HIGH SPEED DATA CONVERTER CIRCUITS IN SI-GE

by

Dirk J. Robinson

A dissertation submitted in partial fulfillment of

the requirements for the degree of

DOCTOR OF PHILOSOPHY

WASHINGTON STATE UNIVERSITY

School of Electrical Engineering and Computer Science

December 2008

To the Faculty of Washington State University:

The members of the committee appointed to examine the dissertation of DIRK J. ROBINSON find it satisfactory and recommend that it be accepted.

_____

Chair

_____

_____

ACKNOWLEDGMENTS

I would first like to thank my wife Michelle for her steady support and encouragement throughout my extended graduate education. My new daughter Katelyn has also helped me to finish on time by showing me pure determination in learning to crawl and walk.

I also appreciate the time my committee members have taken to ensure my success in graduate school. I enjoyed working with Dr. La Rue and his style of teaching by working closely with his students. Finally, I would like to thank the Center for the Design of Analog and Digital Integrated Circuit (CDADIC), for providing the majority of the funding for this work.

HIGH SPEED DATA CONVERTER CIRCUITS IN SI-GE

Abstract

by Dirk J. Robinson, Ph.D.

Washington State University

December 2008

Chair: George S. La Rue

Analog-to-Digital Converters (ADCs) with high speed and high accuracy play a critical role in electronics testing, scientific instrumentation, and digital communication systems. The performance of ADCs in this performance corner are limited by the input Track-and-Hold Amplifier (THA) that is used. Consequently, this input THA must have excellent thermal noise and jitter performance.

Past THA designs were optimized for low signal distortion at the expense of voltage headroom, power consumption, thermal noise, and bandwidth. Current technology scaling makes it desirable to allow increased signal distortion, which will be compensated for in the digital domain. With the lower voltage headroom limitations of modern processes, digital distortion compensation will become essential in order to achieve ever higher sampling rates at high-accuracy.

The purpose of this research was to design and fabricate an ADC using a THA based on this principle. Digital distortion compensation methods unique to the architecture of this ADC were also investigated. This document describes the design and performance of the ADC and evaluates the compensation algorithms.

TABLE OF CONTENTS

Page

Appendix

Page

LIST OF TABLES

LIST OF FIGURES

Figure Page

NOMENCLATURE

ADC Analog-to-Digital Converter.

DAC Digital-to-Analog Converter.

BDMS-THA Bipolar-Driven MOS-Switch THA. A THA design using passive CMOS switches whose gates are driven by NPN bipolar transistors. Developed during this work.

BDMS-ADC A pipelined ADC designed incorporating the BDMS-THA in the first pipeline stage. Developed during this work.

ENOB Effective-Number-Of-Bits. The number of bits of an ideal quantizer that would result in a quantization noise power equal to the power in a given ADC's distortion and noise.

GUI A Graphical User Interface, is a part of a software program that allows a user to interact with the software in a graphical, event-driven manner.

LSB Least Significant Bit. The voltage step equivalent to a unit change in the digital word, for either an ADC or DAC.

LUT Look-Up Table. An array of values used for compenstation that are indexed by part of the binary result of an ADC.

PCB Printed Circuit Board.

SINAD SIgnal-to-Noise-And-Distortion ratio. The ratio of power in the noise and distortion to that of the signal.

SFDR Spurious-Free Dynamic Range. When a pure sinusoid is input to an ADC or DAC, this is the ratio of output power at the input frequency to the power level at the next highest frequency point (spur).

SNR Signal-to-Noise ratio. The ratio of power in the thermal and quantization noise to the power in the signal.

THA Track-and-Hold Amplifier. Holds an analog signal at a constant level for a certian time interval so that is may be processed.

Chapter 1

Introduction

Analog-to-Digital Converters (ADCs) with high speed and high accuracy play a critical role in electronics testing, scientific instrumentation, and digital communication systems. There is a constant push to improve the trade-off between noise, distortion, bandwidth and power of these converters. This work investigates methods to improve this trade-off and presents a new high-performance ADC that takes advantage of these methods.

1.1 Motivation

In the high-speed ADC frontier, an ADC's performance is generally limited by the performance of the input THA. Even single-stage flash ADC architectures require an input THA in order to achieve modest resolution. The input THA reduces sampling time uncertainty, which becomes a limiting factor for signals approaching 1 GHz. By holding the signal constant for a length of time, the THA also relieves the ADC system of gain roll-off and hysteresis effects including non-linear settling.

We investigated two THA architectures that are typically used at high speeds, the Diode-Bridge THA[1] and the Emitter-Switch THA[2]. We found that these conventional architectures required large voltage headroom due to device stacking. This made them power-inefficient and difficult to design using modern process technologies with their reduced device breakdown voltages.

1.2 Objectives and Approach

By allowing increased distortion in hold mode, we were able to develop a THA architecture that operated with much lower voltage headroom. In the ADC performance trade-off between noise, distortion, bandwidth and power, distortion is the term with potential for improvement with post-processing. With this in mind, the THA design will prioritize noise, bandwidth, power and headroom over distortion.

We placed our low-headroom THA design as the input THA to a four-stage pipelined ADC. The first stage of this ADC has a similar performance trade-off compared with the THA. The DAC in this stage must have similar noise performance to the THA. The residual amplifier also must match the noise performance of the THA, but at an increased gain-bandwidth-product.

To meet these goals, the ADC also was designed to prioritize the uncorrectable performance attributes over distortion. As such, the THA design centers on the use of open-loop amplifiers and implementing distortion correction in the digital domain.

1.2.1 Designing With Open-Loop Amplifiers

In order to optimize the design of the ADC for noise, bandwidth and power, we implemented the residual amplifiers and buffers as open-loop amplifiers. Although they have increased distortion compared to closed-loop amplifiers, open-loop amplifiers have significantly better noise performance[3]. The high gain-bandwidth-product requirement of these amplifiers was also a factor in choosing to use open-loop designs.

Although open-loop amplifiers were used in all four stages of the design, they contribute significant distortion only in the first stage. Still, the distortion of the system is made more difficult to model, as it now depends on the first stage binary output as well as the THA input history.

1.2.2 Moving Compensation to Digital Domain

Past designs, including our own[4], have used analog feed-forward techniques to compensate for distortion. However, these methods tend to increase power consumption and voltage headroom, and often add parasitics to critical signal paths.

It is well known that technology scaling can not be expected to yield the same level of performance enhancements to analog circuitry as it does for digital[5, 6, 7, 8]. Consequently, the impetus for mixed-signal designs is to move as much as possible from the analog domain to the digital. By applying this concept to distortion compensation, better overall performance can be achieved. This also makes it possible to design ADCs that better match the headroom requirements of modern processes.

Chapter 2

Background

High-speed THAs are typically characterized separately from ADCs using a master-slave arrangement [2]. The second THA in the master-slave arrangement acts like a frequency down mixer, making it easy to see the hold-mode distortion characteristics of the first THA. However, thermal noise and sampling time jitter cannot be directly measured by the master-slave arrangement. Since distortion can be corrected to some extent using digital post-processing, THA performance ultimately limited by random noise added to each sample, requiring a different test arrangement.

With the above considerations in mind, the best way to evaluate the performance of a THA is to sample its output directly with an ADC. In principle, the ADC could operate in a sub-sampling mode, at a much lower clock rate than the THA. By operating the ADC at a low frequency, the parasitics of the test setup become negligible, so the ADC could be placed off-chip. However, in this mode, the accumulated time jitter between the clock and input sources becomes a problem. To prevent accumulated jitter, we designed an ADC to operate at the same speed as the THA and placed it on the same die.

The primary function of THAs is as input samplers to high-performance ADCs. By designing the full ADC, we can better understand how the performance of the THA relates to the overall system performance.

2.1 THA Comparison

The Emitter-Switched THA and Diode-Bridge THA are commonly used as high-speed samplers. They can be constructed using only NPN transistors and resistors, giving them exceptionally fast hold transitions. Consequently, they have excellent jitter performance. However, they each require large voltage headroom. We developed the "Bipolar-Driven MOS-Switch THA" (BDMS-THA) in order to address the voltage headroom issue and achieve a better noise-power-bandwidth trade-off.

