half-latches, and the regenerative latch could benefit from a reset scheme and appropriate additional clock phase. For a higher sampling-rate converter, an improved latch design would probably be necessary.

Considerable power was dissipated in the two reference resistor strings. Two strings were used to simplify layout and signal routing. Clever layout and signal routing may be able to eliminate one of the strings, saving twenty-five percent of the power consumption.

No attempt was made to reduce the power consumption of the digital logic. Although the digital logic only accounted for a small fraction of the total power dissipation, many low-power digital logic families exist, and improvements could be made.

As with all circuits that need to drive signals off chip, the output drivers consumed significant power. In this design, their power consumption was simply ignored in the calculations. Improved total performance would be achieved with better output drivers. Standard low-voltage differential signaling (LVDS) would be an attractive option.

Appendix A

Analog-to-Digital Converter Testing

Analog-to-digital converters are essential building blocks in modern electronic systems. They form the critical link between front-end analog transducers and back-end digital computers that can efficiently implement a wide variety of signal-processing functions. The wide variety of digital-signal-processing applications leads to the availability of a wide variety of analog-to-digital (A/D) converters of varying price, performance, and quality.

Ideally, an A/D converter encodes a continuous-time analog input voltage, V_{IN} , into a series of discrete N -bit digital words that satisfy the relation

$$V_{IN} = V_{FS} \sum_{k=0}^{N-1} \frac{b_k}{2^{k+1}} + \epsilon$$

where V_{FS} is the full-scale voltage, b_k are the individual output bits, and ϵ is the quantization error. This relation can also be written in terms of the least significant bit (LSB) or quantum voltage level

$$V_Q = \frac{V_{FS}}{2^N} = 1 \text{ LSB}$$

as

$$V_{IN} = V_Q \sum_{k=0}^{N-1} b_k 2^k + \epsilon$$

A plot of this ideal characteristic for a 3-bit A/D converter is shown in Figure A.1.




Figure A.1: Ideal analog-to-digital converter characteristic.

Each unique digital code corresponds to a small range of analog input voltages. This range is 1 LSB wide (the “code width”) and is centered around the “code center.” All input voltages resolve to the digital code of the nearest code center. The difference between the analog input voltage and the corresponding voltage of the nearest code center (the difference between the solid and dashed lines in Figure A.1) is the quantization error. Since the A/D converter has a finite number of output bits, even an ideal A/D converter produces some quantization error with every sample.

A.1 A/D Converter Figures of Merit

The number of output bits from an analog-to-digital converter do not fully specify its behavior. Real A/D converters can differ from ideal behavior in many ways. While static imperfections, such as gain and offset, are easy to quantify, the success of many signal-processing applications depends on the dynamic behavior of the A/D converter. Ultimately, the application determines the requirements, and A/D converter resolution may not be either necessary or sufficient to specify the required performance. In many cases, the quality of the A/D converter must be tested for the specific application.

The wide variety of analog-to-digital converter applications leads to a large number of figures of merit for specifying performance. These figures of merit include accuracy, resolution, dynamic range, offset, gain, differential nonlinearity, integral nonlinearity, signal-to-noise ratio, signal-to-noise-and-distortion ratio, effective number of bits, spurious-free dynamic range, intermodulation distortion, total harmonic distortion, effective resolution bandwidth, full-power bandwidth, full-linear bandwidth, aperture delay, aperture jitter, transient response, and overvoltage recovery.

These specifications can be loosely divided into three categories — static parameters, frequency-domain dynamic parameters, and time-domain dynamic parameters — and are defined in this section.

A.1.1 Static Parameters

Static parameters are the A/D converter specifications that can be tested at low speed, or even with constant voltages. These specifications include accuracy, resolution, dynamic range, offset, gain, differential nonlinearity, and integral nonlinearity.

Accuracy

Accuracy is the total error with which the A/D converter can convert a known voltage, including the effects of quantization error, gain error, offset error, and nonlinearities. Technically, accuracy should be traceable to known standards (for example, NIST), and is generally a “catch-all” term for all static errors.

Resolution

Resolution is the number of bits, N , out of the A/D converter. The characteristic in Figure A.1 shows a 3-bit A/D converter. Probably the most noticeable specification, resolution determines the size of the least significant bit, and thus determines the dynamic range, the code widths, and the quantization error.

Dynamic Range

Dynamic range is the ratio of the smallest possible output (the least significant bit or quantum voltage) to the largest possible output (full-scale voltage), mathematically $20 \log_{10} 2^N \approx 6N$.




Figure A.2: Analog-to-digital converter characteristic, showing offset and gain error.

Offset Error

Offset error is the deviation in the A/D converter's behavior at zero. The first transition voltage should be 1/2 LSB above analog ground. Offset error is the deviation of the actual transition voltage from the ideal 1/2 LSB. Offset error is easily trimmed by calibration. Compare the location of the first transitions in Figures A.1 and A.2.

Gain Error

Gain error is the deviation in the slope of the line through the A/D converter's end points at zero and full scale from the ideal slope of $2^N/V_{FS}$ codes-per-volt. Like offset error, gain error is easily corrected by calibration. Compare the slope of the dashed lines in Figures A.1 and A.2.




Figure A.3: Analog-to-digital converter characteristic, showing nonlinearity errors and a missing code. The dashed line is the ideal characteristic, and the dotted line is the best fit.

Differential Nonlinearity

Differential nonlinearity (DNL) is the deviation of the code transition widths from the ideal width of 1 LSB. All code widths in the ideal A/D converter are 1 LSB wide, so the DNL would be zero

everywhere. Some datasheets list only the “maximum DNL.” See the wide range of code widths illustrated in Figure A.3.

Integral Nonlinearity

Integral nonlinearity (INL) is the distance of the code centers in the A/D converter characteristic from the ideal line. If all code centers land on the ideal line, the INL is zero everywhere. Some datasheets list only the “maximum INL.” See the deviations of the code centers from the ideal line in Figure A.3.

Note that there are two possible ways to express maximum INL, depending on your definition of the “ideal line.” Datasheet INL numbers can be decreased by quoting maximum INL from a “best fit” line instead of the ideal line. In Figure A.3, the ideal line (shown dashed) exhibits an INL of 1 LSB, while the best-fit line (shown dotted) exhibits an INL of half that size. While this prevarication underestimates the effect of INL on total accuracy, it is probably a better reflection of the A/D converter’s linearity in AC-coupled applications.

Missing Codes

Missing codes are output digital codes that are not produced for any input voltage, usually due to large DNL. In some converters, missing codes can be caused by non-monotonicity of the internal D/A. The large DNL in Figure A.3 causes code 100 to be “crowded out.”

A.1.2 Resolution and Quantization Noise

Quantization error due to the finite resolution N of the A/D converter limits the signal-to-noise ratio. Even in an ideal A/D converter, the quantization of the input signal creates errors that behave like noise. All inputs within $\pm 1/2$ LSB of a code center resolve to that digital code. Thus, there will be a small difference between the code center and the actual input voltage due to this

quantization. If assume that this error voltage is uncorrelated and distributed uniformly, we can calculate the expected rms value of this “quantization noise.”




Figure A.4: Quantization error voltage for ideal analog-to-digital converter.

The signal-to-noise ratio due to quantization can be directly calculated. The range of the error voltage is the quantum voltage level (the least significant bit)

$$V_Q = \frac{V_{FS}}{2^N} = 1 \text{ LSB}$$

Finite amplitude resolution introduces a quantization error between the analog input voltage and the reconstructed output voltage. Assuming this quantization error voltage is uniformly distributed over the code width from $-1/2$ LSB to $+1/2$ LSB, the expectation value of the error voltage is

$$E\{\epsilon^2\} = \frac{1}{V_Q} \int_{-\frac{1}{2}V_Q}^{+\frac{1}{2}V_Q} \epsilon^2 d\epsilon = \frac{1}{V_Q} \left[\frac{\epsilon^3}{3} \right]_{-\frac{1}{2}V_Q}^{+\frac{1}{2}V_Q} = \frac{V_Q^2}{12}$$

This quantization noise is assumed to be uncorrelated and broadband. Using this result, the maximum signal-to-noise ratio for a full-scale input can be calculated. The rms value of a full-scale

peak-to-peak amplitude V_F is

$$V_{\text{rms}} = \frac{V_{FS}}{2\sqrt{2}} = \frac{2^N V_Q}{2\sqrt{2}}$$

thus the signal-to-noise ratio is

$$\text{SNR} = 20 \log \left(\frac{V_{\text{rms}}}{\sqrt{E(\epsilon^2)}} \right) = 20 \log(2^N \sqrt{1.5}) = 6.02N + 1.76 \text{ dB}$$

when the noise is due only to quantization. Using this result, we can tabulate the ideal signal-to-noise ratio for several A/D converter resolutions in Figure A.5.

resolution	signal-to-noise ratio
6 bits	37.9 dB
8 bits	49.9 dB
10 bits	62.0 dB
12 bits	74.0 dB
14 bits	86.0 dB
16 bits	98.1 dB

Figure A.5: Ideal signal-to-noise ratio due to quantization versus resolution.

Figure A.6 shows the fast Fourier transform (FFT) spectrum of a sine wave sampled by an ideal 10-bit A/D converter. The noise floor in this figure is due only to quantization noise. For an M -point FFT, the average value of the noise contained in each frequency bin is $10 \log_2(M/2)$ dB below the rms value of the quantization noise. This “FFT process gain” is why the noise floor appears 30 dB lower than the quantization noise level listed in Figure A.5.

A.1.3 Frequency-Domain Dynamic Parameters

All real analog-to-digital converters have additional noise sources and distortion processes that degrade the performance of the A/D converter from the ideal signal-to-noise ratio calculated above. These imperfections in the dynamic behavior of the A/D converter are quantified and reported in

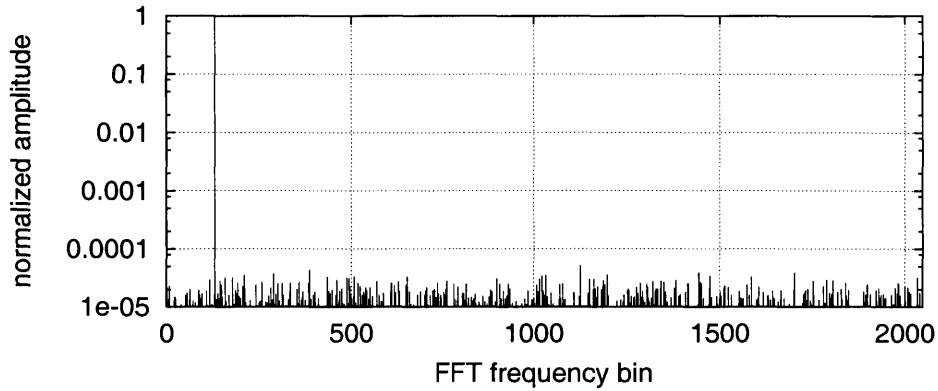


Figure A.6: Quantization noise floor for an ideal 10-bit A/D converter (4096 point FFT).

a variety of ways.

Signal-to-Noise-and-Distortion Ratio

Signal-to-noise-and-distortion ratio (S/N+D, SINAD, or SNDR) is the ratio of the input signal amplitude to the rms sum of all other spectral components. For an M -point FFT of a sine wave test, if the fundamental is in frequency bin m (with amplitude A_m), the SNDR can be calculated from the FFT amplitudes

$$\text{SNDR} = 10 \log \left[A_m^2 \left(\sum_{k=1}^{m-1} A_k^2 + \sum_{k=m+1}^{M/2} A_k^2 \right)^{-1} \right]$$

To avoid any spectral leakage around the fundamental, often several bins around the fundamental are ignored. The SNDR is dependent on the input-signal frequency and amplitude, degrading at high frequency and power. Measured results are often presented in plots of SNDR versus frequency for a constant-amplitude input, or SNDR versus amplitude for a constant-frequency input.

Effective Number of Bits

Effective number of bits (ENOB) is simply the signal-to-noise-and-distortion ratio expressed in bits rather than decibels by solving the “ideal SNR” equation

$$\text{SNR} = 6.02N + 1.76 \text{ dB}$$

for the number of bits N , using the measured SNDR

$$\text{ENOB} = \frac{\text{SNDR} - 1.76 \text{ dB}}{6.02 \text{ dB/bit}}$$

In the presentation of measured results, ENOB is identical to SNDR, with a change in the scaling of the vertical axis.

Spurious-Free Dynamic Range

Spurious-free dynamic range (SFDR) is the ratio of the input signal to the peak spurious or peak harmonic component. Spurs can be created at harmonics of the input frequency due to nonlinearities in the A/D converter, or at subharmonics of the sampling frequency due to mismatch or clock coupling in the circuit. The SFDR of an A/D converter can be larger than the SNDR. Measurement of SFDR can be facilitated by increasing the number of FFT points or by averaging several data sets. In both cases, the noise floor will improve, while the amplitude of the spurs will stay constant. The spectrum of an A/D converter with significant harmonic spurs is shown in Figure A.7. Because SFDR is often slew-rate dependent, it will be a function of input frequency and magnitude [62]. The maximum SFDR often occurs at an amplitude below full scale.




Figure A.7: A/D converter with significant nonlinearity, showing poor SFDR and THD. Note that high frequency harmonics of the input signal are aliased down to frequencies below the Nyquist frequency.

Total Harmonic Distortion

Total harmonic distortion (THD) is the ratio of the rms sum of the first five harmonic components (or their aliased versions, as in Figure A.7) to the input signal

$$\text{THD} = 10 \log \left(\frac{V_2^2 + V_3^2 + V_4^2 + V_5^2 + V_6^2}{V_1^2} \right)$$

where V_1 is the amplitude of the fundamental, and V_n is the amplitude of the n -th harmonic.

Intermodulation Distortion

Intermodulation distortion (IMD) is the ratio of the amplitudes of the sum and difference frequencies to the input signals for a two-tone test, sometimes expressed as “intermod-free dynamic range (IFDR)” See the FFT spectrum in Figure A.8. For second-order distortion, the IMD would be

$$\text{IMD} = 10 \log \left(\frac{V_+^2 + V_-^2}{V_1^2 + V_2^2} \right)$$

where V_1 and V_2 are the rms amplitudes of the input signals, and V_+ and V_- are the rms amplitudes of the sum and difference intermodulation products. See Figure A.8.




Figure A.8: Two-tone IMD test with second-order nonlinearity, showing (from left) $F_2 - F_1$ product, F_1 input, F_2 input, $2F_1$ product, $F_1 + F_2$ product, and $2F_2$ product.

Effective Resolution Bandwidth

Effective resolution bandwidth (ERBW) is the input-signal frequency where the SNDR of the A/D converter has fallen by 3 dB (0.5 bit) from its value for low-frequency input signals.

Full-Power Bandwidth

In amplifiers, full-power bandwidth (FPBW) is the maximum frequency at which the amplifier can reproduce a full-scale sinusoidal output without distortion (sometimes calculated as slew-rate divided by $\pm 2\pi V_{\max}$), or where the amplitude of full-scale sinusoid is reduced by 3 dB. Using this definition for A/D converters can result in optimistic frequencies where the SNDR is severely degraded. Some manufacturer report the full-power bandwidth as the frequency where the amplitude of the *reconstructed* input signal is reduced by 3 dB [63].

Full-Linear Bandwidth

Full-linear bandwidth is the frequency where the slew-rate limit of the input sample-and-hold begins distorting the input signal by some specified amount ([64] uses 0.1 dB)

A.1.4 Time-Domain Dynamic Parameters

Aperture Delay

Aperture delay is the delay from when the A/D converter is triggered (perhaps the rising edge of the sampling clock) to when it actually converts the input voltage into the appropriate digital code.

Aperture delay is also sometimes called aperture time.

Aperture Jitter

Aperture jitter is the sample-to-sample variation in the aperture delay. The rms voltage error caused by rms aperture jitter decreases the overall signal-to-noise ratio, and is a significant limiting factor in the performance of high-speed A/D converters [3].




Figure A.9: Effects of aperture jitter.

If we assume that the input waveform is a sinusoid

$$V_{IN} = V_{FS} \sin \omega t$$

then the maximum slope of the input waveform is

$$\left. \frac{dV_{IN}}{dt} \right|_{\max} = \omega V_{FS}$$

which occurs at the zero crossings. If there is an rms error in the time at which we sample (aperture jitter, t_a) during this maximum slope, then there will be an rms voltage error of

$$V_{\text{rms}} = \omega V_{FS} t_a = 2\pi f V_{FS} t_a$$

Since the aperture time variations are random, these voltage errors will behave like a random noise source. Thus the signal-to-jitter-noise ratio

$$\text{SJNR} = 20 \log \left(\frac{V_{FS}}{V_{\text{rms}}} \right) = 20 \log \left(\frac{1}{2\pi f t_a} \right)$$

The SJNR for several values of the jitter t_a is shown in Figure A.9.

Transient Response

Transient response is the settling time for the A/D converter to full accuracy (to within $\pm 1/2$ LSB) after a step in input voltage from zero to full scale

Overvoltage Recovery

Overvoltage recovery is the settling time for the A/D converter to full accuracy after a step in input voltage from outside the full scale voltage (for example, from $1.5V_F$ to $0.5V_F$)

A.2 Dynamic A/D Converter Testing Methods

A wide variety of tests have been developed to measure dynamic specifications. Many of these tests rely on Fourier analysis using the discrete Fourier transform (DFT) and the fast Fourier transform (FFT), as well as other mathematical models.

A.2.1 Testing with the Fast Fourier Transform

The simplest frequency-domain tests use the direct application of the fast Fourier transform. Taking the FFT of the output data while driving the A/D converter with a single, low distortion sine wave, the SNDR, ENOB, SFDR, and THD can easily be calculated. It is useful to take these measurements at several input amplitudes and frequencies, and plot the results. Taking data for high input frequencies allows the full-power and full-linear bandwidths to be calculated.

Two more tests are completed while driving the A/D converter with an input composed of two sine waves of different frequencies. The FFT of this test result is used to calculate the IMD (for second-order and third-order products) and the two-tone SFDR.

A.2.2 Signal Coherence for FFT Tests

One potential problem in using Fourier analysis is solved by paying close attention to the coherence of the sampled waveform in the data record. Unless the sampled data contains a whole number of periods of the input waveform, spectral leakage of the input frequency can obscure the results [65]. Figure A.10 shows the FFT of a 4096-point data record containing 127.5 periods of a perfect sine wave. This spectrum should contain a single impulse in frequency, but the incomplete cycle of the sine wave at the end of the data record causes the spectrum to be broadly smeared.¹

In addition, the number of periods of the input waveform in the sample record should not be

¹This smearing can be solved by windowing [66], but it is easier to simply use a whole number of cycles.




Figure A.10: FFT of non-integer number of cycles (4096 points, 127.5 cycles).

a non-prime integer sub-multiple of the record length. For example, using a power-of-two for both the number of periods and the number of samples results in repetitive data. Figure A.11 shows the FFT of a 4096-point data record that contains 128 periods of a sine wave. In this case, the first 32 samples of the data record are simply repeated 128 times. This input vector does a poor job of exercising the device under test, and the quantization noise is concentrated in the harmonics of the input frequency rather than uniformly distributed across the Nyquist bandwidth [67].




Figure A.11: FFT of even divisor number of cycles (4096 points, 128 cycles).

Figure A.12 shows the FFT of a 4096-point data record that contains 127 periods of a perfect sine wave. Since 127 is both odd and prime, there are no common factors between the number of

input periods and the number of samples in the data record. Because of these choices, there are no repeating patterns and every sample in the data record is unique. This FFT correctly shows the frequency impulse due to the input sine wave and the noise floor due to quantization, without any measurement-induced artifacts.




Figure A.12: FFT of non-divisor prime number of cycles (4096 points, 127 cycles).

A.2.3 Histogram Test for Linearity

The linearity (INL and DNL) of the A/D converter can be determined with a histogram test [68, 69]. A histogram of output digital codes is recorded for a large number of samples for an input sine wave. The results are compared to the number of samples expected from the theoretical sine-wave probability-density function. If the input sine wave is

$$V = A \sin \omega t$$

then the probability-density function is

$$p(V) = \frac{1}{\pi \sqrt{A^2 - V^2}}.$$

The probability of a sample being in the range (V_a, V_b) is found by integration

$$P(V_a, V_b) = \int_{V_b}^{V_a} p(V) dV = \frac{1}{\pi} \left(\arcsin \frac{V_b}{A} - \arcsin \frac{V_a}{A} \right)$$

The difference between the measured and expected probability of observing a specific output code is a function of that specific code width, which can be used to calculate the differential nonlinearity.

Due to the statistical nature of this test, the length of the data record for the histogram test can be quite large [69]. The number of samples, M , required for a high-confidence measurement is

$$M = \frac{\pi 2^{N-1} Z_{\alpha/2}^2}{\beta^2}$$

where N is the number of bits, β is the DNL-measurement resolution, and for a 99% confidence

level, $Z_{\alpha/2} = 2.576$. Thus for a 6-bit A/D converter, and a DNL-measurement resolution of 0.1 LSB

$$M = \frac{\pi 2^{N-1} Z_{\alpha/2}^2}{\beta^2} = \frac{\pi 2^5 (2.576)^2}{(0.1)^2} = 67,000$$

If the sample size is large enough, this test will work with an asynchronous input signal, however a coherent input signal with unique samples (odd and prime frequency ratio, as explained above) is also possible. Sampling an input signal harmonically related to the sampling frequency would create the same problems that occur with the FFT tests. Figure A.13 shows the histogram for a 6-bit A/D converter using an insufficient number of points, while Figure A.14 shows a properly smooth histogram result from using enough data.




Figure A.13: Histogram of 1000 points for 6-bit A/D converter (insufficient).

Once the histogram data is collected, the experimental code widths are from the measured data H_k using the expected probability [65]

$$P_k = \frac{1}{\pi} \left[\arcsin \left(\frac{(k+1)V_Q}{A} \right) - \arcsin \left(\frac{kV_Q}{A} \right) \right]$$




Figure A.14: Histogram of 100,000 points for 6-bit A/D converter.

The differential nonlinearity is thus

$$\text{DNL}_k = 1 \text{ LSB} \left(\frac{H_k}{P_k} - 1 \right)$$

However, this method is sensitive to errors in the measurement of the input sine wave amplitude A .

A better method is to calculate the INL and DNL from a cumulative histogram [69]. First, the offset is found by equating the number of positive samples M_p and the number of negative samples M_n

$$M_n = \sum_{k=1}^{2^{N-1}} H_k \quad M_p = \sum_{k=2^{N-1}+1}^{2^N} H_k$$

$$V_{OS} = \frac{A\pi}{2} \sin \left(\frac{M_p - M_n}{M_p + M_n} \right)$$

Second, the transition voltages are found from

$$V_j = -A \cos \left(\frac{\pi}{M} \sum_{k=0}^j H_k \right)$$

Once the transition voltages are known, the INL and DNL follow

$$\text{INL}_j = \frac{V_j - V_1}{1 \text{ LSB}} \quad \text{DNL}_j = \frac{V_{j+1} - V_j}{1 \text{ LSB}}$$

This method makes the amplitude A a linear factor in the calculations, which reduces the sensitivity to errors in A , and makes the final result easily normalizable.

A.2.4 Sine Wave Curve Fit for ENOB

Another way to calculate the effective number of bits is by using a sine wave curve fit [67]. A least-squared-error sine wave is fit to the measured data, and compared to the input waveform. The resulting rms error of the curve fit, E , is a measure of the effective number of bits lost due to A/D converter error sources. The effective number of bits (ENOB) can be calculated from

$$\text{ENOB} = N - \log_2 \left(\frac{E_{\text{rms}}}{V_Q/\sqrt{12}} \right)$$

Since the frequency of the input and the output must be the same, only the amplitude, offset, and phase of the output sinusoid must be found.

General Form

The fixed-frequency sine-wave curve-fit method [70] is used to find the best fit. The sinusoid to fit to the output data is

$$x_n = A \cos(\omega t_n) + B \sin(\omega t_n) + C$$

where t_n are the sample times, ω is the input frequency (in radians/second) and A , B , and C are the fit parameters. The squared error between the samples y_n and the curve fit is

$$E = \sum_{k=1}^M (y_k - x_k)^2 = \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C]^2$$

The error is minimized by setting the partial derivatives with respect to the fit parameters to zero

$$\begin{aligned} 0 = \frac{\partial E}{\partial A} &= -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C] \cos(\omega t_k) \\ 0 = \frac{\partial E}{\partial B} &= -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C] \sin(\omega t_k) \end{aligned}$$

$$0 = \frac{\partial E}{\partial C} = -2 \sum_{k=1}^M [y_k - A \cos(\omega t_k) - B \sin(\omega t_k) - C]$$

Defining $\alpha_k = \cos(\omega t_k)$ and $\beta_k = \sin(\omega t_k)$ and rearranging terms gives a set of linear equations

$$\begin{aligned} \sum_{k=1}^M y_k \alpha_k &= A \sum_{k=1}^M \alpha_k^2 + B \sum_{k=1}^M \alpha_k \beta_k + C \sum_{k=1}^M \alpha_k \\ \sum_{k=1}^M y_k \beta_k &= A \sum_{k=1}^M \alpha_k \beta_k + B \sum_{k=1}^M \beta_k^2 + C \sum_{k=1}^M \beta_k \\ \sum_{k=1}^M y_k &= A \sum_{k=1}^M \alpha_k + B \sum_{k=1}^M \beta_k + CM \end{aligned}$$

These equations can be expressed as a single linear equation $Y = UX$

$$\sum_{k=1}^M \begin{bmatrix} y_k \alpha_k \\ y_k \beta_k \\ y_k \end{bmatrix} = \left(\sum_{k=1}^M \begin{bmatrix} \alpha_k^2 & \alpha_k \beta_k & \alpha_k \\ \alpha_k \beta_k & \beta_k^2 & \beta_k \\ \alpha_k & \beta_k & 1 \end{bmatrix} \right) \begin{bmatrix} A \\ B \\ C \end{bmatrix}$$

This linear equation has a solution $X = U^{-1}Y$. The total squared error from above is

$$E = \sum_{k=1}^M [y_k - A\alpha_k - B\beta_k - C]^2$$

and can be rewritten as

$$E = \sum_{k=1}^M [y_k^2 - 2Ay_k\alpha_k - 2By_k\beta_k - 2Cy_k + A^2\alpha_k^2 + 2AB\alpha_k\beta_k + 2AC\alpha_k + B^2\beta_k^2 + 2BC\beta_k + C^2]$$

$$E = \left(\sum_{k=1}^M y_k^2 \right) - 2X^T Y + X^T U X$$

The rms error of the curve fit is then

$$E_{\text{rms}} = \sqrt{\frac{E}{M}}$$

So the effective number of bits (ENOB) is

$$\text{ENOB} = N - \frac{1}{2} \log_2 \left(\frac{12E}{V_Q^2 M} \right)$$

Simplified Coherent Form

One advantage of the curve-fit method of finding the effective number of bits is, unlike the methods that use the FFT and SNDR, it does not require windowing or coherent sampling. However, if coherent sampling is used, such that the expected data record has an integer number of sine-wave cycles, the linear equation above simplifies significantly. The sums over α_k , β_k , and $\alpha_k \beta_k$ vanish, and

$$\sum_{k=1}^M \alpha_k^2 = \sum_{k=1}^M \beta_k^2 = \frac{M}{2}$$

The linear equation becomes

$$\sum_{k=1}^M \begin{bmatrix} y_k \alpha_k \\ y_k \beta_k \\ y_k \end{bmatrix} = \begin{bmatrix} M/2 & 0 & 0 \\ 0 & M/2 & 0 \\ 0 & 0 & M \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix}$$

which has simple, closed-form solutions

$$A = \frac{2}{M} \sum_{k=1}^M y_k \alpha_k \quad B = \frac{2}{M} \sum_{k=1}^M y_k \beta_k \quad C = \frac{1}{M} \sum_{k=1}^M y_k$$

The least squared error simplifies to

$$\begin{aligned}
E &= \sum_{k=1}^M \left[y_k^2 - 2Ay_k\alpha_k - 2By_k\beta_k - 2Cy_k + A^2\alpha_k^2 + 2AB\alpha_k\beta_k + 2AC\alpha_k + B^2\beta_k^2 + 2BC\beta_k + C^2 \right] \\
&= \left(\sum_{k=1}^M y_k^2 \right) - MA^2 - MB^2 - 2MC^2 + \frac{M}{2}A^2 + 0 + 0 + \frac{M}{2}B^2 + 0 + MC^2 \\
&= \left(\sum_{k=1}^M y_k^2 \right) - \frac{M}{2}(A^2 + B^2 + 2C^2)
\end{aligned}$$

The effective number of bits (ENOB) remains

$$\text{ENOB} = N - \frac{1}{2} \log_2 \left(\frac{12E}{V_Q^2 M} \right)$$

Comparison of Methods

Clearly, calculating the effective number of bits using an FFT and the signal-to-noise-and-distortion ratio should produce the same result as the sine-wave curve-fit method. Figure A.15 shows part of a 4096-point quantized sine wave with added noise, plotted with the best-fit sine wave, calculated as described above. The effective number of bits from the curve fit is 3.02. Figure A.16 shows the FFT of the same data. The effective number of bits calculated from SNDR is 3.01, which matches nicely with the curve-fit result.




Figure A.15: Partial data plot showing sine wave and added noise, along with best-fit sine wave. ENOB calculated from curve fit is 3.01.




Figure A.16: Partial FFT showing sine wave and added noise. ENOB calculated from SNDR is 3.02.

A.2.5 Overall Noise and Aperture Jitter

Overall noise of an A/D converter can be measured by grounding the input of the A/D converter (in a unipolar A/D converter, the input should be connected to a mid-scale constant voltage source) and accumulating a histogram. Only the center code bin should have counts in it. Any spread in the histogram around the center code bin is caused by noise in the A/D converter. This test can also be used to determine the offset of the A/D converter as in the histogram test,

Aperture jitter is measured by repeatedly sampling the same voltage of the input waveform. For example, a sine wave input is used, and the A/D converter is triggered to repeatedly sample the positive-slope zero crossing. If the input sine wave and the sampling clock are generated from phase-locked sources, there should be no spread in the output digital codes from this measurement. However, a real A/D converter will produce a spread in output codes due to aperture jitter.

The aperture jitter is calculated from a histogram of output codes produced from this measurement. For an input sine wave sampling at the zero crossings, the aperture jitter is

$$t_a = \frac{V_{\text{rms}}}{2\pi f A}$$

where A is the amplitude and f is the frequency of the input sine wave. As the amplitude of the input increases, the slope at the zero crossings increases, and the spread of output codes should proportionally increase due to aperture jitter.

Appendix B

Noise Sources in Bulk CMOS

The noise behavior of bulk CMOS devices is dominated primarily by two noise sources: thermal noise and flicker ($1/f$) noise. Other sources that are sometimes present in the noise spectrum are shot noise, generation/recombination noise, and “popcorn” noise. Of these sources, thermal noise and shot noise are physically fundamental to the operation of the device and are always present. The quality of the manufactured device (the number of defects in the bulk silicon, in the gate oxide, and in the various interfaces) determines the level of generation/recombination noise and popcorn noise. It is probable that flicker noise appears through both quality-dependent and fundamental noise processes.

B.1 Noise Notation

Following van der Ziel’s [71] lead, common practice in the literature uses two-level subscripts, for example, S_{I_D} to denote the spectral density of the drain current. Unfortunately, this author can’t stand the sight of two-level subscripts, so he reluctantly uses a different notation for noise variables.

In IEEE notation, as described in the Standard [72], variables are named with the following convention:

- Incremental small signal quantities are expressed in all lower case (i_d, v_{ds})
- Quiescent large signal quantities are expressed in all upper case (I_D, V_{DS})
- Instantaneous total values are expressed as a lower case variable with an upper case subscript ($i_D = I_D + i_d, v_{DS} = V_{DS} + v_{ds}$)
- Root-mean-square values are expressed as an upper case variable with a lower case subscript (I_d, V_{ds})

Note that all three editions of Gray and Meyer [73, page 31] label the total variable incorrectly (incorrectly using the form reserved for RMS values¹). Roberge [74, page 3] uses the upper case variable with a lower case subscript to denote complex quantities, such as Laplace transforms, but uses the notation $I_d(s)$ to reinforce the fact that I_d is a function of the complex variable $s = \sigma + j\omega$.

Here, the convention for noise spectral density variables closely follows the notation used by Motchenbacher and Fitchen [75, page 2], using the IEEE standard for mean-square quantities. Thus, we use a capital letter squared for the current or voltage, with a three letter subscript, starting with n for “noise”, followed by either e for excess noise, f for flicker noise, s for shot noise, or t for thermal noise, and ending with either g for gate-referred, d for drain-referred, or b for values in the body. A roman type face is used in the subscript to take advantage of kerning and ligatures so it looks nice. For example,

I_{ned}^2 excess noise current at the drain in units of A^2/Hz

I_{nfd}^2 flicker noise current at the drain in units of A^2/Hz

I_{nsb}^2 shot noise current in the body in units of A^2/Hz

I_{nsd}^2 shot noise current at the drain in units of A^2/Hz

I_{ntd}^2 thermal noise current at the drain in units of A^2/Hz

V_{neg}^2 excess noise voltage referred to the gate in units of V^2/Hz

V_{nfg}^2 flicker noise voltage referred to the gate in units of V^2/Hz

V_{nsb}^2 shot noise voltage in the body in units of V^2/Hz

V_{ntg}^2 thermal noise voltage referred to the gate in units of V^2/Hz

¹Unfortunately, this mistake has probably miseducated a generation of electrical engineers.

It is noted that there is some redundancy in this notation when used with MOSFETs (the input at the gate is always a voltage, and the output at the drain is always a current), but some redundancy in writing is well-known to be helpful for the reader [76].

For device elements, physical parameters are expressed in all upper case (C_{ox}, W, L), while incremental model parameters are expressed in all lower case (g_m, r_o). However, some widely used notation that does not correspond to the above rules is nevertheless still used in this paper, because it is so widely used (like I_{ii}). Thus, in addition to the above noise variables, the following are used:

- α_H Hooge's empirical "constant"
- C_D depletion region capacitance
- C_I inversion layer capacitance
- C_{SS} surface states capacitance
- Δf noise bandwidth
- f_c Lorentzian corner frequency
- I_{ii} impact ionization current
- K_F empirical flicker noise coefficient
- μ_l mobility due to lattice scattering only
- N number of free carriers
- N_{ST} number of surface trap states
- r_{js} incremental resistance of the body/source junction

Other variables are explained as they are introduced.

B.2 Flicker Noise

Flicker noise dominates the noise spectrum at low frequency. Flicker noise was first observed in vacuum tubes over seventy-five years ago [77]. It gets its name from the anomalous “flicker” that was seen in the plate current. Flicker noise is also commonly called $1/f$ noise, because the noise spectrum varies as $1/f^\alpha$, where the exponent α is very close to unity ($\alpha = 1 \pm 0.2$).




Figure B.1: Plot of noise amplitude versus frequency (on linear scales) as measured in two dissimilar vacuum tubes by Johnson in 1925, from [77, Figure 7]. The vertical axis has been normalized by the expected shot noise amplitude.

Fluctuations with a $1/f$ power law have been observed in practically all electronic materials and devices, including homogenous semiconductors, junction devices, metal films, liquid metals, electrolytic solutions, and even superconducting Josephson junctions. In addition it has been observed in mechanical, biological, geological, and even musical systems. No entirely satisfactory physical explanation has been developed, and in fact, available evidence seems to suggest that the

origins of flicker noise in different devices may be quite different [78, page 143].

Two competing models have appeared in the literature to explain flicker noise: the McWhorter number fluctuation theory and the Hooge mobility fluctuation theory. There is experimental evidence to support both theories, and thus the literature is mostly split into two camps over the issue.

B.2.1 The McWhorter Model (Number Fluctuations)

McWhorter, working with germanium filaments at MIT Lincoln Laboratory in 1957 [79], proposed that flicker noise is primarily a surface effect. He cites a number of experiments that showed that the $1/f$ noise in germanium is dependent on the ambient atmosphere of the filament. He writes, “recent results now leave little doubt that the noise is predominantly, if not entirely, a surface phenomenon.”