Project	Technology	Sample Rate	Power	Distortion	Noise
T. Baum [2]	Silicon Bipolar (Emitter-Switch)	1 GSps	490 mW	10 bits	(not reported)
R. Yu [1]	GaAs (Diode Bridge)	1 GSps	800 mW	11 bits	(not reported)
WSU	SiGe BiCMOS	1 GSps	400 mW	9 bits (sim)	>11 bits (meas)

Table 2.1 THA Comparison

Table 2.1 shows a comparison of the BDMS-THA developed at WSU with competitive bipolar-only THAs. The power reported does not include the clock buffer for each of the THAs. For a fair comparison of the three THA architectures, we designed and fabricated each in the Jazz Semiconductor 0.18 μm BiCMOS process. Table 2.2 shows the simulated performance metrics for these THAs. Fig. 2.1 and Fig. 2.2 show the simulated THD and gain roll-off performance of the three THAs designs.

	BDMS	Switched-Emitter	Diode Bridge
Signal Swing, Differential	1.8 V	3.2 V	2 V
Hold Noise (RMS)	70 μV, 12.6 bits	450 μV, 10.8 bits	366 μV, 10.4 bits
Gain Roll-Off	0.55 dB	0.20 dB	0.22 dB
Power Consumption	0.4 W	1.1 W	1.1 W
Power Supplies	+3.6 V, +1.8 V	+4.8 V, -3.0 V	+4V, -4V
Feedthrough at 500 MHz	-52.4 dB	-98.5 dB	-69.5 dB

Table 2.2 Simulated THA Performance Comparison for Our Designs

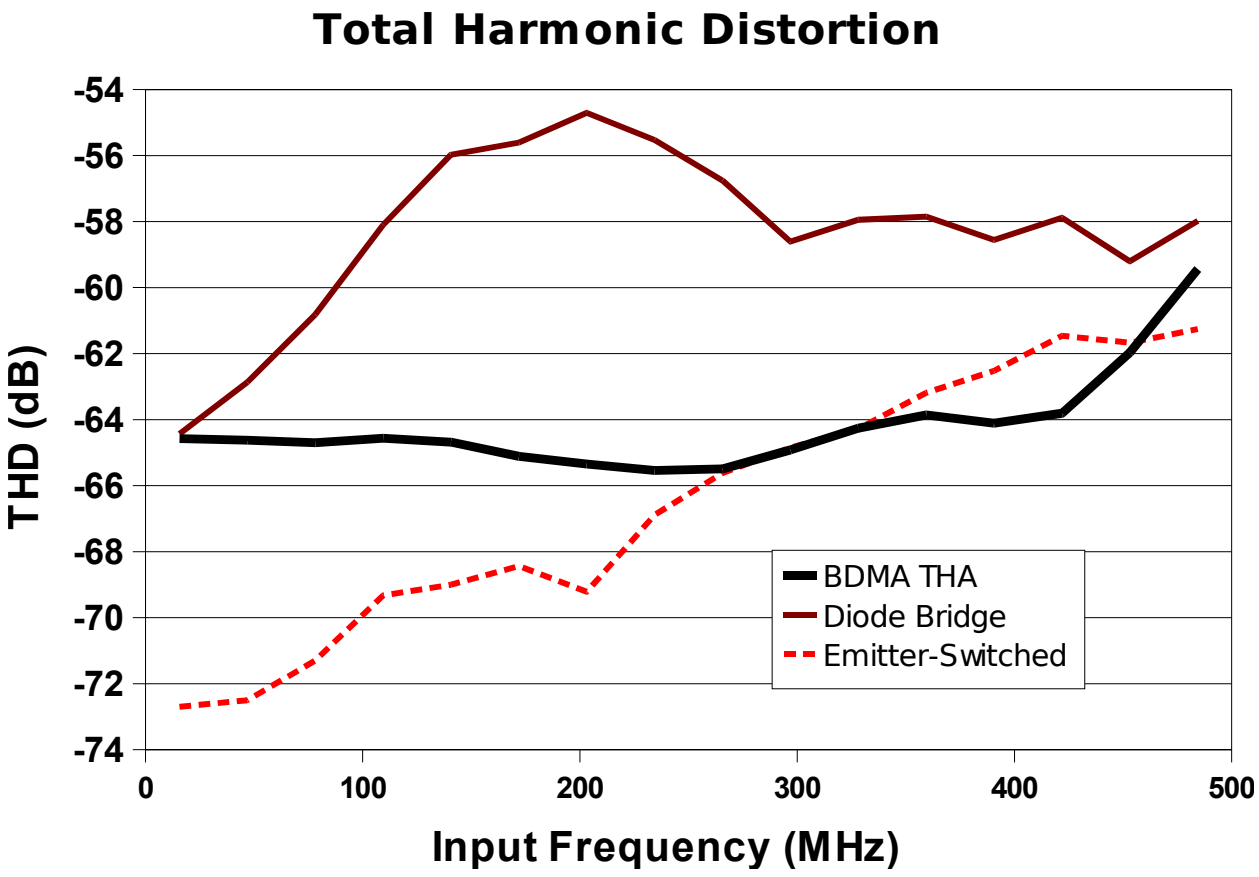

Figure 2.1 Simulated THA Distortion

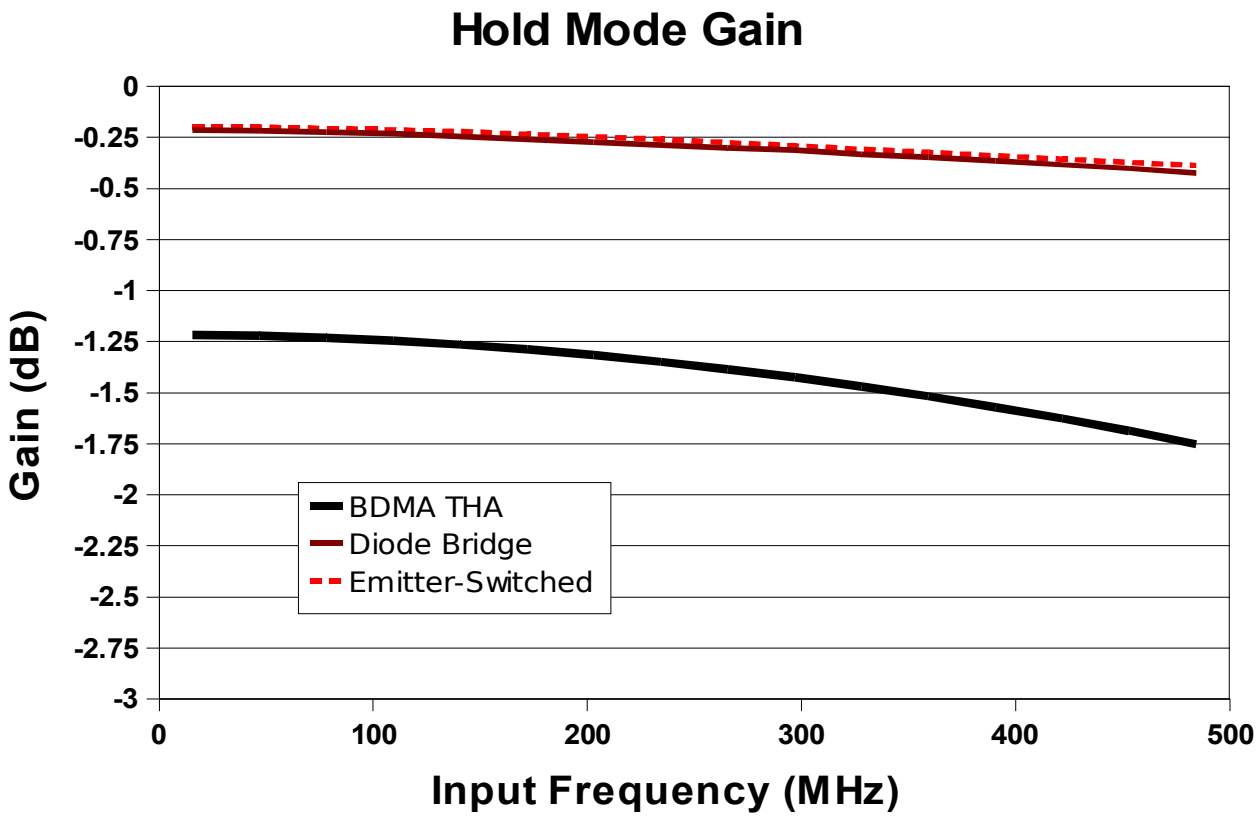

Figure 2.2 Simulated THA Gain Rolloff

Unfortunately, the fabricated version of the Switched-Emitter THA could not retain the signal in hold mode. One concern is that the transistors used have a collector-emitter breakdown voltage of 3.6 V, which is significantly lower than the power supply voltage range. However, a simulation of ramping up the power supplies from zero shows that all transistors are kept in the safe operating region. We suspect that the problem is related to excessive collector-to-substrate biasing.