Thus, the McWhorter number fluctuation (Δn) theory states that flicker noise is generated by fluctuations in the number of carriers due to charge trapping in surface states. McWhorter obtained the necessary $1/f$ spectrum by assuming that the time constant τ of the surface states varied with a $1/\tau$ distribution. Christensson et al [80, 81] were the first to apply the McWhorter theory to MOSFETs, using the assumption that the necessary time constants are caused by the tunneling of carriers from the channel into traps located inside the oxide.

A number of other applications of the McWhorter theory to MOSFETs have been done. Das and Moore [82] reviewed and compared a number of these theories and found that the basic assumptions behind any explanation affects the interpretation of experimental results. They said, “theoretical calculations of the MOSFET drain noise current have been performed by many investigators, which have led to different results, mainly due to the difference in the method of attack and nature of assumptions.”

Reimbold [83] developed the McWhorter model further, taking into account all the capacitive components of the small-signal equivalent circuit, so as to account for all transistor operating regimes. This development was done to fit Reimbold's measurements in weak inversion. Ghibaudo [84] showed a shortcut through Reimbold's work and came up with the same result, where the spectral density of the drain current is

$$I_{\text{nfd}}^2 = \frac{K_F q^4 I_D^2}{n^2 k T W L f} \frac{N_{ST}}{(C_{\text{ox}} + C_{SS} + C_D + C_I)^2}.$$

where N_{ST} is the surface trap density, C_{SS} is the interface state capacitance, C_D is the depletion capacitance, and C_I is the inversion layer capacitance. The bias dependence of the flicker noise is folded into the effective mobility, μ_{eff} , rewriting the above equation as

$$I_{\text{nfd}}^2 = \left(\frac{\mu_{\text{eff}}}{\mu_0} \right)^2 \frac{K_F q^4 I_D^2}{n^2 k T W L f} \frac{N_{ST}}{(C_{\text{ox}} + C_{SS} + C_D + C_I)^2}$$

where μ_0 is the low-field mobility and μ_{eff} has a gate-voltage dependence defined as

$$\mu_{\text{eff}} = \frac{\mu_0}{1 + \theta(V_{GS} - V_T)}$$

where θ is the mobility attenuation factor [42, page 146].

B.2.2 The Hooge Model (Mobility Fluctuations)

In his paper “ $1/f$ noise is no surface effect,” F. N. Hooge [85] proposed that $1/f$ noise is essentially a bulk phenomenon. Working with metal films [86], he championed an empirical relation for $1/f$ noise in terms of resistance fluctuations, where the spectral density of the resistance is

$$R_n^2 = \frac{\alpha_H R^2}{N f}$$

where N is the total number of free carriers in the bulk, and α_H is known as “Hooge’s constant” an empirical parameter with value about 2×10^{-3} . This equation fit his data for metal films very well (see his plots in Figure B.2). Based on these results [87], he said “Investigations of noise... proved that the fluctuations in the conductivity are due to fluctuations in mobility and not in the number of charge carriers.” Hooge summarized experimental support for his mobility fluctuation ($\Delta\mu$) theory in [88] and some theoretical support (a development of a phonon scattering theory) was provided by Jindal and van der Ziel [89].

Vandamme [90] developed Hooge’s theory for MOSFET devices, starting from the lattice scattering theory presented in [87], calculating

$$I_{\text{nfd}}^2 = \frac{\mu_0}{\mu_{\text{eff}}} \left(\frac{\mu_0}{\mu_l} \right)^2 \frac{q\alpha_H I_D^2}{WLC_{\text{ox}}(V_{GS} - V_T)} \frac{1}{f}$$

where μ_l is the mobility if only lattice scattering were present. Again, significant bias dependences can be explained by changes in the assorted mobility parameters.




Figure B.2: Flicker noise in ten metal films, from [86, Figure 1]. The solid lines correspond to a value of $\alpha_H = 2 \times 10^{-3}$ for Hooge's empirical parameter.

B.2.3 Combined Models

In an effort to fit as much published data as possible, some authors have combined the number fluctuation and mobility fluctuation theories into a single model. The main idea is that charges, when trapped, cause correlated surface mobility fluctuations (due to inversion layer carriers scattering off the trapped charges). Jayaraman and Sodini [91] were the first to propose this model². They also showed that oxide band bending due to nonuniform oxide trap distributions give rise to a gate voltage dependence of the flicker noise coefficient, and of the $1/f^\alpha$ exponent. However, from their data, they were unable to determine if the gate voltage dependence they observed was due to correlated mobility fluctuations, or to nonuniform oxide trap distributions (with a corresponding higher trap density), or both.

With an eye on improving the flicker noise model in circuit simulators such as SPICE, the unified model idea was developed further by Hung et al [92, 93]. By examining the channel current modulation due to the trapping of a single electron (popcorn noise in a very small device), they were able to quantify the contribution of the correlated surface mobility fluctuations. This approach led to

$$I_{\text{info}}^2 = \frac{kTI_D^2}{\gamma fWL} \left(\frac{1}{N} + \alpha'\mu \right)^2 N_T(E_{fn}).$$

where γ is the attenuation coefficient of the electron wave function into the oxide, α' is their scattering coefficient, and $N_T(E_{fn})$ is the number of traps at the quasi-Fermi level. However, their model is geared towards numerical simulation.

This recent literature supports number fluctuations as the primary source of flicker noise in MOSFETs. Responding to this evidence, Hooge [94] has recently suggested that “in semiconductors

²Although I have seen a similar theory attributed earlier to van der Ziel by Buckingham [78, page 137] but have been unable to locate the reference: A. van der Ziel. The oxide trap model of $1/f$ noise in MOSFETs. *Proc. Symp. on $1/f$ fluctuations, Orlando, Florida.* 1980.

there always is mobility $1/f$ noise with an α value of about 10^{-4} ” but that “other types of $1/f$ noise may be present and may dominate the mobility noise.”

Other authors, apparently in a peacemaking role, have suggested that $1/f$ noise in NMOS devices is caused by number fluctuations while in PMOS devices the noise is due to mobility fluctuations [95, 96]. This dichotomy makes sense if the NMOS transistors are surface channel devices (where the carriers interact more readily with Si-SiO₂ surface states), and the PMOS transistors are buried channel devices³ (where the carriers interact only with the lattice). However, differences in the energy trap density between NMOS and PMOS transistors could also account for the differences in observed bias dependence [97].

B.2.4 Simple Empirical Models

Many of the empirical models given by textbooks are much simpler than the theoretical and unified models above. Unfortunately, some of them are also questionable.

Tsividis [42, page 343] says that “in a common [bulk] CMOS fabrication process,” one may find an input-referred voltage due to flicker noise of

$$V_{\text{nfg}}^2 = \frac{K_F}{C_{\text{ox}}^2 WL} \frac{1}{f}$$

with K_F equal to $5 \times 10^{-9} \text{ fC}^2/\mu\text{m}^2$ for NMOS devices and $2 \times 10^{-10} \text{ fC}^2/\mu\text{m}^2$ for buried channel PMOS devices. This equation is an empirical relation, and K_F is not entirely independent of operating point (and is not the same in weak inversion and strong inversion). This model is implemented by early versions of SPICE, and is generally considered to be oversimplified.

³In modern deep-submicron technologies ($0.35 \mu\text{m}$ and below), both NMOS and PMOS transistors are surface channel devices, so this dichotomy may no longer exist.

B.3 Thermal Noise

Thermal noise is the voltage fluctuations caused by the random Brownian motion of electrons in a resistive medium. It is broadband white noise, and it gets worse with increasing resistance and temperature. The spectral density of the thermal noise across a resistor with resistance R is given by

$$V_{\text{nt}}^2 = 4kTR$$

A fifty ohm resistor has about 1 nV/ $\sqrt{\text{Hz}}$ of thermal noise (modeled as a noise voltage source in series with the resistor).

The thermal noise in the channel of a MOSFET in strong inversion is

$$I_{\text{ntg}}^2 = 4kT \left(\frac{2g_m}{3} \right)$$

where g_m is the small-signal transconductance at the bias point.

B.4 Shot Noise

Shot noise is caused by the fact that current flowing across a junction isn't smooth, but is comprised of individual electrons arriving at random times. This non-uniform flow gives rise to broadband white noise that gets worse with increasing average current. The spectral density of the shot noise associated with a junction current I is given by

$$I_{\text{ns}}^2 = 2qI.$$

A one milliamp current has about 2 pA/ $\sqrt{\text{Hz}}$ of shot noise (modeled as a noise current source in parallel with the junction).

MOSFETs in subthreshold exhibit shot noise (instead of thermal noise) due to the current flowing in the channel. This noise source has the standard form

$$I_{\text{nsd}}^2 = 2qI_D.$$

since we are in subthreshold, the effective noise voltage at the gate can be written as

$$V_{\text{nsd}}^2 = \frac{I_{\text{nsd}}^2}{g_m^2} = \frac{2q}{I_D} \left(\frac{nkT}{q} \right)^2.$$

For low drain-to-source voltage ($V_{DS} < 5kT/q$), the shot noise grows in value [98], approaching

$$I_{\text{nsd}}^2 = 4qI_{\text{sat}}$$

where I_{sat} is the saturation current of the transistor in subthreshold, defined where the drain current is

$$I_D = I_{\text{sat}} \left[1 - \exp \left(-\frac{qV_{DS}}{kT} \right) \right].$$

For complete equations for I_{sat} see [42, page 139].

B.5 Generation/Recombination Noise

In addition to the flicker noise caused by traps in the oxide, trapping centers in the bulk of the device can cause generation/recombination (G/R) noise. The trapping of carriers by these traps causes fluctuations in the number of carriers, and thus fluctuation in the resistance. The spectral

density of the resistance fluctuation is [94]

$$R_n^2 = \frac{\sigma^2}{N^2} \frac{4\tau R^2}{1 + \omega^2 \tau^2}$$

where τ is the trap relaxation time, and the variance σ^2 is given by

$$\frac{1}{\sigma^2} = \frac{1}{N} + \frac{1}{N_O} + \frac{1}{N_E}$$

where N_O is the number of occupied traps and N_E is the number of empty traps. If there is more than one kind of trap, the equations are significantly more complicated. Note that G/R noise creates a Lorentzian noise spectrum.

B.6 Popcorn Noise

Popcorn noise, sometimes called burst noise or random-telegraph-signal (RTS) noise, is a discrete modulation of the channel current caused by the capture and emission of a channel carrier. See Figure B.3.




Figure B.3: Typical popcorn noise, showing discrete levels of channel current modulation due to the trapping and release of a single carrier, for three different bias conditions, from [99, Figure 1].

B.7 kT/C Noise

kT/C noise⁴ really isn't a fundamental noise source, but is actually thermal noise in the presence of a filtering capacitor.

Consider a passive low-pass filter composed of a resistor and a capacitor, with the output voltage measured across the capacitor. When the input is shorted to ground, the broadband thermal noise of the resistor will be shaped by the low-pass filter. The spectral density of the thermal noise of the resistor is

$$V_{\text{nt}}^2 = 4kTR.$$

In order to find the total output noise, we must integrate the above spectral density over the noise bandwidth, Δf . Expressed in Hertz, this noise bandwidth is

$$\Delta f = \frac{1}{2\pi} \int_0^\infty \frac{d\omega}{1 + (\omega RC)^2} = \frac{1}{2\pi} \frac{\pi}{2RC} = \frac{1}{4RC}$$

Thus, the total output noise voltage that is measurable across the capacitor, in volts, is simply

$$v_{no} = \sqrt{V_{\text{nt}}^2 \Delta f} = \sqrt{\frac{kT}{C}}.$$

This result is not a new noise source, but a repackaging of thermal noise. It is simply the result of low-pass filtering the noise from a resistor. It doesn't matter if the filter above is a discrete resistor and capacitor, or the channel resistance and drain capacitance of a MOSFET. Note that to include thermal noise and kT/C noise in a single noise calculation would be double-counting.

⁴CCD designers [100], who count electrons instead of volts, talk instead about kTC noise ($Q_n^2 = kTC$, usually expressed in units of electrons-squared per Hertz).

Appendix C

Reported Noise Sources in SOI CMOS

A source of additional noise exists in SOI transistors. This excess noise was first reported by Chen et al [101] (using the unfortunate terminology “noise overshoot”). It consists of a single Lorentzian hump: a flat plateau followed by a $1/f^2$ slope that appears above the background of flicker noise and thermal noise, sometimes significantly so. See Figure C.1 for a sample plot of the noise spectral density.

Although the physical source of this noise is not (yet) completely understood (or universally agreed upon), there is no shortage of attempted explanations. Chen et al attributed the phenomenon to traps at the buried oxide interface. Due to the Lorentzian-like noise spectrum, the noise was also considered as the result of charge fluctuations caused by generation/recombination events in the depletion region of the transistor, as has been seen in silicon-on-sapphire technology [102].

This appendix presents three recently proposed theories of the excess noise source in SOI transistors. The first paper, “Noise contribution of the body resistance in partially depleted SOI MOSFET’s,” written by Faccio et al [50] is discussed in Section C.1. The second paper, “Shot-noise-induced excess low-frequency noise in floating-body partially depleted SOI MOSFET’s,” written




Figure C.1: Noise spectral density for an SOI transistor, showing the Lorentzian hump when the drain voltage is above the kink onset voltage, from [101, Figure 2].

by Wei et al [51] is discussed in Section C.2. The third paper, “Floating body induced pre-kink excess low-frequency noise in submicron SOI CMOSFET technology,” written by Tseng et al [52] is discussed in Section C.3.

C.1 Noise Contribution of the Body Resistance

Faccio et al [50, 103] measured excess noise in body contacted H-gate transistors (see Figure C.2), and identified its origin to be the thermal noise of the body film. Although the body resistance was small enough to avoid the kink effect, it was still large enough to produce excess noise. These measurements were made in saturation, but with a drain-to-source voltage below the kink onset voltage ($V_{DS} = 0.8V$).




Figure C.2: SOI noise spectrum, referred to the gate, in H-gate transistors observed by Faccio et al, showing $1/f$ noise, white (thermal) noise and excess noise components, from [50, Figure 2].

Noticing the Lorentzian shape of the excess noise curve, Faccio et al considered generation/recombination noise as a possible source. However, they note that theory states G/R noise should be strongly sensitive to drain bias, and they did not observe such a dependence. Also, they observed the wrong dependence on gate area.

Their model for the noise is based on the white thermal noise of the body resistance, R_B , being converted to noise current at the drain by the body transconductance, g_{mb}

$$I_{ned}^2 = 4kTR_B g_{mb}^2$$

From the body voltage, the gate oxide and the buried oxide are seen in parallel to ground (see the schematic of their model in Figure C.3). This combination creates a low pass filter with the body resistance that gives the Lorentzian shape to the excess noise curve, which, referred to the gate, is

$$V_{\text{neg}}^2 = \frac{4kTR_B}{1 + (f/f_c)^2} \left(\frac{g_{mb}}{g_m} \right)^2$$

with corner frequency

$$f_c = \frac{1}{2\pi R_B (C_{\text{ox}} + C_{\text{box}})}$$




Figure C.3: Schematic of the noise model by Faccio et al, showing the noisy body resistance and the filtering capacitors due to C_{ox} and C_{box} , from [50, Figure 6].

In order to support this theory Faccio completed a number of experiments. It was observed that the noise level and cutoff frequency were strongly dependent on the backgate voltage. To test if the body resistance demonstrated a similar dependence, a special transistor was constructed with independent contacts to the body to allow rough measurements of the sheet resistance (see the layout in Figure C.4). For changes in backgate and body bias, the measured resistance changed over three orders of magnitude. The change in magnitude of the measured body resistance for different biases compared well to the change in magnitude of the measured body-referred noise voltage.

Different bias conditions (backgate biased with body grounded versus body biased with backgate grounded) that resulted in similar measured body resistance produced similar body-referred noise measurements to within twenty percent.




Figure C.4: Layout of special transistor, showing independent contacts to the body allowing resistance measurements, from [50, Figure 7].

Also, signals were injected into the body terminal to measure the frequency response of the transistor. The low-pass filter created by the body resistance and the oxide capacitances was effective for all signals injected through the body terminal.

C.2 Shot Noise Induced Excess Low Frequency Noise

Wei et al [51] identify the origin of the excess noise to be the shot noise associated with impact ionization current and body-to-source diode current. This shot noise, which would otherwise be negligible, is amplified by the floating body effect.

There are two currents in the body that give rise to this noise source. One results from source-to-body diode leakage, and the other is due to impact ionization (same as in the kink effect). Note that these currents are roughly equal to each other due to current balancing in the body.

$$I_{\text{nsb}}^2 = 2qI_{ii} + 2qI_{SB} \approx 4qI_{ii}$$

These currents flow through the impedance from the body to ground (see the circuit model in Figure C.5) and change the body potential

$$V_{\text{nsb}}^2 = \frac{I_{\text{nsb}}^2 r_{js}^2}{1 + (2\pi f r_{js} c_{eq})^2}$$

where c_{eq} is the sum of the body-to-channel and body-to-backgate depletion capacitances and the body-to-source and body-to-drain junction capacitances, and r_{js} is the incremental resistance of the body-to-source diode,

$$r_{js} = \frac{nkT}{qI_{ii}}.$$

This fluctuation in body potential changes the threshold voltage of the device, which causes the excess noise in the drain current.

$$I_{\text{ned}}^2 = V_{\text{nsb}}^2 \left(\frac{di_D}{dv_{BS}} \right) = \frac{4qI_{ii} r_{js}^2 g_{mb}^2}{1 + (2\pi f r_{js} c_{eq})^2}$$

This explanation is supported by the coincidence of the corner frequency in the Lorentzian noise spectrum and the AC output impedance at

$$f_c = \frac{1}{2\pi r_{js} c_{eq}}.$$

Wei et al note that this model is valid for strong inversion only, and it does not apply to body-contacted transistors (since an effective body contact would prevent the floating body effect).

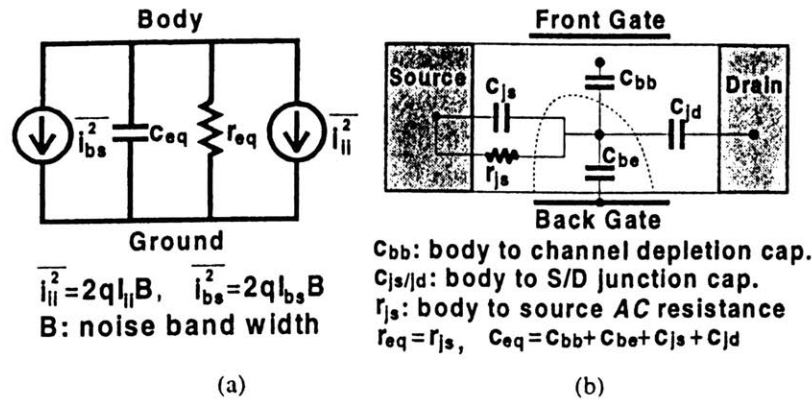


Figure C.5: Schematic of noise model by Wei et al, showing low pass filter created by the incremental source-to-body junction resistance and the capacitances to the channel, from [51, Figure 2].

They report very good matching between their proposed noise model using extracted device parameters and experimental results. When the drain bias is above the kink onset voltage, the Lorentzian noise spectrum is clearly visible in their data.

In contrast to the observations of Faccio et al [50], the noise spectrum here was dependent on transistor bias. A larger drain bias gave a higher corner frequency but a smaller noise magnitude since the value of r_{js} decreases with increasing V_{DS} (which corresponds to increasing I_{ii}).

The excess noise exhibits similar characteristics as the output impedance. Their measurements showed that the corner frequency in the Lorentzian coincided with the corner frequencies on the AC

output impedance. The frequency dispersion in the AC output impedance results from the floating body effect, and is also determined by r_{js} and c_{eq} . Thus, they conclude that the coincidence of the corner frequencies in the noise spectrum and the output impedance indicates that the floating body effect is the physical mechanism which amplifies that shot noise in the body and leads to the excess low-frequency noise in the drain current.

In their article, they assert that the excess noise is not observable in either the triode regime or the pre-kink saturation regime, but this claim is not supported by the referenced figure. In Figure C.6 some excess noise is clearly visible in the pre-kink regime (perhaps it was ignored since their theory does not account for it). This excess noise could correspond to a Lorentzian hump at very low frequency.




Figure C.6: Spectral density of excess noise, showing increased noise for drain voltages above and below the kink onset voltage, from [51, Figure 7].

Wei et al also report that this source does not exist in fully depleted (FD) transistors, but does occur if the back surface of the film is accumulated (which would correspond to partially depleted operation). They demonstrate this effect using FD transistors and a backgate bias of $V_{BS} = -20V$.

See their noise plots in Figure C.7.




Figure C.7: Spectral density of excess noise for a fully depleted device in full depletion and backside accumulation. The excess noise disappears when the device is fully depleted, from [51, Figure 10].

C.3 Pre-Kink Excess Low Frequency Noise

Ying-Che Tseng has published a number of papers using devices from the Motorola RF-TFSOI Group. His paper [52] builds on and corrects his previous publications [104, 105]. In his previous publications he explained the cause of the excess noise as either generation/recombination noise [104] or as tunneling noise between the gate oxide and the channel [105].

In [104] (written at Arizona State), the authors observe increased noise in the subthreshold regime for an NMOS transistor in the form of two Lorentzian shaped humps in the noise spectral density. They assert that this excess noise is due to generation/recombination centers. Their data for a PMOS transistor is difficult to interpret, showing increased noise at high frequency without any recognizable structure.

In [105] the authors (a different set of coauthors, now at UCLA) claim that the source of the Lorentzian noise is “electronic tunneling noise between the gate oxide and the channel.” They claim that the AC kink enhances tunneling noise with the tunneling time constant equal to $\tau_T = 2\pi/\omega_0$ by a mysterious gain factor of $A(\omega_0)$. They reference the plot of the derivative of the output conductance, shown here in Figure C.8, however, they give no indication of derivation or cause of this gain factor. There seems to be no evidence supporting this explanation.

This most recent paper [52] offers yet another explanation which is closely modeled after the referenced work of Wei et al [106]. However, they report excess noise observed below the kink (such as $V_{DS} = 0.5V$) that cannot be explained by impact ionization current alone. They explain that even without impact ionization current, there is still current flowing into the body from other sources.

They propose that in the pre-kink regime, there is a finite source-to-body junction current due to the drain-to-body leakage current $I_{SB} = I_{DL}$. In the post-kink regime, the source-to-body




Figure C.8: Output conductance plot, showing AC floating body effect. Inset plot shows derivative of AC conductance. Note how the derivative is peaky at the onset of the AC kink, from [105, Figure 2].

junction current is due to impact ionization $I_{SB} = I_{ii}$. In both cases the noise is

$$I_{\text{ned}}^2 = \frac{4qI_{SB}(nV_T)^2 g_{mb}^2}{[1 + (f/f_c)^2](I_R + I_{SB})^2}$$

Note that their corner frequency is

$$f_c = \frac{1}{2\pi r_{js} C_{BB}}$$

where C_{BB} is the total body to ground capacitance and r_{js} is the source/body diode junction resistance

$$r_{js} = \frac{nkT}{qI_{SB}}.$$

Again, it is reported that shifting towards fully depleted operation reduces the excess noise. A good body contact provides another low-resistance discharge path for charge in the body, where the impedance of current fluctuation is reduced to $R_B || r_{js}$.

C.4 Generation/Recombination Noise

Other authors, in addition to Tseng et al, have also reported excess noise due to the kink effect in subthreshold SOI transistors. Valenza [107] observed a Lorentzian spectrum in subthreshold, but explained its source to be generation/recombination noise.

This explanation is clearly the work of one of the coauthors, E. Simoen. The unabashedly prolific Simoen has authored a number of similar papers on the excess noise of SOI transistors [108, 109, 110, 111, 112] maintaining that the source of the excess noise was the trapping and re-emission of the channel carriers.

Simoen developed his generation/recombination theory of the excess noise in [109]. He hastily concluded that since the excess noise has a Lorentzian spectrum, it must have a generation/recombination origin. He explains that the noise originates from number fluctuations due to deep-level traps in the depletion region of the silicon film trapping the carriers produced by impact ionization. This assumption leads to the following expression for the excess noise

$$I_{\text{ned}}^2 = \frac{g_m^2 q^2 N_T w_d}{C_{\text{ox}}^2 W L} \frac{4\tau_2^2 / \tau_1}{1 + (2\pi\tau_2 f)^2}$$

where w_d is the width of the depletion region where the traps contribute, τ_1 is the sum of the capture time constant τ_c and the emission time constant τ_e ($\tau_1 = \tau_c + \tau_e$), and τ_2 is their parallel

$$\tau_2 = \frac{\tau_c \tau_e}{\tau_c + \tau_e}.$$

The traps will only contribute to the noise if they lie in that part of the depletion region, with width w_d , where the ratio of the time constants, τ_c/τ_e is approximately between 0.1 and 10. He then goes to great lengths to fold the bias dependencies into his definition of the capture time constant. In

his 1996 review paper [112] he maintains his faith in this model.

This theory seems to be built on a house of cards, not the least of which is the number of unknown parameters that are left to be adjusted to fit the experimental data. Also ignored by Simoen is the often complex structure of the frequency dependence of generation/recombination noise, such as that published for silicon-on-sapphire. In the noise spectrum of SOS, several Lorentzians with a variety of time constants are visible (see the example plot in Figure C.9). These humps are due to traps associated with the poor quality of the back interface. In contrast, SOI only produces a single hump in the excess noise curve.




Figure C.9: Spectral density of SOS noise, clearly showing multiple Lorentzians with multiple corner frequencies due to generation/recombination noise, from [102, Figure 7].

C.5 Comparison of Theories

Overall, the excess noise theory proposed by Wei et al is the most convincing. This theory is the most straightforward, and seems best supported by the experimental data. The minor modification, proposed by Tseng et al, to explain their observations at drain voltages below the kink is a logical extension.

Faccio et al may have seen an additional noise source, and their proposed theory explains their observations very well. In particular, their observed dependence of the noise on body bias is well explained by their theory and subsequent experiments. However, they also made measurements on a source side body contacted transistor that they could not explain with their theory. From their discussion:

The noise characteristics of a [source side body contact] transistor were anomalous as well. The additional noise component was obvious, causing a hump in the noise spectrum up to the MHz frequency range, but it could not be easily controlled by the backgate potential.

While the problem of substrate resistance noise is known in bulk CMOS [113], no other publications have reported noise attributable to the SOI body resistance.

Appendix D

Simulation Code

For full circuit simulations, the repeating and regular structure of the flash converter was exploited by heavy use of SPICE subcircuits. The SPICE code is broken up into many subcircuits stored in separate include files. The subfiles are joined together by the main system file.

The largest of these SPICE files, the main converter topology, was generated by a Perl script. This was done to exploit the regular structure and to eliminate typing errors.

D.1 SPICE File `asystem.sp`

The top level file `asystem.sp` includes all of the individual subcircuits, and runs the whole simulation.

```
* Analog-to-Digital Converter System (6 bits)

.prot
.include 'xmodel.sp'
.include 'xclocks.sp'
.include 'xbubble.sp'
.include 'xcompare.sp'
.include 'xlatch.sp'
.include 'adc.sp'
.unprot

vdd vdd 0 1

* the ideal 6-bit IDAC
.subckt dac bit0 bit1 bit2 bit3 bit4 bit5 iout
  g0 0 iout pwl(1) bit0 0 0.49,0 0.51,0.015625m
  g1 0 iout pwl(1) bit1 0 0.49,0 0.51,0.03125m
  g2 0 iout pwl(1) bit2 0 0.49,0 0.51,0.0625m
  g3 0 iout pwl(1) bit3 0 0.49,0 0.51,0.125m
  g4 0 iout pwl(1) bit4 0 0.49,0 0.51,0.25m
  g5 0 iout pwl(1) bit5 0 0.49,0 0.51,0.5m
.ends

* sine input -> my adc -> ideal idac -> termination
vsinp vinp 0 sin 0.5 0.5 20e6
vsinm vinm 0 sin 0.5 -0.5 20e6
xadc vinp vinm b0 b1 b2 b3 b4 b5 adc
xdac b0 b1 b2 b3 b4 b5 vout dac
rdac vout 0 1000

.option post nomod
.tran 0.2ns 500ns
.print tran v(vout)
.measure tran current avg i(vdd) from=0ns to=500ns
.end
```

D.2 SPICE File xcompare.sp

The comparator is the heart of the flash design. See Chapter 5.




Figure D.1: Single-ended auto-zeroed comparator.

* *comparator subcircuit*

```
.subckt cswit 1 2 control cbar
  ms1 1 cbar 2 vdd PMOS L=0.25u W=8u
  ms2 1 control 2 0 NMOS L=0.25u W=8u
.ends
```

```
.subckt compare inp inm out
  xsw1 inp sum ph12 ph12b cswit
  xsw2 inm sum ph12b ph12 cswit
  xsw3 ain out ph12 ph12b cswit
  ca sum ain 50f
  m1 out ain vdd vdd PMOS L=0.5u W=4u
  m2 out ain 0 0 NMOS L=0.5u W=2u
.ends
```

10

D.3 SPICE File xlatch.sp

The latch holds the previous output value during the next conversion cycle. It consists of a precharging positive-feedback latch, followed by two half-latch circuits.

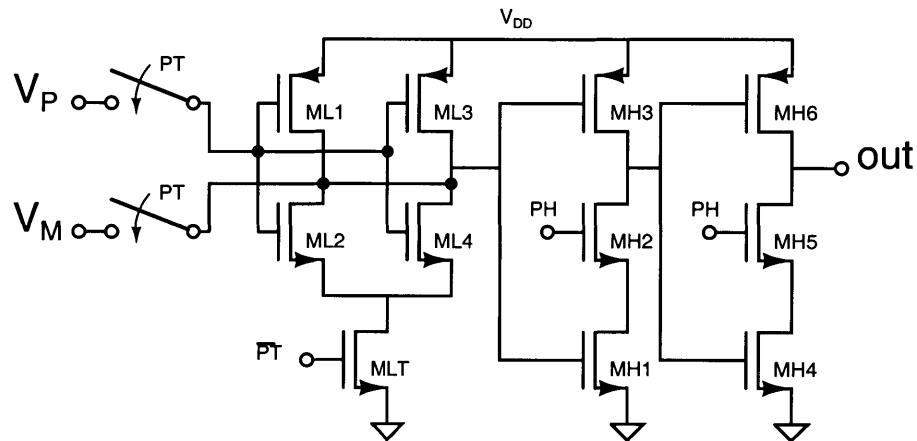


Figure D.2: Regenerative latch and doubled half-latch.

** Latch*

```
.subckt lswit 1 2 control cbar
  ms1 1 cbar 2 vdd PMOS L=0.25u W=8u
  ms2 1 control 2 0 NMOS L=0.25u W=8u
.ends
```

```
.subckt latch inp inm out
```

** latch input section*

```
xlsw1 inp outp ph23 ph23b lswit
xlsw2 inm outm ph23 ph23b lswit
ml1 outp outm vdd PMOS L=0.5u W=2.8u
ml2 outp outm tail NMOS L=0.5u W=1.8u
ml3 outm outp vdd PMOS L=0.5u W=2.8u
ml4 outm outp tail NMOS L=0.5u W=1.8u
mlt tail ph23b 0 NMOS L=0.25u W=3.8u
```

** half latch output section*

```
mh1 1 outp 0 NMOS L=0.25u W=1.8u
mh2 bin ph1 1 NMOS L=0.25u W=1.8u
mh3 bin outp vdd PMOS L=0.25u W=1.8u
mh4 2 bin 0 NMOS L=0.25u W=1.8u
mh5 out ph1 2 NMOS L=0.25u W=1.8u
mh6 out bin vdd PMOS L=0.25u W=1.8u
```

```
.ends
```

10

20

D.4 SPICE File xbubble.sp

The bubble circuit implements the majority function. This implementation saves a NAND gate by using the AND-IN (a_i) and AND-OUT (a_o) signals.

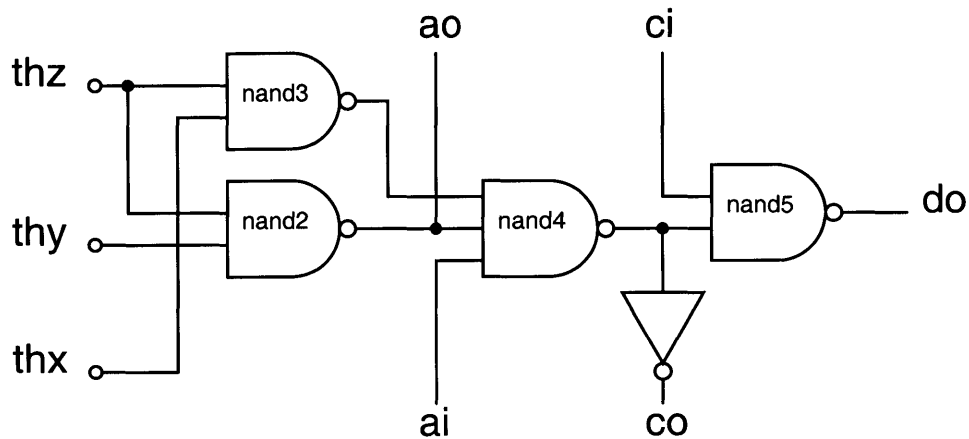


Figure D.3: Bubble-error correction majority circuit.

** thermometer code bubble correction logic*

.subckt inv in out

m1 out in vdd PMOS L=0.25u W=1.8u
m2 out in 0 NMOS L=0.25u W=0.8u

.ends

.subckt nand2 a b out

m1 out a vdd PMOS L=0.25u W=1.8u
m2 out b vdd PMOS L=0.25u W=1.8u
m3 out a mid NMOS L=0.25u W=1.8u
m4 mid b 0 NMOS L=0.25u W=1.8u

.ends

.subckt nand3 a b c out

m1 out a vdd PMOS L=0.25u W=1.8u
m2 out b vdd PMOS L=0.25u W=1.8u
m3 out c vdd PMOS L=0.25u W=1.8u
m4 out a mid NMOS L=0.25u W=2.8u
m5 mid b mid NMOS L=0.25u W=2.8u
m6 mid c 0 NMOS L=0.25u W=2.8u

.ends

.subckt bubble thx thy thz ai ao ci co do

xnand2 thx thz v nand2
xnand3 thy thz ao nand2

10

20

```
xnand4 ai v ao t nand3  
xnand5 ci t do nand2  
xinv t co inv  
.ends
```

30

D.5 SPICE File `xclocks.sp`

There are three main clock signals that drive the A/D converter, as shown in Figure D.4.




Figure D.4: The three primary clock phases and their purposes.

The subcircuit file `xclocks.sp` produces these necessary clock signals from a pair of differential quadrature sine wave inputs. It also produces the complement of phases PH12 and PH23. Proper functionality of the circuit depends on the timing of these input waveforms.

The clock input circuit, shown in Figure D.5, buffers the input sine waves and produces proper square waves to drive the clock generation logic. The clock generation logic, shown in Figure D.6, produces the five necessary clock signals. The final inverters in the clock generation are distributed throughout the layout as sixty-three small inverter circuits.

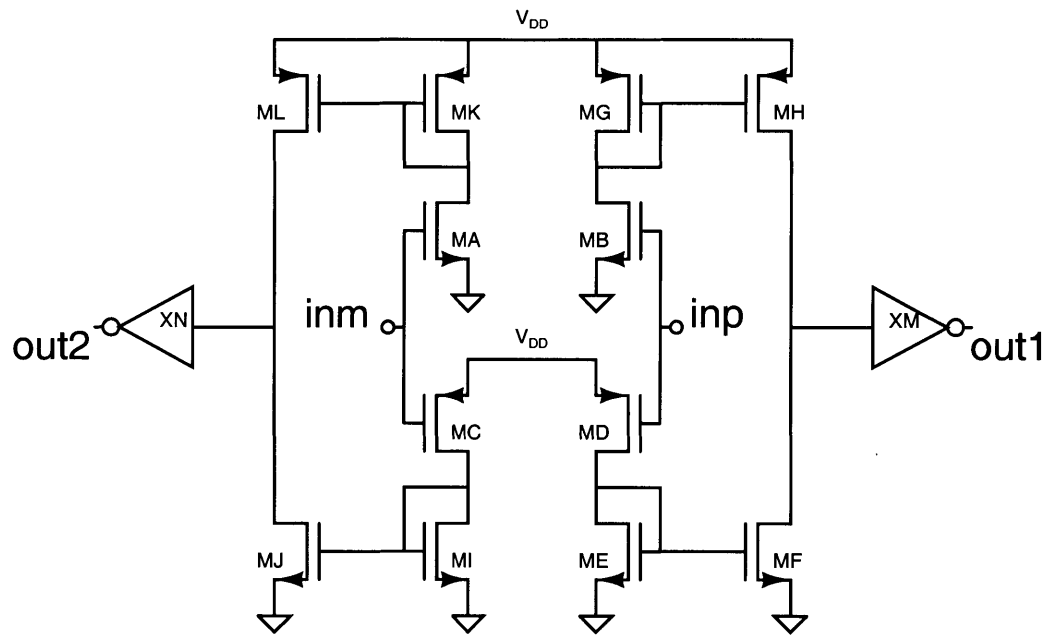


Figure D.5: Differential clock input buffer.