It should be noted that analog distortion compensation circuitry was added to the Emitter-Switch THA, which forced a higher voltage headroom requirement. Whereas, the BDMS-THA relies on digital distortion compensation. Thus, the power/noise trade-off of the Emitter-Switch THA might compare more favorably with the BDMS-THA if it were designed with digital compensation in mind.

The Diode Bridge THA also requires a high power supply range because of device stacking. In simulation, this THA performs with inherently low-distortion, without requiring much compensation circuitry. However, the fabricated version of this THA showed noise and distortion levels of only about 7 effective bits. The THA was tested using a custom-designed probe card, which may have contributed to the observed distortion level. This poor result was another motivating factor to develop a THA and ADC together on the same die to facilitate testing.

2.2 ADC Performance

Besides power and sampling rate, the three most important metrics of a high-speed ADC are thermal noise, sampling time jitter and distortion. Noise and jitter can only be improved in post-processing at the expense of signal bandwidth. However, if the distortion characteristics of the ADC can be sufficiently modeled, the distortion may be corrected by digital post-processing.

2.2.1 Quantization Noise

The quantization noise limit for an ideal N-bit ADC is [9]:

$$V_{Q(rms)} = \frac{V_{LSB}}{\sqrt{12}} \tag{2.1}$$

Thermal voltage noise adds in quadrature to quantization noise. Consider an ADC with equal thermal and quantization RMS noise voltages. If the ADC has a thermal noise level equivalent to N-bits, quantization will reduce the SNR by $\sqrt{1^2 + 1^2}$, resulting in N-0.5 effective bits.

This design, like many pipelined ADCs, was designed to have one extra bit of resolution beyond the thermal noise limit. It requires little additional power to increase the quantization resolution for pipelined architectures, if thermal noise performance remains constant. In this case, if the ADC has a thermal noise level equivalent to N-bits, quantization will reduce the SNR by $\sqrt{1^2 + \frac{1}{2}^2}$, resulting in N-0.16 effective bits. This allows us to neglect quantization noise, and also prevents concern about round-off error in the distortion compensation algorithm.

2.2.2 Thermal Noise

Thermal noise tends to limit the resolution of all ADCs up to the gigahertz range [10]. Using Eq. 2.1 and $V_{LSB} = \frac{V_{RANGE}}{2^N}$, in order to achieve N_t-bit performance with respect to thermal noise, we must limit the thermal noise to:

$$V_{t(rms)} = \frac{V_{RANGE}}{2^{N_t}\sqrt{12}} \tag{2.2}$$

We can also express the thermal noise in terms of an equivalent input resistance. The thermal noise generated by a resistor is:

$$V_{t(rms)} = \sqrt{4kTRf} = 4\,\mathrm{nV}\sqrt{\frac{R}{1\,\mathrm{k\Omega}}\frac{f}{1\,\mathrm{Hz}}}\Bigg|_{T=300\,\mathrm{k}} \tag{2.3}$$

where k, T, R, f are the Boltzmann constant, temperature, resistance and noise bandwidth, respectively. On the rightmost side of Eq. 2.3 is a useful approximation for room temperature applications.

For example, consider an ADC with 12 effective bits with respect to thermal noise, an input range of 1 V, and an input noise bandwidth of 1 GHz. From Eq. 2.2, this ADC has an input referred voltage noise of 70 μV. Using Eq. 2.3, we obtain an equivalent input noise resistance of only 310 Ω.

In a pipelined ADC, the first stage THA, DAC and differencing amplifier are all major contributors to thermal noise. In order to achieve noise performance similar to the above example, the use of MOS devices in active mode is not practical for any of these components due to their excessive noise and input capacitance. Thus, the only MOS components used in the signal path in our ADC design are used as passive switches, while the remainder of the design uses bipolar NPN type amplifiers.

2.2.3 Sampling Time Jitter

Sampling time jitter, or uncertainty, begins to limit THA and ADC performance at high frequencies. The SNR due to an RMS time jitter of $\epsilon_{\text{rms}}$is given by $\text{SNR} = -20\log_{10}(2\pi f \epsilon_{\text{rms}})$, where f is the input frequency [11].

In addition to jitter from the external clock source, noise sources in the ADC's clock buffer contribute to jitter. ADCs in the gigahertz range generally use a sinusoidal clock input so the input noise of the first clock buffer stage must also be low.

For example, in an 11-bit ADC, the quantization limit for SNR is 68 dB. If we assume a maximum input frequency of 500 MHz, we must limit the jitter to 126.7 fs in order to achieve 11-bits effective bits with respect to jitter.

2.2.4 Distortion

The Signal-to-Noise-and-Distortion (SINAD) ratio is the most comprehensive metric of ADC performance. It is the ratio of power in the desired signal to power after the signal has been subtracted. As such, it includes components from quantization, thermal noise, jitter as well as distortion.

For ADCs used in communication, the Spurious-Free Dynamic Range (SFDR) is also important. It represents the ratio of power at the signal frequency to the power level at the next highest peak (spur) in frequency space. The spur is usually a second or third harmonic of the input.

With a model of the distortion characteristics of an ADC, is is possible to improve both the SINAD and SFDR.

Chapter 3

ADC Design

The ADC was designed as a four-stage pipeline, as shown in Fig. 3.1. To save power, the second and third stages were scaled to use approximately 25% of the power of the first stage.

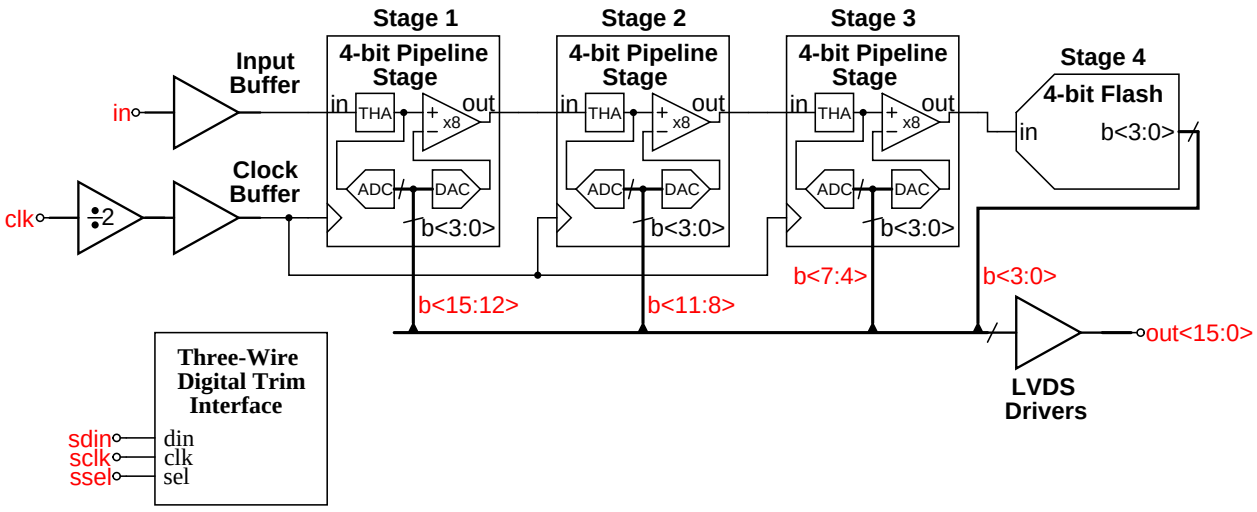

Figure 3.1 Pipelined ADC

Error correction is achieved by having approximately one bit of overlap per stage[12]. This lessens the accuracy requirement of the flash ADC in each stage to an accuracy of about 5 bits. Since this allows less than full resolution settling of the ADC input, the conversion can occur immediately upon entering hold mode. This helps eliminate an additional THA per pipeline stage, saving significant power and added noise.

3.1 Pipeline Stage in Detail

Fig. 3.2 shows a simplified schematic of the first three pipeline stages. Stages 2 and 3 use a very simple THA like the one shown, although the design is fully differential.

After the THA, the analog signal splits into two paths: the analog path on top and the digital path below. In the digital path, the signal is quantized by a 4-bit flash ADC with minimal delay after the THA enters hold-mode. The 15 comparator outputs of the ADC feed directly into 15 equally sized current sources, which make up the 4-bit DAC. At the end of the digital path is a quantized version of the input, which is subtracted from the analog path. The result, or residual, is then fed into the next pipeline stage.