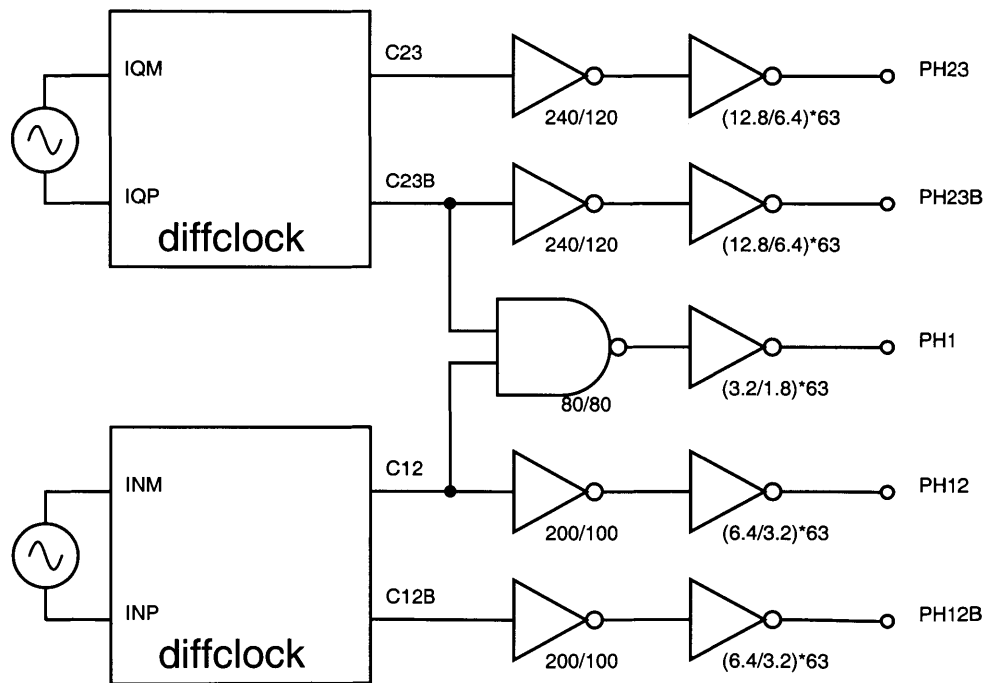


Figure D.6: Clock phase generator and distribution buffers.

** global clocks*

.global vdd ph1 ph12 ph12b ph23 ph23b
.param speed = 5e8

vinp inp 0 **sin** 0.5 0.25 'speed'
vinm inm 0 **sin** 0.5 -0.25 'speed'
viqp iqp 0 **sin** 0.5 0.25 'speed' 0 0 -90
viqm iqm 0 **sin** 0.5 -0.25 'speed' 0 0 -90

10

.subckt dcinv in out
 m1 out in vdd vdd PMOS L=0.25u W=120u
 m2 out in 0 0 NMOS L=0.25u W=60u
.ends

.subckt diffclock inp inm out1 out2
 ma 8 inm 0 0 NMOS L=0.25u W=80u
 mb 6 inp 0 0 NMOS L=0.25u W=80u
 mc 9 inm vdd vdd PMOS L=0.25u W=80u
 md 7 inp vdd vdd PMOS L=0.25u W=80u
 me 7 7 0 0 NMOS L=0.25u W=80u
 mf ou1 7 0 0 NMOS L=0.25u W=80u
 mg 6 6 vdd vdd PMOS L=0.25u W=80u
 mh ou1 6 vdd vdd PMOS L=0.25u W=80u
 mi 9 9 0 0 NMOS L=0.25u W=80u
 mj ou2 9 0 0 NMOS L=0.25u W=80u
 mk 8 8 vdd vdd PMOS L=0.25u W=80u
 ml ou2 8 vdd vdd PMOS L=0.25u W=80u
 xm ou1 out1 dcinv
 xn ou2 out2 dcinv
.ends

20

30

.subckt ckib in out
 m1 out in vdd vdd PMOS L=0.25u W=240u
 m2 out in 0 0 NMOS L=0.25u W=120u
.ends

.subckt ckic in out
 m1 out in vdd vdd PMOS L=0.25u W=200u
 m2 out in 0 0 NMOS L=0.25u W=100u
.ends

40

.subckt ckid in out
 m1 out in vdd vdd PMOS L=0.25u W=12.8u
 m2 out in 0 0 NMOS L=0.25u W=6.4u
.ends

.subckt ckie in out
 m1 out in vdd vdd PMOS L=0.25u W=3.2u
 m2 out in 0 0 NMOS L=0.25u W=1.8u
.ends

50

.subckt ckif in out
 m1 out in vdd vdd PMOS L=0.25u W=6.4u
 m2 out in 0 0 NMOS L=0.25u W=3.2u
.ends

```

.subckt clknand a b out
  m1 out a vdd vdd PMOS L=0.25u W=80u
  m2 out b vdd vdd PMOS L=0.25u W=80u
  m3 out a mid 0 NMOS L=0.25u W=80u
  m4 mid b 0 0 NMOS L=0.25u W=80u
.ends

xdca inp inm c12 c12b diffclock
xdcb iqp iqm c23 c23b diffclock

xc1 c12 ca ckib
xc2 c12b cb ckib
xc3 c12 c23b cc clknand
xc4 c23 cd ckic
xc5 c23b ce ckic

xc6 ca ph12 ckid m=63
xc7 cb ph12b ckid m=63
xc8 cc ph1 ckie m=63
xc9 cd ph23 ckif m=63
xc0 ce ph23b ckif m=63

ccl1 ph12 0 3.3pF
ccl2 ph12b 0 3.3pF
ccl3 ph23 0 1.1pF
ccl4 ph23b 0 1.4pF
ccl5 ph1 0 0.4pF

```

60

70

80

D.6 Program adc.pl

The perl program `adc.pl` writes the flash ADC, properly interconnecting all the occurrences of the comparators, latches, and bubble logic blocks. It completes the design by writing out the final decoder ROM at the transistor level. For a sample four-bit output file, see Section D.7.

```
# adc.pl
$version = 2.2;
#
#
# This handy perl script creates a flash adc spice input file
# set the number of bits in $nob below
# usage: perl adc.pl > adc.sp
#
# VERSION upgrade history:
#   version 1.0 first
#   version 1.1 subcircuits moved to .include files
#   version 1.2 added ROM diagram
#   version 1.3 pared down bubble circuit by one gate
#       bubble now requires carry_in and nand_in
#   version 1.4 replaced $size/2 input OR gates with
#       proper tree of 2-input NOR/NAND gates
#   version 1.5 R_ref string center-tapped for LSB fix
#   version 1.6 implemented encoding ROM for output
#   version 2.0 comparator now fully differential
#   version 2.1 added some pretty-printing and indenting
#
# compare -> latch -> bubble -> rom
#
# Kent Lundberg, March/April 2000

$nob = 4;
$rref = "1";
$size = 2**$nob;

print STDERR "Thank you for coming to Loews. ";
print STDERR "Sit back and relax. Enjoy the show!\n";
print "* ",$nob,"-bit FLASH adc, ",$size-1," comparators.\n";
print "* produced by adc.pl version $version\n";

# main adc subcircuit, loop adds appropriate number of digital outputs
print "\n* main adc subcircuit, with ",$nob," digital outputs\n\n";
print ".subckt adc vinp vinm";
for ($loop = 0; $loop < $nob; $loop++) {
    print " bit",$loop;
}
print "\n* vinp is the plus input, vinm is the minus input\n\n";

# comparator array and reference resistor string
# first (last) resistor requires connection to ground (vref)
print "* reference resistor string and comparator array\n";
print " vrep0 rep0 0 0\n";
```

```

print " vrem0 rem0 0 1\n";
for ($loop = 1; $loop < $size; $loop++) {
  print "* comparator number $loop\n";
  print " rrepa",$loop," rep",$loop-1," midp",$loop," ",$rref,"\n";
  print " rrepb",$loop," midp",$loop," rep",$loop," ",$rref,"\n";
  print " rrema",$loop," rem",$loop-1," midm",$loop," ",$rref,"\n";
  print " rremb",$loop," midm",$loop," rem",$loop," ",$rref,"\n";
  print " xcomp",$loop," vinp midp",$loop," lp",$loop," compare\n";
  print " xcompm",$loop," vinm midm",$loop," lm",$loop," compare\n";
}
print " vrep1 rep",$loop-1," 0 1\n";
print " vrem1 rem",$loop-1," 0 0\n\n";

# dump out the latches, too
print "* latches\n";
for ($loop = 1; $loop < $size; $loop++) {
  print " xlat",$loop," lp",$loop," lm",$loop," th",$loop," latch\n";
}
print "\n";

# bubble correction logic array
# bottom gate thx connected to th0 = H
# bottom and-in connected to a0 = NOT(th1)
# bubble thx thy thz ai ao ci co do
# top gate thz connected to thzz = L
# top carry-in connected to czz = H
print "* bubble correction logic array\n";
print " vh0 th0 0 1\n";
print " xith1 th1 a0 inv\n";
for ($loop = 1; $loop < $size; $loop++) {
  print " xsub",$loop," th",$loop-1," th",$loop," th",$loop+1," a",$loop-1;
  print " a",$loop," c",$loop+1," c",$loop," d",$loop," bubble\n";
}
print " rthz th",$loop," 0 10\n";
print " vhz c",$loop," 0 1\n\n";

# diagrams are almost as good as comments
$sizh = $size/4;
print "\n* short connection diagram\n";
for ($loop = 0; $loop < $nob; $loop++) {
  $cell = "0"x(2**$loop)."1"x(2**$loop);
  $bin[$loop] = $cell x 2**($nob-$loop-1);
  $out[$loop] = pack("b$sizh",$bin[$loop]);
  print "* bit$loop = ",unpack("h$sizh",$out[$loop])," ",$bin[$loop],"\n";
}
print "\n";

# final output binary encode using a ROM
print "* final output binary encode using a ROM\n";
for ($loopx = 0; $loopx < $nob; $loopx++) {
  print "* ROM transistors for output bit$loopx\n";
  for ($loopy = 1; $loopy < $size; $loopy++) {
    if ($loopy&2**$loopx) {
      print " mrom",$loopx*100+$loopy," rom",$loopx;
      print " d",$loopy," vdd vdd PMOS L=0.25u W=15u\n";
    }
  }
  print " mromn$loopx rom$loopx vdd 0 0 NMOS L=0.25u W=3.8u\n";
  print " cload$loopx bitb$loopx 0 0.05pF\n";
}

```



```
# buffer the output with some inverters
print " mbitx$loopx bitb$loopx rom$loopx 0 0 NMOS L=0.25u W=0.8u\n";
print " mbitb$loopx bitb$loopx rom$loopx vdd vdd PMOS L=0.25u W=1.8u\n";
print " mbitc$loopx bit$loopx bitb$loopx 0 0 NMOS L=0.25u W=0.8u\n";
print " mbitd$loopx bit$loopx bitb$loopx vdd vdd PMOS L=0.25u W=1.8u\n";
}

#
# This line has to end the output file, dude.
#

print ".ends\n\n* Th-th-th-th-th-that's all, Folks!\n";
```

110

120

D.7 Sample adc.pl Output File

Here is a sample output file from program `adc.pl`, for a four-bit flash converter. (The SPICE listing for the six-bit converter is not reproduced here because it is, of course, four times longer than the six-bit version.)

```
* 4-bit FLASH adc, 15 comparators.
* produced by adc.pl version 2.2
*
* main adc subcircuit, with 4 digital outputs

.subckt adc vinp vinm bit0 bit1 bit2 bit3
* vinp is the plus input, vinm is the minus input

* reference resistor string and comparator array
vrep0 rep0 0 0
vrem0 rem0 0 1
* comparator number 1
rrep1 rep0 midp1 1
rrepb1 midp1 rep1 1
rema1 rem0 midm1 1
remb1 midm1 rem1 1
xcompp1 vinp midp1 lp1 compare
xcompm1 vinm midm1 lm1 compare
* comparator number 2
rrep2 rep1 midp2 1
rrepb2 midp2 rep2 1
rema2 rem1 midm2 1
remb2 midm2 rem2 1
xcompp2 vinp midp2 lp2 compare
xcompm2 vinm midm2 lm2 compare
* comparator number 3
rrep3 rep2 midp3 1
rrepb3 midp3 rep3 1
rema3 rem2 midm3 1
remb3 midm3 rem3 1
xcompp3 vinp midp3 lp3 compare
xcompm3 vinm midm3 lm3 compare
* comparator number 4
rrep4 rep3 midp4 1
rrepb4 midp4 rep4 1
rema4 rem3 midm4 1
remb4 midm4 rem4 1
xcompp4 vinp midp4 lp4 compare
xcompm4 vinm midm4 lm4 compare
* comparator number 5
rrep5 rep4 midp5 1
rrepb5 midp5 rep5 1
rema5 rem4 midm5 1
remb5 midm5 rem5 1
xcompp5 vinp midp5 lp5 compare
```

xcompm5 vinm midm5 lm5 compare
 * *comparator number 6*
 rrepa6 rep5 midp6 1
 rrepb6 midp6 rep6 1
 rrema6 rem5 midm6 1 50
 rremb6 midm6 rem6 1
 xcompp6 vinp midp6 lp6 compare
 xcompm6 vinm midm6 lm6 compare
 * *comparator number 7*
 rrepa7 rep6 midp7 1
 rrepb7 midp7 rep7 1
 rrema7 rem6 midm7 1
 rremb7 midm7 rem7 1
 xcompp7 vinp midp7 lp7 compare
 xcompm7 vinm midm7 lm7 compare 60
 * *comparator number 8*
 rrepa8 rep7 midp8 1
 rrepb8 midp8 rep8 1
 rrema8 rem7 midm8 1
 rremb8 midm8 rem8 1
 xcompp8 vinp midp8 lp8 compare
 xcompm8 vinm midm8 lm8 compare
 * *comparator number 9*
 rrepa9 rep8 midp9 1
 rrepb9 midp9 rep9 1 70
 rrema9 rem8 midm9 1
 rremb9 midm9 rem9 1
 xcompp9 vinp midp9 lp9 compare
 xcompm9 vinm midm9 lm9 compare
 * *comparator number 10*
 rrepa10 rep9 midp10 1
 rrepb10 midp10 rep10 1
 rrema10 rem9 midm10 1
 rremb10 midm10 rem10 1
 xcompp10 vinp midp10 lp10 compare 80
 xcompm10 vinm midm10 lm10 compare
 * *comparator number 11*
 rrepa11 rep10 midp11 1
 rrepb11 midp11 rep11 1
 rrema11 rem10 midm11 1
 rremb11 midm11 rem11 1
 xcompp11 vinp midp11 lp11 compare
 xcompm11 vinm midm11 lm11 compare
 * *comparator number 12*
 rrepa12 rep11 midp12 1 90
 rrepb12 midp12 rep12 1
 rrema12 rem11 midm12 1
 rremb12 midm12 rem12 1
 xcompp12 vinp midp12 lp12 compare
 xcompm12 vinm midm12 lm12 compare
 * *comparator number 13*
 rrepa13 rep12 midp13 1
 rrepb13 midp13 rep13 1
 rrema13 rem12 midm13 1
 rremb13 midm13 rem13 1 100
 xcompp13 vinp midp13 lp13 compare
 xcompm13 vinm midm13 lm13 compare
 * *comparator number 14*
 rrepa14 rep13 midp14 1

```

rrepb14 midp14 rep14 1
rrema14 rem13 midm14 1
rremb14 midm14 rem14 1
xcomp14 vinp midp14 lp14 compare
xcompm14 vinm midm14 lm14 compare
* comparator number 15
rrep15 rep14 midp15 1
rrepb15 midp15 rep15 1
rrema15 rem14 midm15 1
rremb15 midm15 rem15 1
xcomp15 vinp midp15 lp15 compare
xcompm15 vinm midm15 lm15 compare
vrep1 rep15 0 1
vrem1 rem15 0 0

* latches
xlat1 lp1 lm1 th1 latch
xlat2 lp2 lm2 th2 latch
xlat3 lp3 lm3 th3 latch
xlat4 lp4 lm4 th4 latch
xlat5 lp5 lm5 th5 latch
xlat6 lp6 lm6 th6 latch
xlat7 lp7 lm7 th7 latch
xlat8 lp8 lm8 th8 latch
xlat9 lp9 lm9 th9 latch
xlat10 lp10 lm10 th10 latch
xlat11 lp11 lm11 th11 latch
xlat12 lp12 lm12 th12 latch
xlat13 lp13 lm13 th13 latch
xlat14 lp14 lm14 th14 latch
xlat15 lp15 lm15 th15 latch

* bubble correction logic array
vh0 th0 0 1
xith1 th1 a0 inv
xbub1 th0 th1 th2 a0 a1 c2 c1 d1 bubble
xbub2 th1 th2 th3 a1 a2 c3 c2 d2 bubble
xbub3 th2 th3 th4 a2 a3 c4 c3 d3 bubble
xbub4 th3 th4 th5 a3 a4 c5 c4 d4 bubble
xbub5 th4 th5 th6 a4 a5 c6 c5 d5 bubble
xbub6 th5 th6 th7 a5 a6 c7 c6 d6 bubble
xbub7 th6 th7 th8 a6 a7 c8 c7 d7 bubble
xbub8 th7 th8 th9 a7 a8 c9 c8 d8 bubble
xbub9 th8 th9 th10 a8 a9 c10 c9 d9 bubble
xbub10 th9 th10 th11 a9 a10 c11 c10 d10 bubble
xbub11 th10 th11 th12 a10 a11 c12 c11 d11 bubble
xbub12 th11 th12 th13 a11 a12 c13 c12 d12 bubble
xbub13 th12 th13 th14 a12 a13 c14 c13 d13 bubble
xbub14 th13 th14 th15 a13 a14 c15 c14 d14 bubble
xbub15 th14 th15 th16 a14 a15 c16 c15 d15 bubble
rthz th16 0 10
vhz c16 0 1

* short connection diagram
* bit0 = aaaa 0101010101010101
* bit1 = cccc 0011001100110011
* bit2 = 0f0f 0000111100001111
* bit3 = 00ff 0000000011111111