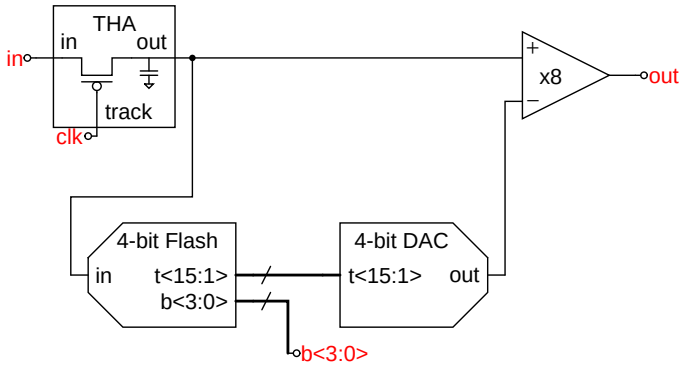

Figure 3.2 Typical Pipeline Stage

There are four such stages, although the final stage only consists of a 4-bit Flash. In the ideal case, where the gain of the residual amplifier is fixed at 8, there is exactly one bit of overlap between stages. Thus, we can apply the error correction and obtain the ADC result by a simple

binary summation:

$$\begin{array}{cccccccccccccc} & b_{15} & b_{14} & b_{13} & b_{12} & & & & & & & & & \\ + & & & & b_{11} & b_{10} & b_9 & b_8 & & & & & & \\ + & & & & & & & b_7 & b_6 & b_5 & b_4 & & & \\ + & & & & & & & & & & b_3 & b_2 & b_1 & b_0 \\ \hline & d_{12} & d_{11} & d_{10} & d_9 & d_8 & d_7 & d_6 & d_5 & d_4 & d_3 & d_2 & d_1 & d_0 \end{array} \tag{3.1}$$

Where $b_{15:12}, b_{11:8}, b_{7:4}, b_{3:0}$ are the outputs from stages 1, 2, 3, 4 and $d_{12:0}$ is the ADC result. In practice, the inter-stage gains are subject to process variation. Obtaining the ADC result for this case will be addressed in the following chapter.

As shown, the THA, ADC, DAC and residue amplifier must all settle in a single clock cycle. Many pipeline ADC designs use an additional THA to create an additional half-cycle delay in the analog path[12, 13]. We elected to omit the extra THA to reduce the noise of the stage. This is made possible by our fast ADC-DAC combination, which is in turn made possible by the use of error correction in each stage.

3.1.1 The Bipolar Driven MOS-Switch THA

The top level architecture of the BDMS-THA is shown in Fig. 3.3. The design is partitioned into a gate driving circuit constructed with high-speed NPN bipolar transistors, and a switching network of PMOS devices. The first stage THA must have low jitter, making it moderately complex and power hungry. Subsequent stages deal with a stationary level at reduced resolution, so a simple single-transistor switch with binary gate levels suffices.

3.1.1.1 MOS Switches

Fig. 3.4 shows the heart of the BDMS-THA. Transistors M1 and M3 have an on-resistance of around $50\,\Omega$, as do their differential counterparts, M2 and M4. When in track mode, the signal

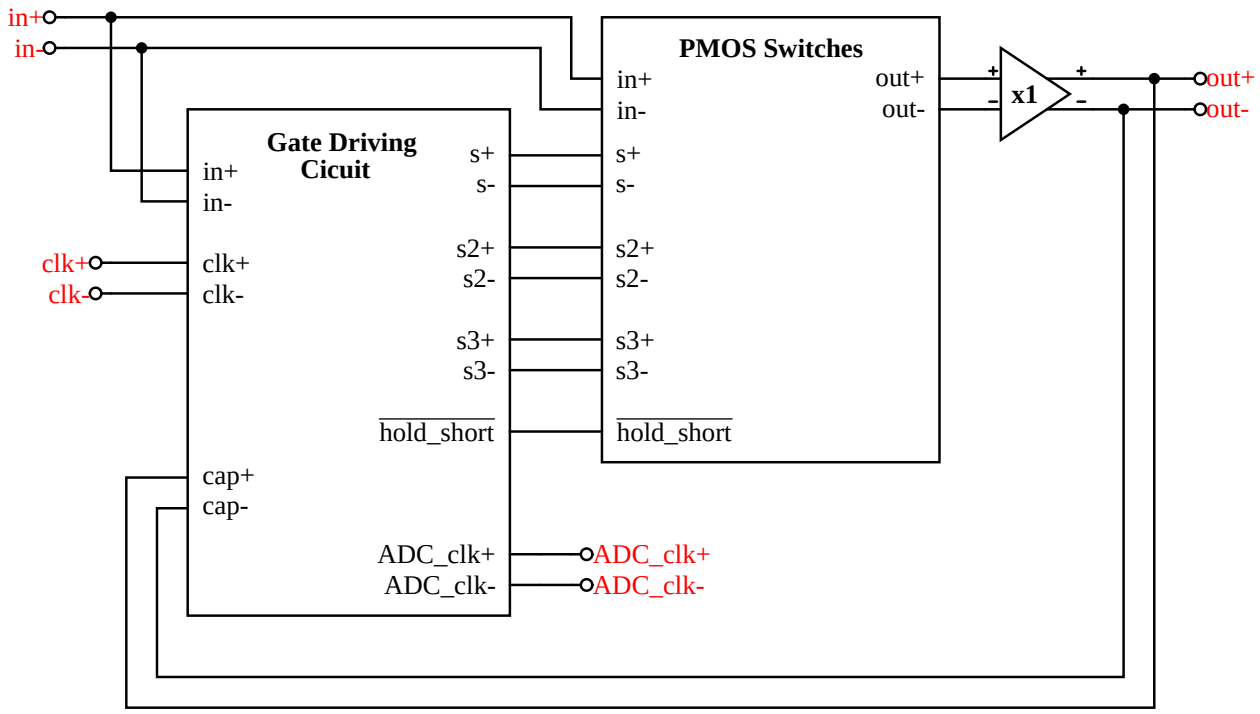

Figure 3.3 First Stage Track and Hold

at the gates of these transistors tracks the input. This maintains a nearly constant on-resistance, keeping the distortion manageable.

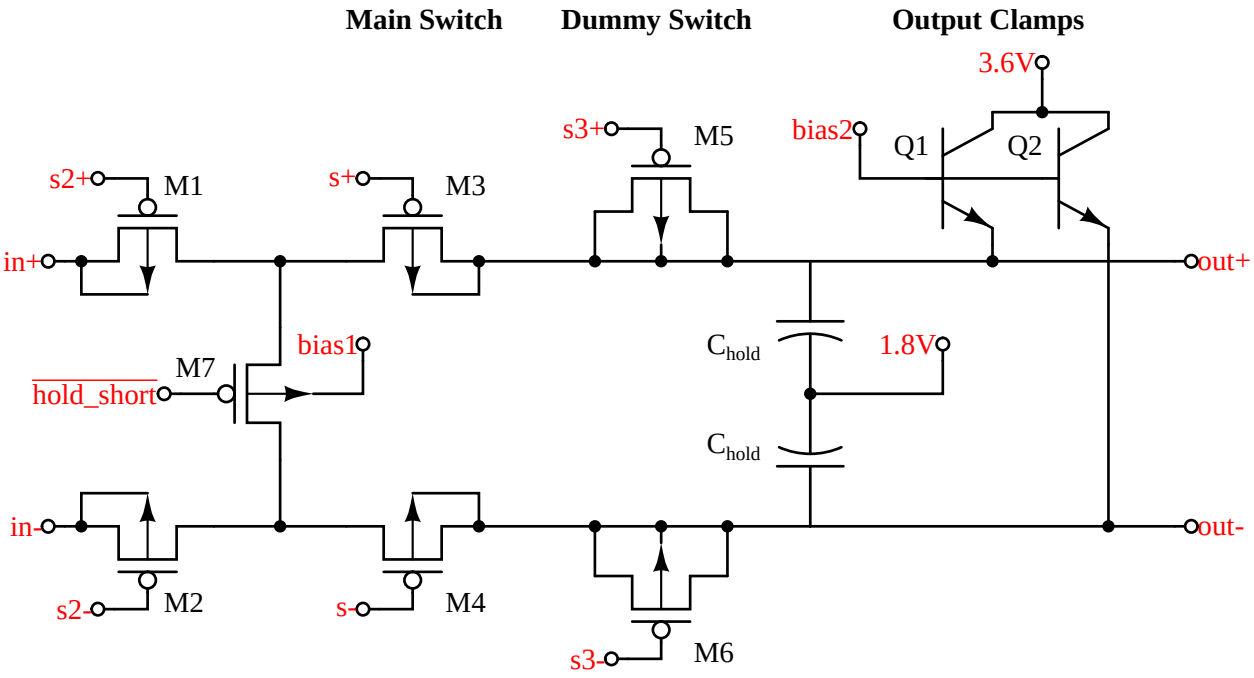

Figure 3.4 MOS Switches in the Main THA

The on-resistance of each path forms a low pass filter with the hold capacitor. We use this to limit the noise bandwidth to around 1 GHz. Note that the noise bandwidth of a first order filter is about 1.6 times greater than the filters 3 dB frequency[3]. Unfortunately, this also limits the settling time in track mode, creating a gain roll-off with increasing input frequency.

When entering hold mode, M3 and M4 are the first to be turned off when their gates are pulled high. These transistors are considered the main switch, as their transition time determines the jitter of the whole ADC. Next, M1 and M2 are switched off, and M7 is turned on to short the intermediate signal nodes together. This greatly reduces feed-through in hold mode. At the same time, the gates of M5 and M6 are pulled high in order to offset most of the charge injection.

Although the BDMS-THA operates at a power supply range of 3.6 V, it uses PMOS transistors with 1.8 V breakdown for improved performance. Output clamps were added to ensure the signal would stay in range to protect these transistors.

3.1.1.2 Gate Driving Circuit

The basic building block of the gate driving circuit is a maximum value function shown in Fig. 3.5. It enables the gates in the THA to switch from signal-following to holding with the speed of an NPN emitter follower. Also, note that with PMOS switches, entering hold mode is a positive transition and consequently not current limited. Although PMOS typically requires about double the gate capacitance of NMOS to achieve a given on-resistance, the transition speed is primarily determined by the NPN devices.

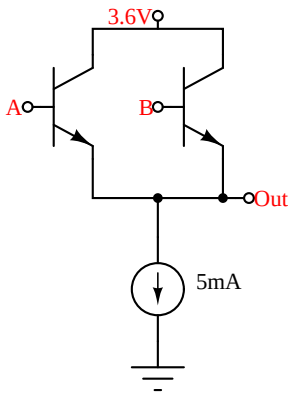

Figure 3.5 Maximum Value Circuit (ECL Logic OR)

The entire gate driving circuit is shown in Fig. 3.6. With an input of 1 GHz, the clock buffer and gate driving circuit have a combined timing jitter of 64 fs at the s+/s- ouput pair, in simulation. The total jitter at the hold capacitors was simulated at 99 fs.

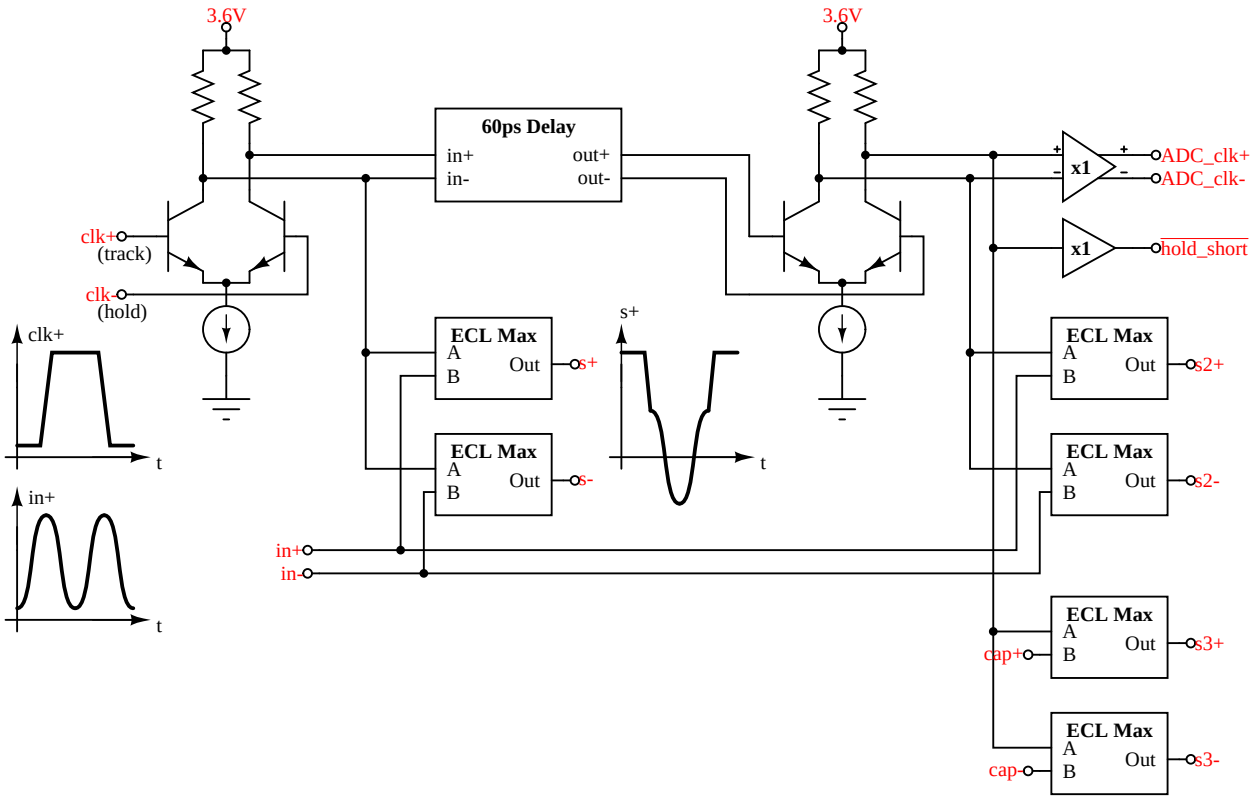

Figure 3.6 Gate Driving Circuit

3.1.2 Differential 4-bit Flash

All four pipeline stages use the 4-bit flash ADC shown in Fig. 3.7. Thanks to using a single bit of overlap for error correction between the stages, the ADC need only have an accuracy of around 5-bits. Thus, the ADC need not consume much power, and can be efficiently used without scaling or other modification for all stages.

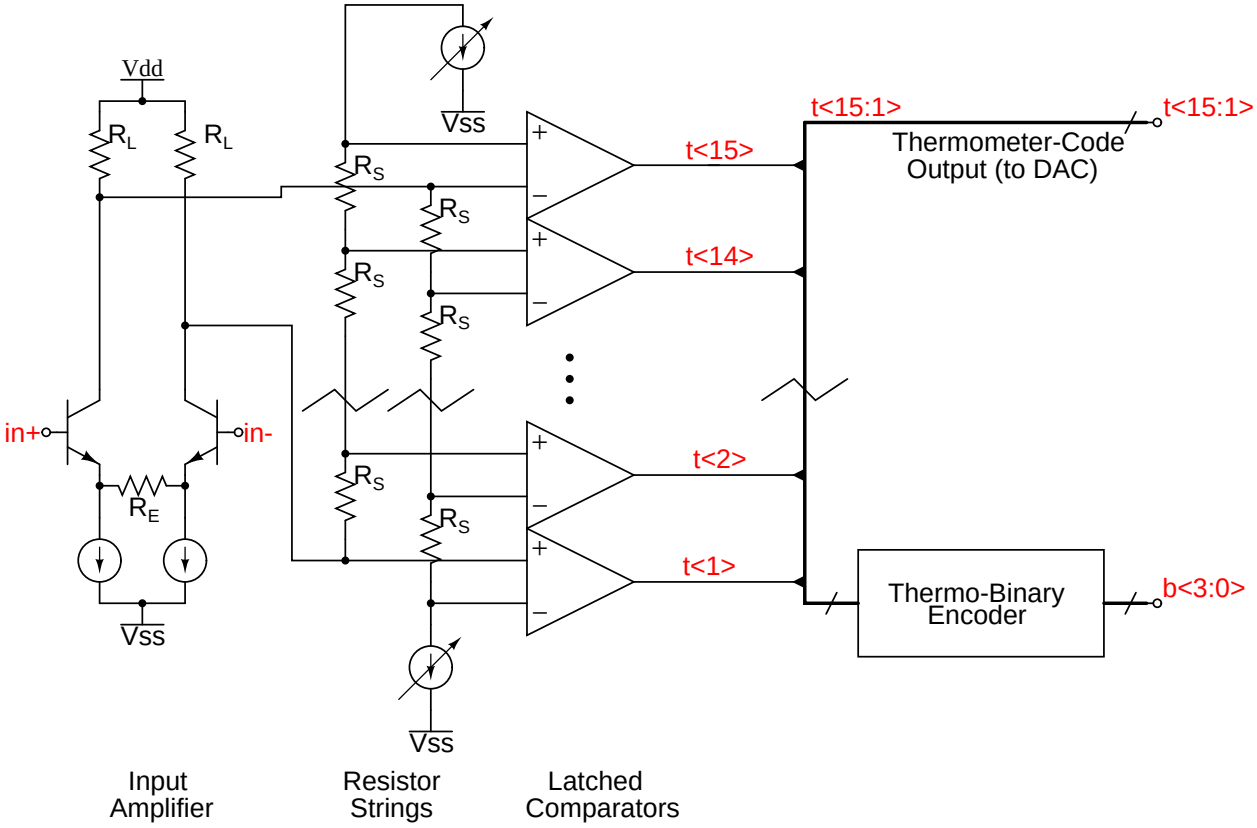

Figure 3.7 Differential 4-bit Flash

The ADC architecture uses flash architecture based on a pair of resistor strings. The concept seems an obvious extension of a single-ended resistor string flash, but we have not yet encountered it in the literature. The dual resistor string flash is fully differential all the way to the thermometer outputs, which feed a fully differential DAC. Another advantage of the design is that no reference