```

110

120

130

140

150

160

* final output binary encode using a ROM

* ROM transistors for output bit0

mrom1 rom0 d1 vdd vdd PMOS L=0.25u W=15u
mrom3 rom0 d3 vdd vdd PMOS L=0.25u W=15u
mrom5 rom0 d5 vdd vdd PMOS L=0.25u W=15u
mrom7 rom0 d7 vdd vdd PMOS L=0.25u W=15u
mrom9 rom0 d9 vdd vdd PMOS L=0.25u W=15u
mrom11 rom0 d11 vdd vdd PMOS L=0.25u W=15u
mrom13 rom0 d13 vdd vdd PMOS L=0.25u W=15u
mrom15 rom0 d15 vdd vdd PMOS L=0.25u W=15u
mromn0 rom0 vdd 0 0 NMOS L=0.25u W=3.8u
cload0 bitb0 0 0.05pF
mbitx0 bitb0 rom0 0 0 NMOS L=0.25u W=0.8u
mbitb0 bitb0 rom0 vdd vdd PMOS L=0.25u W=1.8u
mbitc0 bit0 bitb0 0 0 NMOS L=0.25u W=0.8u
mbitd0 bit0 bitb0 vdd vdd PMOS L=0.25u W=1.8u

170

180

* ROM transistors for output bit1

mrom102 rom1 d2 vdd vdd PMOS L=0.25u W=15u
mrom103 rom1 d3 vdd vdd PMOS L=0.25u W=15u
mrom106 rom1 d6 vdd vdd PMOS L=0.25u W=15u
mrom107 rom1 d7 vdd vdd PMOS L=0.25u W=15u
mrom110 rom1 d10 vdd vdd PMOS L=0.25u W=15u
mrom111 rom1 d11 vdd vdd PMOS L=0.25u W=15u
mrom114 rom1 d14 vdd vdd PMOS L=0.25u W=15u
mrom115 rom1 d15 vdd vdd PMOS L=0.25u W=15u
mromn1 rom1 vdd 0 0 NMOS L=0.25u W=3.8u
cload1 bitb1 0 0.05pF
mbitx1 bitb1 rom1 0 0 NMOS L=0.25u W=0.8u
mbitb1 bitb1 rom1 vdd vdd PMOS L=0.25u W=1.8u
mbitc1 bit1 bitb1 0 0 NMOS L=0.25u W=0.8u
mbitd1 bit1 bitb1 vdd vdd PMOS L=0.25u W=1.8u

190

* ROM transistors for output bit2

mrom204 rom2 d4 vdd vdd PMOS L=0.25u W=15u
mrom205 rom2 d5 vdd vdd PMOS L=0.25u W=15u
mrom206 rom2 d6 vdd vdd PMOS L=0.25u W=15u
mrom207 rom2 d7 vdd vdd PMOS L=0.25u W=15u
mrom212 rom2 d12 vdd vdd PMOS L=0.25u W=15u
mrom213 rom2 d13 vdd vdd PMOS L=0.25u W=15u
mrom214 rom2 d14 vdd vdd PMOS L=0.25u W=15u
mrom215 rom2 d15 vdd vdd PMOS L=0.25u W=15u
mromn2 rom2 vdd 0 0 NMOS L=0.25u W=3.8u
cload2 bitb2 0 0.05pF
mbitx2 bitb2 rom2 0 0 NMOS L=0.25u W=0.8u
mbitb2 bitb2 rom2 vdd vdd PMOS L=0.25u W=1.8u
mbitc2 bit2 bitb2 0 0 NMOS L=0.25u W=0.8u
mbitd2 bit2 bitb2 vdd vdd PMOS L=0.25u W=1.8u

200

210

* ROM transistors for output bit3

mrom308 rom3 d8 vdd vdd PMOS L=0.25u W=15u
mrom309 rom3 d9 vdd vdd PMOS L=0.25u W=15u
mrom310 rom3 d10 vdd vdd PMOS L=0.25u W=15u
mrom311 rom3 d11 vdd vdd PMOS L=0.25u W=15u
mrom312 rom3 d12 vdd vdd PMOS L=0.25u W=15u
mrom313 rom3 d13 vdd vdd PMOS L=0.25u W=15u
mrom314 rom3 d14 vdd vdd PMOS L=0.25u W=15u
mrom315 rom3 d15 vdd vdd PMOS L=0.25u W=15u
mromn3 rom3 vdd 0 0 NMOS L=0.25u W=3.8u
cload3 bitb3 0 0.05pF
mbitx3 bitb3 rom3 0 0 NMOS L=0.25u W=0.8u

220

```
mbitb3 bitb3 rom3 vdd vdd PMOS L=0.25u W=1.8u
mbitc3 bit3 bitb3 0 0 NMOS L=0.25u W=0.8u
mbitd3 bit3 bitb3 vdd vdd PMOS L=0.25u W=1.8u
.ends
```

** Th-th-th-th-th-that's all, Folks!*

Appendix E

Agilent 93000 SOC Tester

E.1 AC Testing with the Agilent 93000

High-speed AC testing of the analog-to-digital converters was completed on an Agilent 93000 SOC Tester.

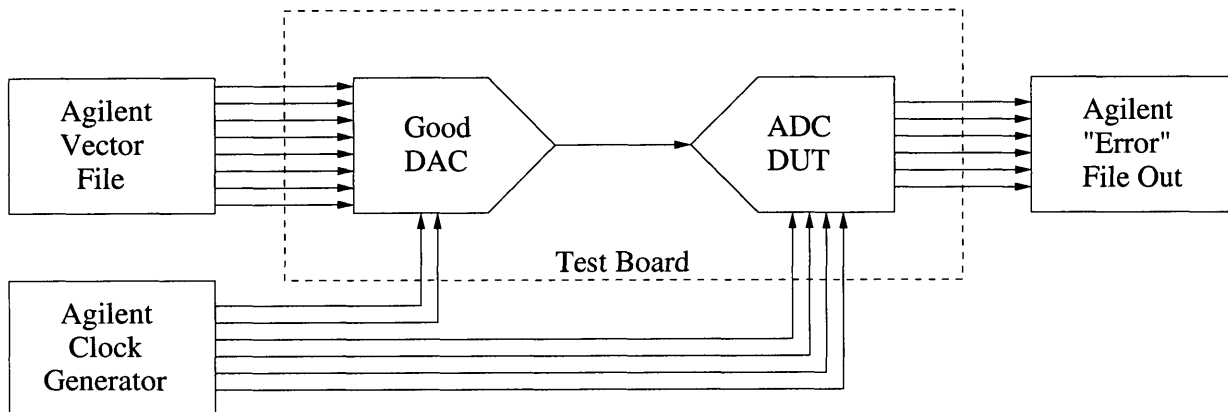


Figure E.1: Block diagram of test board interface to Agilent 93000 tester.

Since MIT Lincoln Laboratory's tester does not have any analog outputs, a test board was constructed that included its own digital-to-analog front end, using a Philips TDA8776A DAC [114]. Software to convert the Agilent output files to plots of S/N+D, SFDR, INL and noise

histograms were written in Octave.

E.2 Agilent Vector Format and Pin Assignments

This file is documentation, rather than a machine-readable input file. It explains the individual columns in the Agilent vector file, documents the mapping from logical names to physical pin numbers, and lists the necessary termination and voltage levels.

Vector File Format and Pin Assignments

This file describes the file format of the vector files used for testing analog-to-digital converters on the Agilent 93000 tester.

The format line, as it appears in the vector file, is:

```
FORMAT clocks bits dacout valid adcout valbar adcbar error strig ;
```

```
#          clocks 10-bit-dac do data-bit-from-adc E S
#          ddaabb 9876543210 pn V 543210 C 543210 E T
# testname enable oooooooooo ii i iiiiii i iiiiii i o
```

Here is a short explanation of each entry on the FORMAT line, as well as which Agilent pins the signals correspond to:

TESTNAME: adctest

CLOCKS ENABLE: six clocks lines drive the test board

```
clocks5 10116 (DAC board clock positive)
clocks4 10115 (DAC board clock negative, complementary to 5)
clocks3 10715 (ADC analog clock a positive)
clocks2 10501 (ADC analog clock a negative, complementary to 3)
clocks1 10503 (ADC digital clock b positive, in quadrature with 3)
clocks0 10505 (ADC digital clock b negative, complementary to 1)
```

DAC BITS INPUT: ten bit lines, Agilent outputs (10)

```
bits9 10105
bits8 10106
bits7 10107
bits6 10108
bits5 10109
bits4 10110
bits3 10111
bits2 10112
bits1 10113
bits0 10114
```

DAC OUT: analog output from dac, Agilent inputs (2), ignored

```
dacout1 10103 (DAC out negative)
dacout0 10104 (DAC out positive)
```


VALID: data-ready output from d.u.t.adc, Agilent input, ignored
terminated through 50 ohms to -1.0 volt
valid 10711 (ADC data valid signal)

ADC OUT: digital output from d.u.t. adc, Agilent inputs (6), ignored
terminated through 50 ohms to -1.0 volt
adcout5 10513 (ADC bit 5)
adcout4 10514 (ADC bit 4)
adcout3 10705 (ADC bit 3)
adcout2 10703 (ADC bit 2)
adcout1 10707 (ADC bit 1)
adcout0 10710 (ADC bit 0)

VALID-BAR: complement of data-ready output, Agilent input, ignored
terminated through 50 ohms to -1.0 volt
valbar 10712 (ADC data valid, complement)

ADC-BAR: complements of digital outputs from adc, Agilent inputs (6), ignored
terminated through 50 ohms to -1.0 volt
adcbar5 10515 (ADC bit 5, complement)
adcbar4 10516 (ADC bit 4, complement)
adcbar3 10706 (ADC bit 3, complement)
adcbar2 10704 (ADC bit 2, complement)
adcbar1 10709 (ADC bit 1, complement)
adcbar0 10713 (ADC bit 0, complement)

ERROR: this input pin, tied to ground, is used to create an error for
every line in the vector file, so the Agilent prints out every line.
error 10102

SCOPE TRIGGER: scope trigger, Agilent output
strig 10101

POWER SUPPLIES: see text for schematic of power connections
vee DSP11 -5.2V
vss DSP12 -1.0V
vdd2 DSP13 +1.0V
vcc DSP14 +6.0V

May 22, 2002 -- Kent Lundberg
June 18, 2002 -- Kent Lundberg

E.3 Agilent Output File Format

In high-speed mode, the Agilent 93000 SOC Tester doesn't actually capture the output data from the converter; it only keeps track of the errors. The tester is optimized for the testing of digital systems with deterministic outputs, where the input vector file would contain all of the test inputs and the corresponding expected outputs. The tester would then run a complete test on the digital system, and would only keep track of the failures where the actual output did not match the expected output.

This test method does not apply to the testing of analog-to-digital converters. Due to several factors — offset error, gain error, nonlinearities, aperture jitter, noise — the digital output will not match the expected output.

In low-speed mode, for clock frequencies below 166 MHz, it is possible to capture the output data in JWA files. This method was used to most of the converter testing in Chapter 7.

E.3.1 Agilent JWA to Octave Converter

```
print "# OCTAVE data converted from Agilent vector file\n";
print "# command line was: $0";
foreach $item (@ARGV) {print " $item";}
print "\n";
print "# name: vector\n# type: matrix\n";
print "# rows: 1\n# columns: 4095\n";

while (<>) {
  if (/([HL][HL][HL][HL][HL][HL])/) {
    $bits = $1;
    $bits =~ y/HL/01/; # oops
    #print $bits;
    printf " %d\n",unpack("c",pack("B8","00".$bits));
  }
}

# perl jwa.pl kent1.jwa > kent1.dat
# in gnuplot: plot "kent1.dat"
```

10

E.3.2 Sample Agilent JWA Output File

```
FORMAT bits9 bits8 bits7 bits6
bits5 bits4 bits3 bits2 bits1
bits0 clocks5 clocks4 clocks3 clocks2
clocks1 clocks0 strig adcbar5 adcbar4
adcbar3 adcbar2 adcbar1 adcbar0 adcout5
adcout4 adcout3 adcout2 adcout1 adcout0
dacout1 dacout0 error valbar valid;

R1 adctst 01111111111111111111XXXXXXHHHHHHXXHXL;
R1 adctst 100110001111111110XXXXXXHHHLHXXHXL;
R1 adctst 101100001011111110XXXXXXHLLHHHXXHXL;
R1 adctst 110001101011111110XXXXXXLHHHLHXXHXL;
R1 adctst 110110011111111110XXXXXXLHLLHHHXXHXL;
R1 adctst 111010011111111110XXXXXXLHLLHXXHXL;
R1 adctst 111101011111111110XXXXXXLHHHLHXXHXL;
R1 adctst 111110101111111110XXXXXXLHLLHHHXXHXL;
R1 adctst 111111111111111110XXXXXXLHLLHXXHXL;
R1 adctst 111111011111111110XXXXXXLHLLHXXHXL;
R1 adctst 111101101111111110XXXXXXLHLLHXXHXL;
R1 adctst 111010111011111110XXXXXXLHLLHXXHXL;
R1 adctst 1101100001111110XXXXXXLHLLHXXHXL;
R1 adctst 11001001001111110XXXXXXLHLLHHHXXHXL;
R1 adctst 10110011101111110XXXXXXLHLLHHHXXHXL;
R1 adctst 10011011111111110XXXXXXLHHHHXXHXL;
R1 adctst 10000011001111110XXXXXXLHLLHHHXXHXL;
R1 adctst 0110101001111110XXXXXXLHLLHHHXXHXL;
R1 adctst 0101001001111110XXXXXXHHHHHXXHXL;
R1 adctst 0011100001111110XXXXXXHLLHHHXXHXL;
R1 adctst 0010100001111110XXXXXXHLLHHHXXHXL;
R1 adctst 0001011111111110XXXXXXHLLHXXHXL;
R1 adctst 0000101101111110XXXXXXHLLHXXHXL;
R1 adctst 0000001101111110XXXXXXHLLHXXHXL;
R1 adctst 0000000001111110XXXXXXHHHLHXXHXL;
R1 adctst 0000000110111110XXXXXXHHHLHXXHXL;
R1 adctst 0000011111111110XXXXXXHHHLHXXHXL;
R1 adctst 0001001010111110XXXXXXHHHLHXXHXL;
R1 adctst 0010000110111110XXXXXXHHHLHXXHXL;
R1 adctst 0011010000111110XXXXXXHHHLHXXHXL;
R1 adctst 0100100110111110XXXXXXHLLHXXHXL;
R1 adctst 0110000100111110XXXXXXHHHHHXXHXL;
```

E.4 Vector File Generator in Perl

Here is the perl code for the program that generates all of the vector files used in this testing. It produces sine waves, sawtooth waves, square waves, random noise, two-tone inputs, and DC levels. It can produce vector files of any amplitude, length and periodicity for HSPICE, the Agilent 93000 tester, Gnuplot and Octave. The script also produces crude histograms for debugging purposes.

```
# vector.pl
#
# Produces vector files for HSPICE and Agilent tester
# Usage: perl vector.pl format wavetype length cycles pretty > datafile.vec
#
# written 2000,2002 by Kent Lundberg

$bits = 10;
$max = 2**$bits;      # amplitude of waveform (probably 2^n)
$pi = atan2(1,1) * 4; # 3.1415926535...
srand;                # seed the random number generator

@wavename = ("sine","sawtooth","square","noise","two-tone","DC");

$format = $ARGV[0];
$wave = $ARGV[1];
$length = $ARGV[2]; # total number of points (lines in file)
$cycles = $ARGV[3]; # complete cycles of periodic waveform
$amp = $ARGV[4];
$pretty = $ARGV[5];

if ($cycles==0) {
    $loop = 1;
    select(STDERR);
    print "Usage: $0 format wavetype length cycles amplit pretty\n";
    print " formats: 1=hspice, 2=agilent, 3=gnuplot, 4=octave, 5=histog\n";
    print " wavetypes: ";
    foreach $item (@wavename) {print $loop++,"=$item, ";}
    print "\n length: number of data points in file\n";
    print " cycles: number of cycles of periodic waveforms (must be > 0)\n";
    print " amplit: amplitude (for sine, DC), in fraction of full-scale\n";
    print " pretty: if nonzero, prints ascii waveform using comments\n";
    die "\n$0: I need brains!!! and arguments...\n";
}

if ($format == 1) {$hspice=1;}
if ($format == 2) {$agilent=1;}
if ($format == 3) {$gnuplot=1;}
if ($format == 4) {$octave=1;}
if ($format == 5) {$histog=1;}

if ($hspice) {$c = ";";} # vector file comment character
else {$c = "#";}

```

```

print "$c VECTOR FILE\n$c Autogenerated $wavename[$wave-1] wave, ";
print "$length points, $cycles cycles, amplitude = $max.\n";
print "$c Command line was: $0";
foreach $item (@ARGV) {print " $item";}
print "\n";

if ($hspace) { # Proper header for HSPICE vector files
print "; Generating for HSPICE... \n\nname";
for ($loop = $bits; $loop > 0; $loop--) {
    print " bit", $loop-1;
}
print "\n";
print "radix ", "1"x$bits, " ; $bits bits of data\n";
print "vih 1 ", "1"x$bits, " ; logic high = 1.0 volts\n";
print "io ", "i"x$bits, " ; all $bits bits are inputs\n";
print "period 1\ntunit ns\n\n";
$line = 64;
}

if ($agilent) { # Proper header for AGILENT vector files
print "# Generating for Agilent 93000 tester... \n\n";
print "FORMAT clocks bits dacout valid adcout valbar adcbar error strig ;\n\n";
print "# clocks 10-bit-dac do data-bit-from-adc E S\n";
print "# ddaabb 9876543210 pn V 543210 C 543210 E T\n";
print "# tstname enable oooooooooo ii i iiiiii i iiiiii i o\n\n";
$line = 24;
$strig = 1;
}

if ($gnuplot) { # Proper header for GNUPLOT data files
print "# Generating for gnuplot... \n\n";
$line = 64;
}

if ($octave) { # Proper header for OCTAVE data files
print "# Generating for octave... \n";
print "# name: vector\n# type: matrix\n";
print "# rows: 1\n# columns: $length\n";
$pretty = 0;
}

if ($histog) { # Proper header for OCTAVE histogram files
print "# Generating a histogram for octave... \n\n";
print "# name: vector\n# type: matrix\n";
print "# rows: 1\n# columns: $max\n";
$pretty = 0;
}

for ($loop = 0; $loop < $length; $loop++) {
    $index = $loop*$cycles/$length;
    # sine wave
    if ($wave==1) {$val = ($max/2-0.01)*(1+$amp*sin(2*$pi*$index));}
    # sawtooth wave
    if ($wave==2) {$val = ($max*$index) % $max;}
    # square wave
    if ($wave==3) {$val = ($max-1)*((2*$index+1) % 2);}
    # gaussian white noise
    if ($wave==4) {

```

```

    $val = 0;
    for ($lo = 0; $lo < $max; $lo++) {
        $val += rand;
    }
}
# flat PDF noise: $val = rand $max;
# two-tone wave
if ($wave==5) {
    $index = $loop*($cycles+128)/$length;
    $val = ($max/4-0.01)*(2+sin(2*$pi*$index)+sin(2*$pi*$index));
}
# simple DC value
if ($wave==6) {
    $val = ($max-1)*$amp;
}

# There has GOT TO BE an easier, more elegant way to do this conversion!
# Pack $val into an integer, unpack that integer into n bits of
# low-to-high binary, and then reverse that into high-to-low binary.
# This method works on Linux, but not on Solaris. WTF?

$ddata = reverse unpack("b$bits",pack("i",$val));

if ($hspice) {print $ddata;}
if ($agilent) {
    print "R1 adctst 111111 \"$ddata,\" LL L LLLLLL L LLLLLL H \"$strig,\"";
    $strig = 0;
}
if ($gnuplot) {print int($val);}
if ($octave) {print " ",int($val);}
if ($pretty) {print " "x($line*$val/$max+1),"c";}

if ($histog) {$hist[int($val)] += 1;}
else {print "\n";}

# if ($hspice| $agilent| $gnuplot) {print "\n";}
}

if ($histog) { # print data for histogram
    $histmax = 0;
    $digits = int(log($length)/2.3);
    for ($loop = 0; $loop < $max; $loop++) {
        if ($hist[$loop]>$histmax) {$histmax = $hist[$loop];}
    }
    for ($loop = 0; $loop < $max; $loop++) {
        printf "%${digits}d\n",$hist[$loop];
        #print " "x($line*$hist[$loop]/$histmax+1),"c\n";
    }
}
}

```

110

120

130

140

150

E.5 Plots of Available Vectors

Here is a list of the available waveform from vector.pl that can be used as test inputs.

E.5.1 Sine Wave Vector

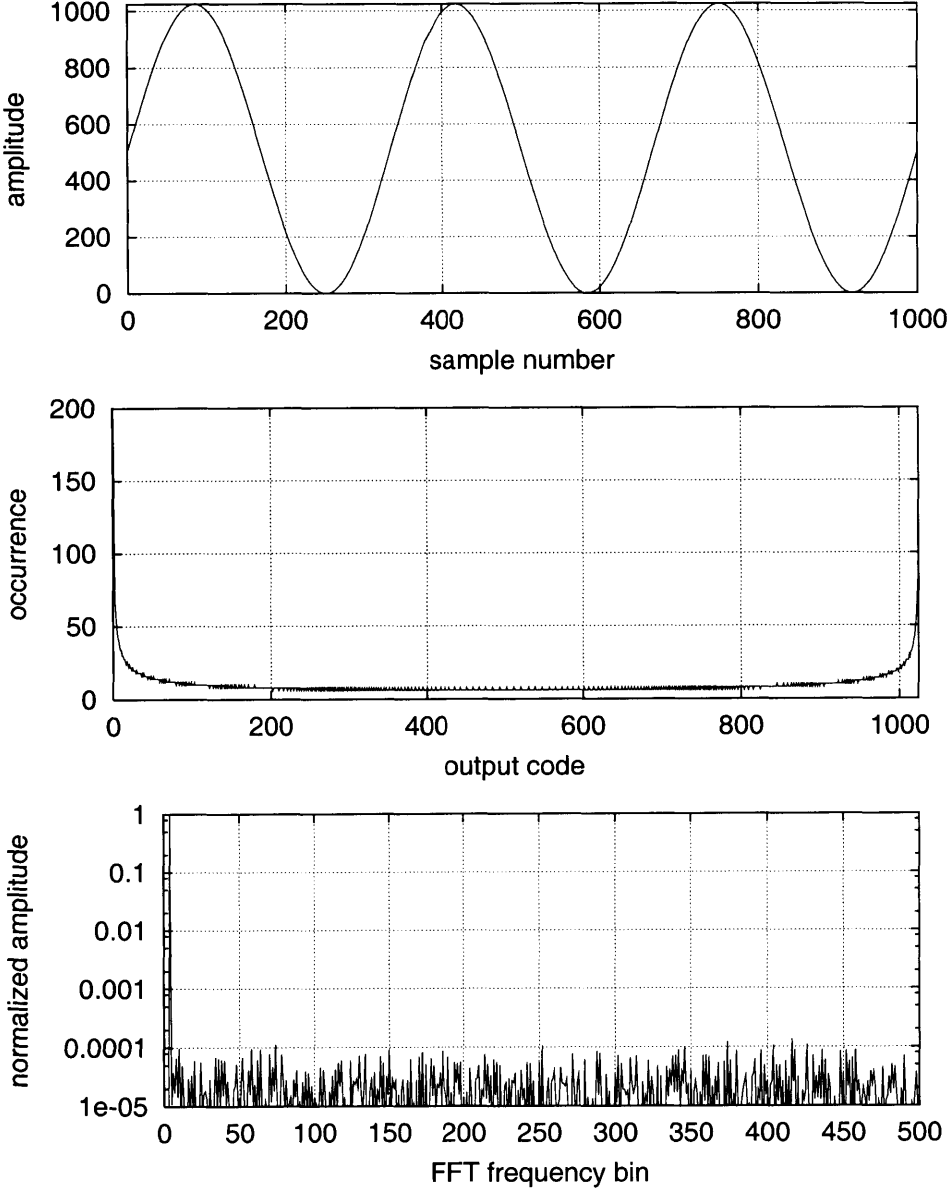


Figure E.2: Sine wave plot, histogram, and FFT from vector file data.

E.5.2 Sawtooth Wave Vector




Figure E.3: Sawtooth wave plot and FFT from vector file data.

E.5.3 Square Wave Vector

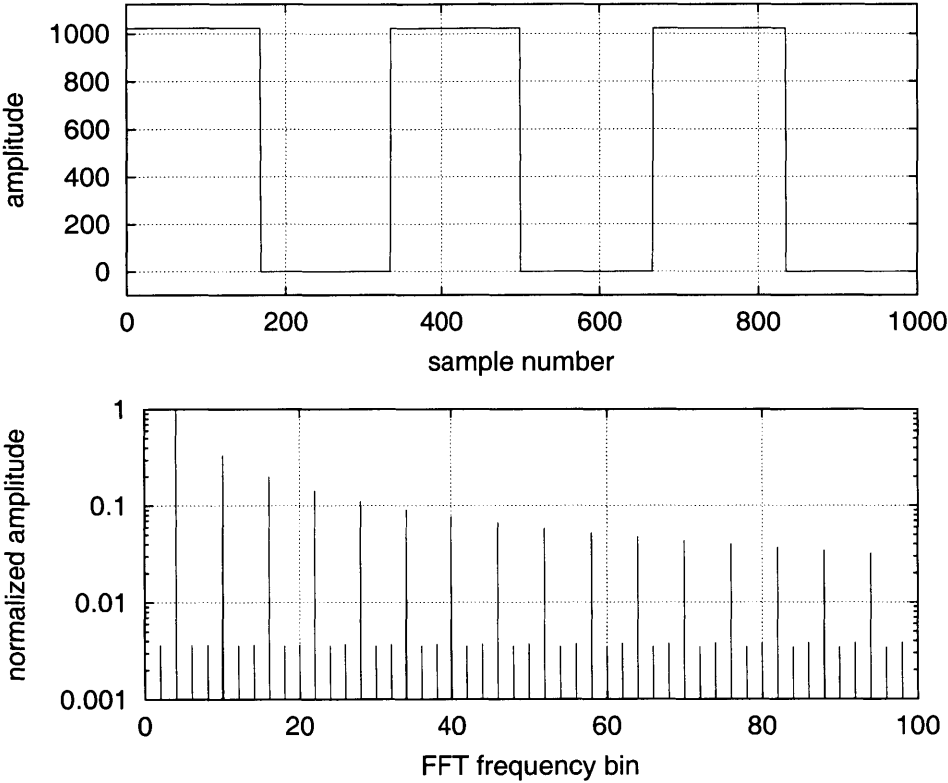


Figure E.4: Square wave plot and FFT from vector file data.

E.5.4 White Noise Vector




Figure E.5: White noise plot, histogram, and FFT from vector file data.

E.5.5 Two-tone Sine Wave Vector




Figure E.6: Two-tone sine wave plot, histogram, and FFT from vector file data.

E.6 Sample HSPICE Vector File

```
; VECTOR FILE
; Autogenerated sine wave, 50 points, 3 cycles, amplitude = 1024.
; Command line was: vector.pl 1 1 50 3 1
; Generating for HSPICE...
```

```
vname bit9 bit8 bit7 bit6 bit5 bit4 bit3 bit2 bit1 bit0
radix 1111111111 ; 10 bits of data
vih 1 1111111111 ; logic high = 1.0 volts
io iiiiiiiiii ; all 10 bits are inputs
period 1
tunit ns
```

```
0111111111 ;
1010111100 ;
1101011110 ;
1111001111 ;
1111111110 ;
1111100110 ;
1110001010 ;
1011110110 ;
1001000000 ;
0110000000 ;
0011010011 ;
0001001111 ;
000001001 ;
000001001 ;
0001001111 ;
0011010011 ;
0110000000 ;
1001000000 ;
1011110110 ;
1110001010 ;
1111100110 ;
1111111110 ;
1111001111 ;
1101011110 ;
1010111100 ;
0111111111 ;
0101000011 ;
0010100001 ;
0000110000 ;
0000000001 ;
0000011001 ;
0001110101 ;
0100001001 ;
0110111111 ;
1001111111 ;
1100101100 ;
1110110000 ;
1111110110 ;
1111110110 ;
1110110000 ;
1100101100 ;
1001111111 ;
0110111111 ;
0100001001 ;
0001110101 ;
0000011001 ;
0000000001 ;
0000110000 ;
0010100001 ;
0101000011 ;
```

E.7 Sample Agilent Vector File

```
# VECTOR FILE
# Autogenerated sine wave, 50 points, 3 cycles, amplitude = 1024.
# Command line was: vector.pl 2 1 50 3 1
# Generating for Agilent 93000 tester...

FORMAT clocks bits dacout valid adcout valbar adcbare error strig ;

#          clocks 10-bit-dac do data-bit-from-adc E S
#          ddaabb 9876543210 pn V 543210 C 543210 E T
# testname enable oooooooooo ii i iiiiii i iiiiii i o

R1 adctest 111111 0111111111 XX L LLLLLL X XXXXXX H 1;      #
R1 adctest 111111 1010111100 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1101011110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1111001111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1111111110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1111100110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1110001010 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1011110110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1001000000 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0110000000 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0011010011 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0001001111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0000001001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0000001001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0001001111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0011010011 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0110000000 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1011110110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1110001010 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1111100110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1111111110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1111001111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1101011110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1010111100 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0111111111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0101000011 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0010100001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0000110000 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0000000001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0000011001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0001110101 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0100001001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0110111111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1001111111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1100101100 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1110110000 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1111110110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1111110110 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1110110000 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1100101100 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 1001111111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0110111111 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0100001001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0001110101 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0000011001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0000000001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0000110000 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0010100001 XX L LLLLLL X XXXXXX H 0;      #
R1 adctest 111111 0101000011 XX L LLLLLL X XXXXXX H 0;      #
```


Colophon

A conscious effort was made to produce this document, and the work described herein, using only free software (that is, free as in speech). While meeting this goal was not entirely possible, most of the software used was *software libre*.

Free Software

This document was produced in L^AT_EX 2 ϵ (which is based on T_EX) with files edited in GNU Emacs. Schematic drawings were produced by Xcircuit, plots were produced by Octave and Gnuplot, and other diagrams were drawn with Xfig. The software shown in Appendices D and E was written in Perl. All of this software was run on a workstation running the SuSE variety of GNU/Linux.

Non-Free Software

Unfortunately, not all of the tools required were available as free software. The non-free tools used included the Silvaco SmartSpice and the Avant! Star-HSPICE simulation programs, the Mentor Graphics IC layout software, and the testing software on the Agilent SOC 93000 tester.

Lgrind, used to pretty-print the Perl and SPICE code in Appendices D and E, is technically non-free due to an idiotic (but likely unintentional) licensing restriction by one of its contributors, Van Jacobson.

Support

As a display of gratitude for making this document possible, \$1000 was given to the Electronic Frontier Foundation¹ and \$500 was given to the Free Software Foundation² and the University of Wisconsin Foundation for Octave³ at the completion of this work.

¹<http://www.eff.org>

²<http://www.fsf.org>

³<http://www.octave.org>

Bibliography

- [1] Frank Goodenough. Analog technology of all varieties dominate ISSCC. *Electronic Design*, 44(4):96–111, February 19, 1996.
- [2] Robert H. Walden. Analog-to-digital converter technology comparison. In *IEEE GaAs IC Symposium, Technical Digest*, pages 217–219, October 1994.
- [3] Robert H. Walden. Analog-to-digital converter survey and analysis. *IEEE Journal on Selected Areas in Communication*, 17(4):539–550, April 1999.
- [4] Govert Geelen. A 6b 1.1GSample/s CMOS A/D converter. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 128–129, February 2001.
- [5] Peter Scholtens and Maarten Vertregt. A 6b 1.6GSample/s flash ADC in 0.18 μ m CMOS using averaging termination. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 168–169, February 2002.
- [6] Joshua W. Phinney. Student Evaluation for MIT Course 6.331: Advanced Circuit Techniques, May 2002. In his evaluation, Mr. Phinney wrote: “Recitation Instructor [Kent Lundberg] is hopelessly pedantic.” This comment was taken as a compliment.
- [7] Declan Dalton, G. Spalding, H. Reyhani, T. Murphy, K. Deevy, M. Walsh, and P. Griffin. A 200-MSPS 6-bit flash ADC in 0.6- μ m CMOS. *IEEE Transactions of Circuits and Systems — II: Analog and Digital Signal Processing*, 45(11):1433–1444, November 1998.
- [8] P. Setty, J. Barner, J. Plany, H. Burger, and J. Sonntag. A 5.75b 350MSample/s or 6.75b 150MSample/s reconfigurable flash ADC for a PRML read channel. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 148–149, February 1998.
- [9] Sanroku Tsukamoto, William G. Schofield, and Toshiaki Endo. A CMOS 6-b, 400-MSample/s ADC with error correction. *IEEE Journal of Solid-State Circuits*, 33(12):1939–1947, December 1998.
- [10] Michael Flynn and Ben Sheahan. A 400-Msample/s, 6-b CMOS folding and interpolating ADC. *IEEE Journal of Solid-State Circuits*, 33(12):1932–1938, December 1998.
- [11] Yuko Tamba and Kazuo Yamakido. A CMOS 6b 500MSample/s ADC for hard disk drive read channel. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 324–325, February 1999.
- [12] Kwangho Yoon, Sungkyung Park, and Wonchan Kim. A 6b 500MSample/s CMOS flash ADC with a background interpolated auto-zeroing technique. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 326–327, February 1999.

- [13] Iuri Mehr and Declan Dalton. A 500-MSample/s 6-bit nyquist-rate ADC for disk-drive read-channel applications. *IEEE Journal of Solid-State Circuits*, 34(7):912–920, July 1999.
- [14] Peter Scholtens. A 2.5 volt 6 bit 600MS/s flash ADC in 0.25 μ m CMOS. In *Proceedings of the European Solid-State Circuits Conference*, pages 196–199, September 2000.
- [15] K. Nagaraj, D. Martin, M. Wolfe, R. Chattopadhyay, A. Pavan, J. Cancio, and T. Viswanathan. A dual-mode 700MSample/s 6-bit 200MSample/s 7-bit A/D converter in a 0.25- μ m digital CMOS process. *IEEE Journal of Solid-State Circuits*, 35(12):1760–1768, December 2000.
- [16] Kouji Sushihara, H. Kimura, Y. Okamoto, K. Nishimura, and A. Matsuzawa. A 6b 800MSample/s CMOS A/D converter. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 428–429, February 2000.
- [17] Michael Choi and Asad A. Abidi. A 6-b 1.3Gsample/s A/D converter in 0.35- μ m CMOS. *IEEE Journal of Solid-State Circuits*, 36(12):1847–1858, December 2001.
- [18] Kevin R. Nary, R. Nubling, S. Beccue, W. Colleran, J. Penney, and K. Wang. An 8-bit 2 gigasample per second analog to digital converter. In *IEEE GaAs IC Symposium, Technical Digest*, volume 17, pages 303–306, October 1995.
- [19] Ken Poulton, K. Knudsen, J. Corcoran, K. Wang, R. Nubling, R. Pierson, M. Chang, P. Asbeck, and R. Huang. A 6-b, 4 GSA/s GaAs HBT ADC. *IEEE Journal of Solid-State Circuits*, 30(10):1109–1118, October 1995.
- [20] Cynthia Baringer, Joe Jensen, Larry Burns, and Bob Walden. 3-bit, 8 GSPS flash ADC. In *Proceedings of the International Conference on Indium Phosphide and Related Materials*, pages 64–67, April 1996.
- [21] Koichi Irie, N. Kusayanagi, T. Kawachi, T. Nishibu, and Y. Matsumori. An 8b 500MS/s full nyquist cascade A/D converter. In *Symposium on VLSI Circuits, Digest of Technical Papers*, pages 77–78, June 1999.
- [22] Ken Poulton, K. Knudsen, J. Kerley, J. Kang, J. Tani, E. Cornish, and M. VanGrouw. An 8-GSa/s 8-bit ADC system. In *Symposium on VLSI Circuits, Digest of Technical Papers*, pages 23–24, June 1997.
- [23] Ken Poulton, R. Neff, A. Muto, W. Liu, A. Burstein, and M. Heshami. A 4GSample/s 8b ADC in 0.35- μ m CMOS. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 166–167, February 2002.
- [24] Fairchild Semiconductor. SPT7760 8-bit, 1 GSPS, flash A/D converter. *Datasheet*, 1995.
- [25] Robert H. Walden. Spreadsheet of data for ADC survey. Available from HRL web site at <http://www.hrl.com/TECHLABS/micro/ADC/adc.html>, July 1999. Includes data for Rockwell RSC-ADC080S.
- [26] Maxim Integrated Products. MAX108 \pm 5V, 1.5GSPS, 8-bit ADC with on-chip 2.2GHz track/hold amplifier. *Datasheet*, 2001.
- [27] Eric A. Vittoz. Future of analog in the VLSI environment. In *Proceedings of the IEEE International Symposium on Circuits and Systems*, pages 1372–1375, May 1990.