voltages are needed. The range and offset of the flash are controlled by two digitally trimmed current sources attached to the resistor strings.

In stages 1 to 3, the ADC range and offset must be trimmed to match the range of the corresponding DAC, so that the stage residue does not exceed the error correction range of the following stage. The ADC range of stage four can be more flexibly set. By decreasing the range of stage four, we can increase the effective number of quantization bits. We took advantage of this fact to compensate for lower than expected residue amplifier gains in the fabricated version of the ADC.

3.1.3 Differential Differencing Amplifier

The differential differencing amplifier (DDA) shown in Fig. 3.8 was used as the residue amplifier in the first three pipeline stages. The amplifier is unique in that it subtracts a differential voltage input from the THA and a differential current input from the DAC. Since the first stage of the DDA is biased by the current output from the DAC, this design saves considerable power. Speed and noise performance are also improved by shortening the DAC signal path.

During signal transitions, the voltage and current inputs can vary enough to saturate the amplifier outputs. Thus, in order to maintain a fast settling time, clamping diodes where added to the first stage outputs. To further improve settling performance, we added a clamping signal to bring the second stage inputs together during track mode.

The DDA is constructed as a two-stage open loop amplifier in order to meet the noise and gain-bandwidth requirements. Unfortunately, this make the amplifier gain quite susceptible to process variation. These gain coefficients must be determined as part of the digital post-compensation process.

In simulation, the differencing amplifier had less than 0.2% Total Harmonic Distortion (THD) over its operating range, and it had a settling time of under 400 ps.

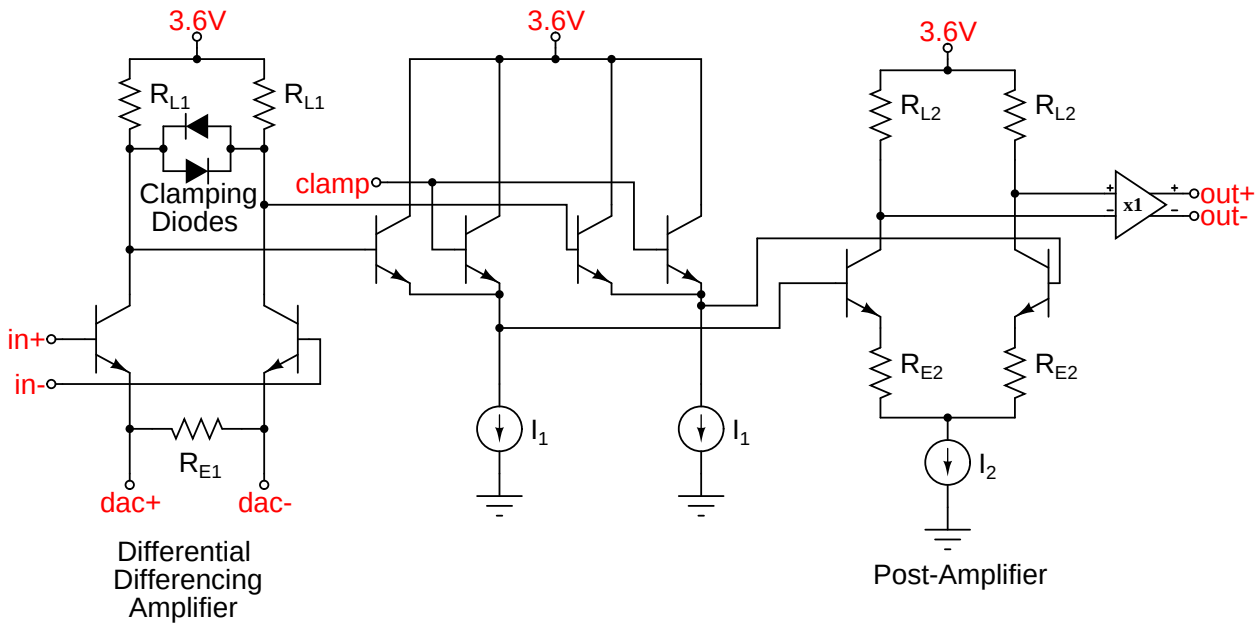

Figure 3.8 Differential Differencing Amplifier

3.1.4 Low Headroom DAC

In order to limit the design to a power supply range of 3.6 V, we needed a low noise DAC with reduced headroom requirement. Our final design for the DAC used a unary-weighted (thermometer-coded) architecture. The schematic for each DAC bit is shown in Fig. 3.9. The DAC bits are directly driven by the 15 differential comparators of the 4-bit FLASH, in order to obtain the shortest possible digital path delay.

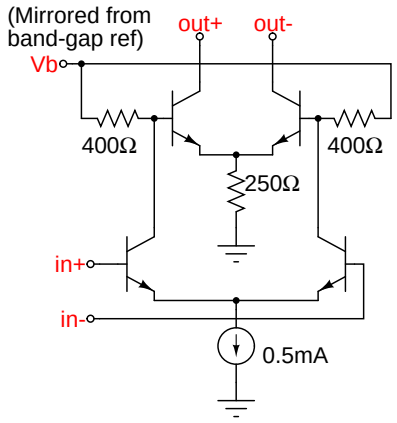

Figure 3.9 Low Headroom DAC bit

The DAC uses a current 2 mA in the main switching branch (upper branch in Fig. 3.9). The degeneration resistor in the main branch drops a relatively large voltage of 500 mV, which is key to the low-noise performance of the DAC. The secondary switch is set to pull the base contacts of the main switch down by only 200 mV. This ensures the main switch transistors stay out of saturation and maintains a fast current-switching time.

3.2 Error Correction by Bit Overlap

Between each of the four pipeline stages, there is nominally a single bit of overlap. Thus, the ADC in each stage need not have 12-bit accuracy. This also allows for incomplete settling at the

ADC input nodes. Distortion within the pipeline stages may also cause the signal to go into the error correction ranges.

Fig. 3.10 illustrates the error ranges for a typical conversion. We will examine this conversion in detail. First, the stage one ADC decides the first stage result is a 3, as indicated by a dot. The value should have been 4, so we end up in the upper error correction range for stage two. Stage two gets the value correct, but stage three makes another error. In this manner, we can obtain a 13-bit accurate result using four stages with only around 5-bits of accuracy each.

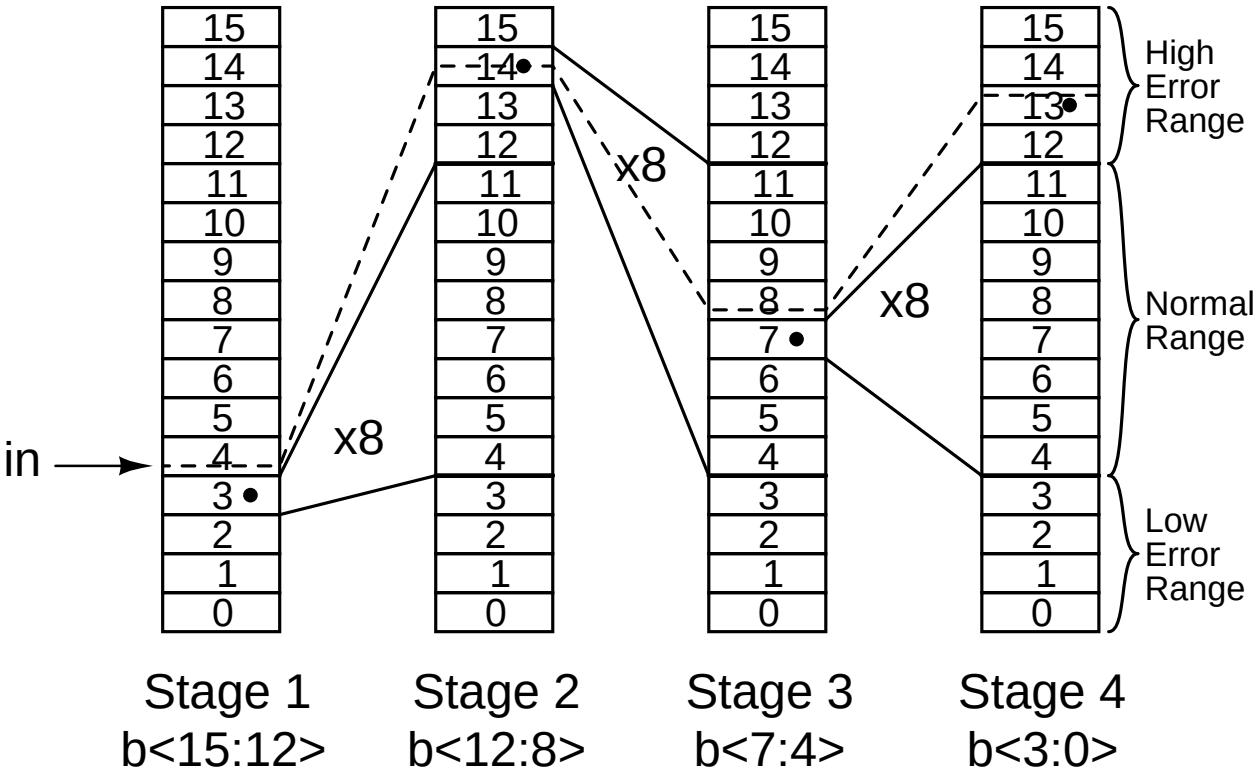

Figure 3.10 Error Correction Example: A Typical Conversion

3.3 Chip Layout

The layout for the BDMS-ADC which we fabricated is shown in Fig. 3.11. The high-speed analog and clock inputs are at the bottom left and bottom right. These inputs have redundant pads,

which allow for removing the wire-bonds and using RF wafer probes at the alternate pad locations. Most of the remaining pins are the 16 pairs of Low-Voltage Differential-Signaling (LVDS) ADC outputs.

The majority of the power consumption occurs in the lower half of the chip, where the clock driver, and first stage are located. The signal path continues upward to stages 2 and 3, which are each scaled down in power. At the top is the end of the signal path, where the final 4-bit flash resides.

3.4 Design Summary

Table 3.1 presents a summary of the simulated thermal noise and power performance of the critical design components. Also shown are total values for the entire chip. The ADC accepts a differential input voltage range of 1.8 V. Using this value in Eq. 2.2, we find that 106 μV of input noise results in 12.2 effective bits with respect to thermal noise.

Component	Thermal Noise (μV_{rms})	Power (mW)
Input Buffer	21.9	189
Clock Buffer	N/A	266
THA 1	66.5	417
DAC 1	60.4	214
Residue Amp 1	51.9	190
Other	-	935
TOTAL	106.0	2211

Table 3.1 Simulated Noise and Power Contributions for Major Components

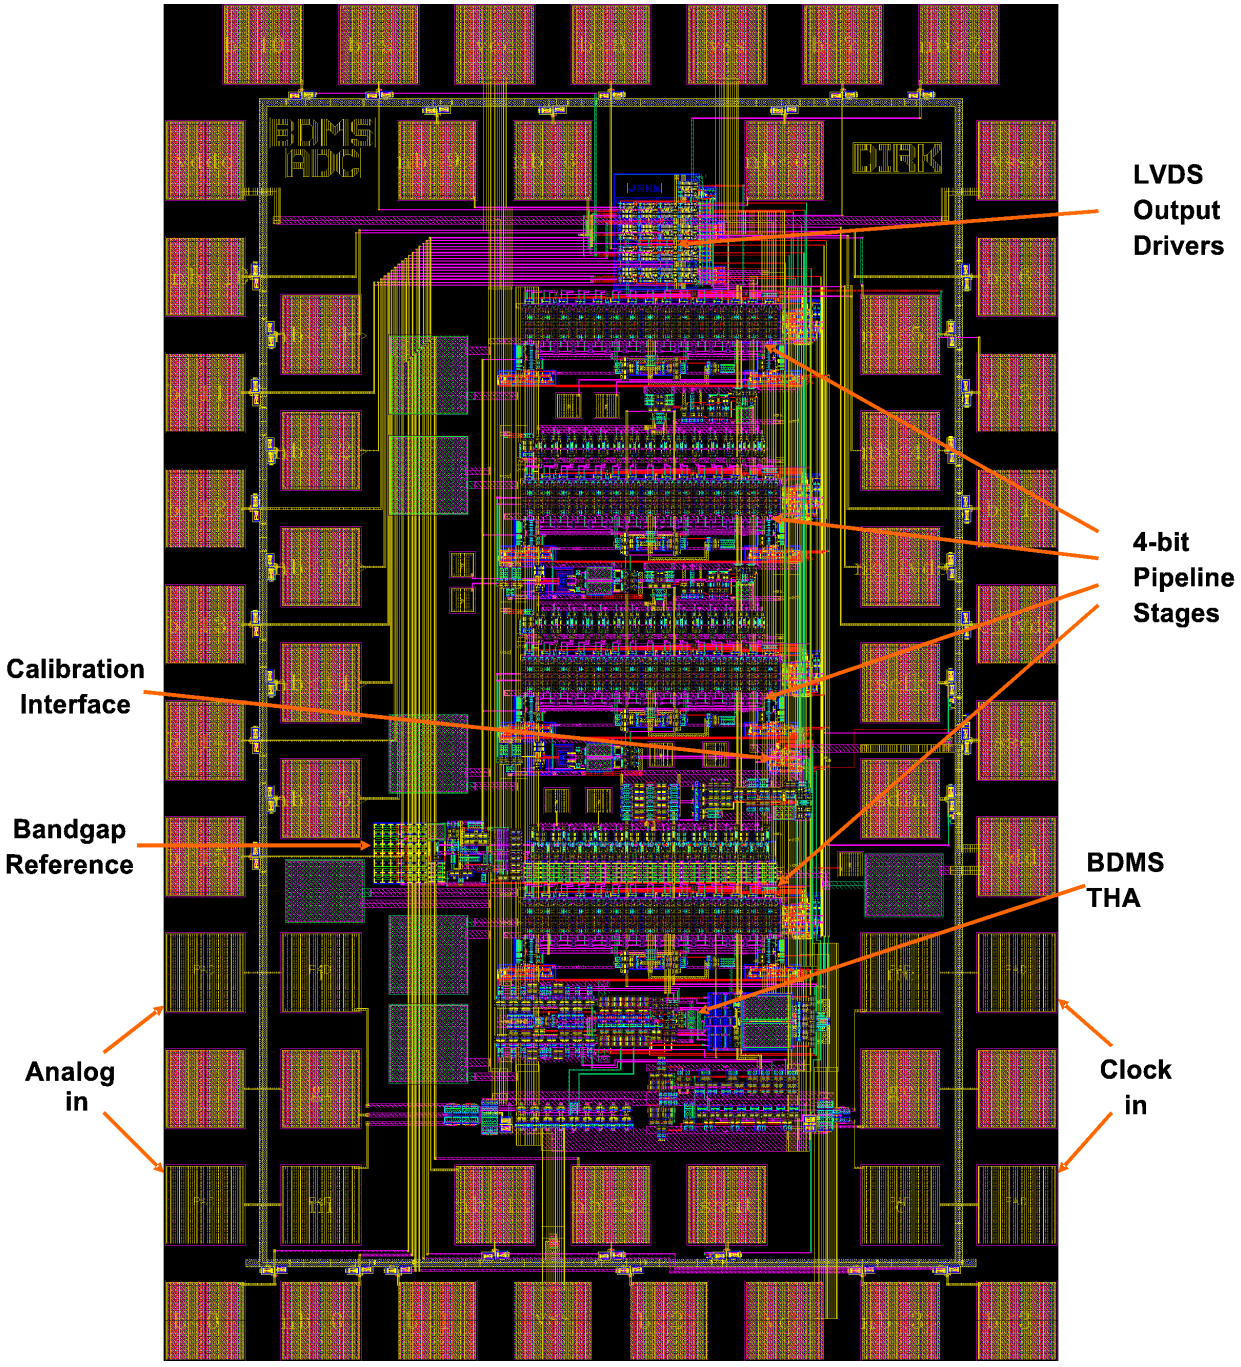

Figure 3.11 ADC Chip Layout

Chapter 4

Digital Distortion Post-Compensation

From the beginning of the design of the BDMS-ADC, we intended to correct the combined distortion of the THA and ADC using digital post-compensation. This freed us to focus on reducing the voltage headroom, power consumption, and sampling noise.