- [28] Thomas Byunghak Cho and Paul R. Gray. A 10 b 20 Msample/s 35 mW pipeline A/D converter. *IEEE Journal of Solid-State Circuits*, 30(3):166–172, March 1995.
- [29] Peter Kinget and Michiel Steyaert. Impact of transistor mismatch on the speed-accuracy-power trade-off of analog CMOS circuits. In *Proceedings of the IEEE Custom Integrated Circuits Conference*, pages 333–336, May 1996.
- [30] Marcel J.M. Pelgrom, Hans P. Tuinhout, and Maarten Vertregt. Transistor matching in analog CMOS applications. In *International Electron Devices Meeting, Technical Digest*, pages 915–918, December 1998.
- [31] M. Steyaert, V. Peluso, J. Bastos, P. Kinget, and W. Sansen. Custom analog low power design: the problem of low voltage and mismatch. In *Proceedings of the IEEE Custom Integrated Circuits Conference*, pages 285–292, May 1997.
- [32] Willy Sansen. Analog circuit design in scaled CMOS technology. In *Symposium on VLSI Circuits, Digest of Technical Papers*, pages 8–11, June 1996.
- [33] Koen Uyttenhove and Michiel Steyaert. Speed-power-accuracy trade-off in high-speed ADC's: what about nano-electronics? In *Proceedings of the IEEE Custom Integrated Circuits Conference*, pages 341–344, May 2001.
- [34] D. Fu, K. Dyer, S. Lewis, and P. Hurst. Digital background calibration of a 10b 40MSample/s parallel pipelined ADC. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 140–141, February 1998.
- [35] K. Dyer, D. Fu, S. Lewis, and P. Hurst. Analog background calibration of a 10b 40MSample/s parallel pipelined ADC. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 142–143, February 1998.
- [36] Kerry Bernstein and Norman J. Rohrer. *SOI Circuit Design Concepts*. Kluwer Academic Publishers, Boston, 2000.
- [37] Jean-Pierre Colinge. *Silicon-on-Insulator Technology: Materials to VLSI*. Kluwer Academic Publishers, Boston, 1997.
- [38] Sorin Cristoloveanu and Sheng S. Li. *Electrical Characterization of Silicon-On-Insulator Materials and Devices*. Kluwer Academic Publishers, Boston, 1995.
- [39] Zachary J. Lemnios and Kaigham J. Gabriel. Low-power electronics. *IEEE Design and Test of Computers*, 11(4):8–13, 1994.
- [40] Jean-Pierre Colinge. Thin-film SOI technology: the solution to many submicrom CMOS problems. In *International Electron Devices Meeting, Technical Digest*, pages 817–820, December 1989.
- [41] Robert F. Pierret. *Semiconductor Device Fundamentals*. Addison-Wesley, Reading, Massachusetts, 1996.
- [42] Yannis P. Tsividis. *Operation and Modeling of the MOS Transistor*. McGraw-Hill, New York, 1987.

- [43] D. Flandre, B. Gentinne, J. P. Eggermont, and P. Jespers. Design of thin-film fully-depleted SOI CMOS analog circuits significantly outperforming bulk implementations. In *Proceedings of the IEEE International SOI Conference*, pages 99–100, October 1994.
- [44] J. P. Colinge, J. P. Eggermont, D. Flandre, P. Francis, and P. G. A. Jespers. Potential of SOI for analog and mixed analog-digital low-power applications. In *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, pages 194–195, February 1995.
- [45] D. Flandre, L. F. Ferreira, P. G. A. Jespers, and J. P. Colinge. Modelling and application of fully depleted SOI MOSFETs for low voltage, low power analogue CMOS circuits. *Solid-State Electronics*, 39(4):455–460, 1996.
- [46] Jean-Pierre Colinge. Fully-depleted SOI CMOS for analog applications. *IEEE Transactions on Electron Devices*, 45(5):1010–1016, May 1998.
- [47] D. Flandre et al. Fully-depleted SOI CMOS technology for low-voltage low-power mixed digital/analog/microwave circuits. *Analog Integrated Circuits and Signal Processing*, 21:213–228, 1999.
- [48] Benard M. Tenbroek, Michael S. L. Lee, William Redman-White, R. John T. Bunyan, and Michael J. Uren. Impact of self-heating and thermal coupling on analog circuits in SOI CMOS. *IEEE Journal of Solid-State Circuits*, 33(7):1037–1046, July 1998.
- [49] O Rozeau, J. Jomaah, S. Haendler, J. Boussey, and F. Balestra. SOI technologies overview for low-power low-voltage radio-frequency applications. *Analog Integrated Circuits and Signal Processing*, 25:93–114, 2000.
- [50] Federico Faccio, Francis Anghinolfi, Erik H. M. Heijne, Pierre Jarron, and Sorin Cristoloveanu. Noise contribution of the body resistance in partially depleted SOI MOSFET's. *IEEE Transactions on Electron Devices*, 45(5):1033–1038, May 1998.
- [51] Wei Jin, Philip C. H. Chan, Samuel K. H. Fung, and Ping K. Ko. Shot-noise-induced excess low-frequency noise in floating-body partially depleted SOI MOSFET's. *IEEE Transactions on Electron Devices*, 46(6):1180–1185, June 1999.
- [52] Ying-Che Tseng, W. Margaret Huang, Vida Ilderem, and Jason C. S. Woo. Floating body induced pre-kink excess low-frequency noise in submicron SOI CMOSFET technology. *IEEE Electron Device Letters*, 20(9):484–486, September 1999.
- [53] James B. Kuo and Shih-Chia Lin. *Low-Voltage SOI CMOS VLSI Devices and Circuits*. Wiley, New York, 2001.
- [54] Andrew Marshall and Sreedhar Natarajan. *SOI Design: Analog, Memory, and Digital Techniques*. Kluwer Academic Publishers, Boston, 2002.
- [55] Eric A. Vittoz. Micropower techniques. In Jose E. Franca and Yannis Tsividis, editors, *Design of Analog-Digital VLSI Circuits for Telecommunications and Signal Processing*, chapter 3, pages 53–96. Prentice-Hall, Englewood Cliffs, 1994.
- [56] Alan B. Grebene. *Bipolar and MOS Analog Integrated Circuit Design*. Wiley, New York, 1984.
- [57] Roubik Greorian. *Introduction to CMOS Op-Amps and Comparators*. Wiley, New York, 1999.

- [58] Jan M. Rabaey. *Digital Integrated Circuits: A Design Perspective*. Prentice Hall, Upper Saddle River, 1996.
- [59] Christopher W. Mangelsdorf. A 400-MHz input flash converter with error correction. *IEEE Journal of Solid-State Circuits*, 25(1):184–191, February 1990.
- [60] Christian C. Enz and Gabor C. Temes. Circuit techniques for reducing the effects of op-amp imperfections: autozeroing, correlated double sampling, and chopper stabilization. *Proceedings of the IEEE*, 84(11):1584–1614, November 1996.
- [61] Yen S. Yee, Lewis M. Terman, and Lawrence G. Heller. A 1 mV MOS comparator. *IEEE Journal of Solid-State Circuits*, 13(3):294–297, June 1978.
- [62] Walter Kester. Measure flash-ADC performance for trouble-free operation. *EDN*, pages 103–114, February 1, 1990.
- [63] Walter Kester. Flash ADCs provide the basis for high-speed conversion. *EDN*, pages 101–110, January 4, 1990.
- [64] Analog Devices Inc. AD678 12-bit 200 kSPS complete sampling ADC. *Datasheet*, 2000.
- [65] Brendan Coleman, Pat Meehan, John Reidy, and Pat Weeks. Coherent sampling helps when specifying DSP A/D converters. *EDN*, pages 145–152, October 15, 1987.
- [66] Fredric J. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66(1):51–83, January 1978.
- [67] Walter Kester. DSP test techniques keep flash ADCs in check. *EDN*, pages 133–142, January 18, 1990.
- [68] Walter Kester. Test video A/D converters under dynamic conditions. *EDN*, pages 103–112, August 18, 1982.
- [69] Joey Doernberg, Hae-Seung Lee, and David A. Hodges. Full speed testing of A/D converters. *IEEE Journal of Solid State Circuits*, 19(6):820–827, December 1984.
- [70] Brian J. Frohring, Bruce E. Peetz, Mark A. Unkrich, and Steven C. Bird. Waveform recorder design for dynamic performance. *Hewlett-Packard Journal*, 39(1):39–48, February 1988.
- [71] Aldert van der Ziel. *Noise in Solid State Devices and Circuits*. Wiley, New York, 1986.
- [72] IRE Symbols Committee et al. IEEE standard letter symbols for semiconductor devices. *IEEE Transactions on Electron Devices*, 11(8):392–397, August 1964.
- [73] Paul R. Gray and Robert G. Meyer. *Analysis and Design of Analog Integrated Circuits*. Wiley, New York, 1993.
- [74] James K. Roberge. *Operational Amplifiers: Theory and Practice*. Wiley, New York, 1975.
- [75] C. D. Motchenbacher and F. C. Fitchen. *Low-Noise Electronic Design*. Wiley, New York, 1973.
- [76] William Strunk, Jr. and E. B. White. *The Elements of Style*. Macmillan, New York, 1959.

- [77] J. B. Johnson. The schottky effect in low frequency circuits. *Physical Review*, 26(1):71–85, July 1925.
- [78] M. J. Buckingham. *Noise in Electronic Devices and Systems*. Ellis Horwood Limited, Chichester, England, 1983.
- [79] A. L. McWhorter. 1/f noise and germanium surface properties. In R. H. Kingston, editor, *Semiconductor Surface Physics*, pages 207–228. University of Pennsylvania Press, Philadelphia, 1957.
- [80] S. Christensson, I. Lundström, and C. Svensson. Low frequency noise in MOS transistors — I. theory. *Solid-State Electronics*, 11:797–812, 1968.
- [81] S. Christensson and I. Lundström. Low frequency noise in MOS transistors — II. experiments. *Solid-State Electronics*, 11:813–820, 1968.
- [82] Makunda B. Das and James M. Moore. Measurements and interpretation of low-frequency noise in FET's. *IEEE Transactions on Electron Devices*, 21(4):247–257, April 1974.
- [83] G. Reimbold. Modified 1/f trapping noise theory and experiments in MOS transistors biased from weak to strong inversion — influence of interface states. *IEEE Transactions on Electron Devices*, 31(9):1190–1198, September 1984.
- [84] Gérard Ghibaudo. A simple derivation of Reimbold's drain current spectrum formula for flicker noise in MOSFETs. *Solid-State Electronics*, 30(10):1037–1038, October 1987.
- [85] F. N. Hooge. 1/f noise is no surface effect. *Physics Letters*, 29A(3):139–140, April 1969.
- [86] F. N. Hooge. Discussion of recent experiments on 1/f noise. *Physica*, 60:130–144, 1976.
- [87] F. N. Hooge and L. K. J. Vandamme. Lattice scattering causes 1/f noise. *Physics Letters*, 66A(4):315–316, May 1978.
- [88] F. N. Hooge. 1/f noise. *Physica*, 83B(1):14–23, May 1976.
- [89] R. P. Jindal and A. van der Ziel. Phonon fluctuation model for flicker noise in elemental semiconductors. *Journal of Applied Physics*, 52(4):2884–2888, April 1981.
- [90] L. K. J. Vandammenad H. M. M. de Werd. 1/f noise model for MOSTs biased in the nonohmic region. *Solid-State Electronics*, 23(4):325–329, April 1980.
- [91] R. Jayaraman and C. Sodini. A 1/f noise technique to extract the oxide trap density near the conduction band edge of silicon. *IEEE Transactions on Electron Devices*, 36(9):1773–1782, September 1989.
- [92] Kwok K. Hung, Ping K. Ko, Chenming Hu, and Yiu C. Cheng. A unified model for the flicker noise in metal-oxide-semiconductor field-effect transistors. *IEEE Transactions on Electron Devices*, 37(3):654–665, March 1990.
- [93] Kwok K. Hung, Ping K. Ko, Chenming Hu, and Yiu C. Cheng. A physics-based MOSFET noise model for circuit simulators. *IEEE Transactions on Electron Devices*, 37(5):1323–1333, May 1990.

- [94] F. N. Hooge. 1/f noise sources. *IEEE Transactions on Electron Devices*, 41(11):1926–1935, November 1994.
- [95] L. K. J. Vandamme, Xiaosong Li, and Dominique Rigaud. 1/f noise in MOS devices, mobility or number fluctuations? *IEEE Transactions on Electron Devices*, 41(11):1936–1945, November 1994.
- [96] Jimmin Chang, A. A. Abidi, and C. R. Viswanathan. Flicker noise in CMOS transistors from subthreshold to strong inversion at various temperatures. *IEEE Transactions on Electron Devices*, 41(11):1965–1971, November 1994.
- [97] John H. Scofield, Nick Borland, and Daniel M. Fleetwood. Reconciliation of different gate-voltage dependencies of 1/f noise in n-MOS and p-MOS transistors. *IEEE Transactions on Electron Devices*, 41(11):1946–1952, November 1994.
- [98] Rahul Sarpeshkar, Tobias Delbrück, and Carver A. Mead. White noise in MOS transistors and resistors. *IEEE Circuit and Devices Magazine*, 9(6):23–29, November 1993.
- [99] B. Jeffrey Gross and Charles G. Sodini. 1/f noise in MOSFETs with ultrathin gate dielectrics. In *IEDM Technical Digest*, pages 881–884, 1992.
- [100] Jerris F. Johnson. Hybrid infrared focal plane signal and noise model. *IEEE Transactions on Electron Devices*, 46(1):96–108, January 1999.
- [101] Jian Chen, Peng Fang, Ping Keung Ko, Chenming Hu, Ray Solomon, Tung-Yi Chan, and Charles G. Sodini. Noise overshoot at drain current kink in SOI MOSFET. In *1990 IEEE SOS/SOI Technology Conference Proceedings*, pages 40–41, 1990.
- [102] P. Gentil and S. Chausse. Low-frequency noise measurements on silicon-on-sapphire (SOS) MOS transistors. *Solid-State Electronics*, 20(11):935–940, November 1977.
- [103] F. Faccio, P. Aspell, P. Jarron, and E. H. Heijne. An additional contribution in the noise spectrum of SOI MOSFETs. In *Proceedings of the 1996 IEEE International SOI Conference*, pages 9–10, October 1996.
- [104] Jeffery A. Babcock, Dieter K. Schroder, and Ying-Che Tseng. Low-frequency noise in near-fully-depleted TFOSI MOSFET's. *IEEE Electron Device Letters*, 19(2):40–43, February 1998.
- [105] Ying-Che Tseng, Wen-Ling M. Huang, Pamela J. Welch, Jenny M. Ford, and Jason C. S. Woo. Empirical correlation between AC kink and low-frequency noise overshoot in SOI MOSFET's. *IEEE Electron Device Letters*, 19(5):157–159, May 1998.
- [106] Wei Jin, Phillip C. H. Chan, Samuel K. H. Fung, and P. K. Ko. A physically-based low-frequency noise model for NFD SOI MOSFET's. In *Proceedings of the 1998 IEEE International SOI Conference*, pages 23–24, October 1998.
- [107] M. Valenza, C. Barros, D. Rigaud, E. Simoen, and C. Claeys. Excess noise behaviour in short n-channel SOI MOSFET's. In *Proceedings of the 14th International Symposium on Noise in Physical Systems and 1/f Fluctuations*, pages 265–269, 1997.
- [108] E. Simoen and C. Claeys. Study of the kink-related excess low frequency noise in SOI NMOS-FET's at room temperature and 77K. In *Proceedings of the Fifth International Symposium on Silicon-on-Insulator Technology and Devices*, pages 127–134. Electrochemical Society, 1992.

- [109] E. Simoen, Ulf Magnusson, Antonio L. P. Rotondaro, and Cor Claeys. The kink-related excess low-frequency noise in silicon-on-insulator MOST's. *IEEE Transactions on Electron Devices*, 41(3):330–339, March 1994.
- [110] E. Simoen and C. Claeys. Correlation between the low-frequency noise spectral density and the static device parameters of silicon-on-insulator MOSFET's. *IEEE Electron Device Letters*, 42(8):1467–1472, August 1995.
- [111] E. Simoen and C. Claeys. Generation-recombination noise sources in silicon-on-insulator MOSFETs. In *Proceedings of the Seventh International Symposium on Silicon-on-Insulator Technology and Devices*, pages 255–262. Electrochemical Society, 1996.
- [112] E. Simoen and C. Claeys. The low-frequency noise behaviour of silicon-on-insulator technologies. *Solid-State Electronics*, 39(7):949–960, July 1996.
- [113] Zhong Yaun Chang and Willy M. C. Sansen. *Low-Noise Wide-Band Amplifiers in Bipolar and CMOS Technologies*. Kluwer Academic Publishers, Boston, 1991.
- [114] Philips Semiconductors. TDA8776A 10-bit, 1000 Msps digital-to-analog converter (DAC). *Datasheet*, 1996.

4935 - 55