In recent literature, there are several methods for using digital post-compensation with ADCs that have monotonically increasing output codes [14, 15, 16]. However, because of the error ranges built in to the BDMS-ADC design, we must extend the previous work to be used with this design.

For this version of the ADC, we apply the post-compensation off-chip to complete records. Different methods of post-compensation were investigated for potential on-chip integration. We will see that an estimate of the input slope is required for best results. For this version, we determined the slope estimate by assuming a sinusoidal input and using a least-squares fit to find the input amplitude, frequency and phase.

4.1 ADC Distortion Model

The correction of nonlinearity and dynamic distortion relies on a simplified distortion model as shown in Fig. 4.1. Since only the first stage must deal with a full resolution analog signal, this stage will be responsible for most of the distortion. As shown, the distortion of the first stage is lumped into two nonlinear elements, with one before and one after the differencing amplifier. We will consider all the other elements to be ideal. We then solve for the intermediate residual values

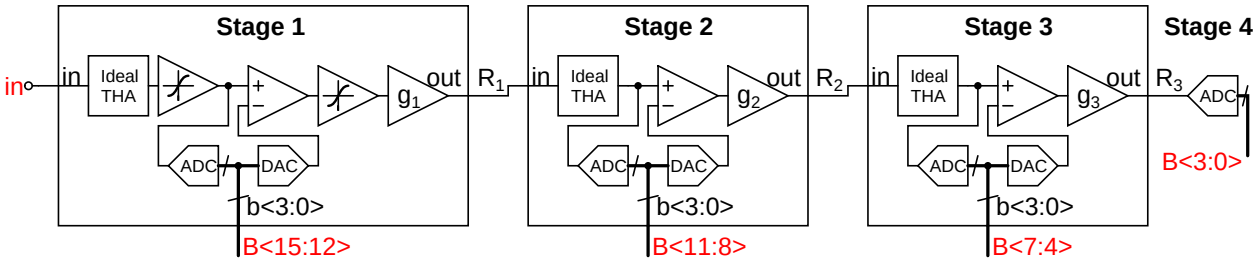

Figure 4.1 Simplified Pipelined ADC Distortion Model

R_1, R_2, R_3:

$$R_1 = g_1 f_2\left(f_1(V_{in}) - B_{15:12}\right) \tag{4.1}$$

$$R_2 = g_2(R_1 - B_{11:8}) \tag{4.2}$$

$$R_3 = g_3(R_2 - B_{7:4}) \tag{4.3}$$

where each residual is in units of the LSB of the stage it is input to, i.e. R_1 is in units of the second stage LSB, etc. Neglecting quantization error, $B_{3:0} = R_3$. Thus, from Eq. 4.2 and Eq. 4.3, we have:

$$R_1 = B_{11:8} + \frac{1}{g_2}\left(B_{7:4} + \frac{1}{g_3}B_{3:0}\right) \tag{4.4}$$

Now, we must invert Eq. 4.1 in order to recover V_{in} from the digital output. Let $f_1^{-1}(), f_2^{-1}()$ represent the inverses of the distortion functions $f_1(), f_2()$. Thus, we can extract V_{in}, in units of the first stage LSB:

$$\begin{aligned} V_{in} &= f_1^{-1}\left(B_{15:12} + f_2^{-1}\left(\frac{1}{g_1}R_1\right)\right) \\ &= f_1^{-1}\left(B_{15:12} + f_2^{-1}\left(\frac{1}{g_1}\left(B_{11:8} + \frac{1}{g_2}\left(B_{7:4} + \frac{1}{g_3}B_{3:0}\right)\right)\right)\right) \end{aligned} \tag{4.5}$$

Theoretically, one could reduce all of the significant sources of distortion in the ADC to a model with just a few parameters. Then, that model could be inverted to find the correct input analog value from the digital result. In practice, we used visual inspection of the corrected data of a distortion model to suggest further refinement to the distortion model.

4.1.1 "gain" model: Finding Inter-stage Gain Coefficients

First, consider an ADC model where the inter-stage gain coefficients are nominally 8 but are actually g_1, g_2, g_3. Assume there are no other sources of distortion, i.e. $f_1(x) = x, f_2(x) = x$ in Eq. 4.5. This will be referred to as the "gain" model. Since the residue amplifiers consist of two-stage open-loop amplifiers, the gain may vary significantly due to process variations. This results in the following simple model:

$$V_{in}(k) = B_{15:12}(k) + \frac{1}{g_1}(B_{11:8}(k) + \frac{1}{g_2}(B_{7:4}(k) + \frac{1}{g_3}B_{3:0}(k))) \qquad \text{(in units of Stage-1 LSBs)} \tag{4.6}$$

where $B_{15:12}, B_{11:8}, B_{7:4}$ and $B_{3:0}$ are the binary outputs from stage one to stage four and $k = 0, 1, \ldots, (N-1)$ are individual sample numbers from a data record of N samples.

We can make Eq. 4.6 linear in terms of the fitting parameters if we make these substitutions: $y_1 = \frac{1}{g_1}, y_2 = \frac{1}{g_1 g_2}$ and $y_3 = \frac{1}{g_1 g_2 g_3}$. This makes it possible to solve using the linear least-squares technique, if we have a data set with known values for V_{in}. In practice, V_{in} is generated by a sinusoidal source. We assume V_{in} is a pure sinusoid with known frequency, but we may only approximately know its amplitude, offset and phase. Thus we must fit the parameters of the input sinusoid at the same time as the model parameters. This results in the linear set of equations:

$$Fit(k) = c_0 + c_1 cos(2\pi f k) + c_2 sin(2\pi f k) \tag{4.7}$$

$$V_{in}(k) = B_{15:12}(k) + y_1 B_{11:8}(k) + y_2 B_{7:4}(k) + y_3 B_{3:0}(k) \tag{4.8}$$

$$Fit(k) = V_{in}(k) \tag{4.9}$$

where f is input frequency normalized to the sample rate, $f = \frac{f_{in}}{f_{clk}}$. Note that since we will use optimization techniques to find parameters on both sides of Eq. 4.9, we must take care to keep the problem constrained. In this case, the problem is constrained by virtue of the fact that $B_{15:12}$ does not have an adjustable scaling coefficient.

Eq. 4.9 can be transformed into the canonical least squares matrix form $A\mathbf{x} = \mathbf{b}$, where $A, \mathbf{x}, \mathbf{b}$ represent the known coefficients, unknown parameters, and data values, respectively. We obtain:

$$A = \begin{bmatrix} -\cos((0)2\pi f) & -\sin((0)2\pi f) & B_{11:8}(0) & B_{7:4}(0) & B_{3:0}(0) \\ -\cos((1)2\pi f) & -\sin((1)2\pi f) & B_{11:8}(1) & B_{7:4}(1) & B_{3:0}(1) \\ & & \vdots & & \\ -\cos((N-1)2\pi f) & -\sin((N-1)2\pi f) & B_{11:8}(N-1) & B_{7:4}(N-1) & B_{3:0}(N-1) \end{bmatrix} \tag{4.10}$$

$$\mathbf{x} = \begin{bmatrix} c_1 \\ c_2 \\ y_1 \\ y_2 \\ y_3 \end{bmatrix}, \qquad \mathbf{b} = \begin{bmatrix} c_0 - B_{15:12}(0) \\ c_0 - B_{15:12}(1) \\ \vdots \\ c_0 - B_{15:12}(N-1) \end{bmatrix}$$

Which has a least squares solution of:

$$\mathbf{x}_{opt} = (A^T A)^{-1} A^T \mathbf{b} \tag{4.11}$$

Note that $A^T A$ is only a 5x5 matrix, so it is feasible to find the solution with on-chip digital circuitry or at least with an on-board micro-controller.

4.1.2 "d1" model: Finding First-Stage DAC Offsets

The next refinement to the ADC model is to account for offset errors in the DAC of the first-stage. This model will be referred to as the "d1 model". This can also be accomplished using the linear least-squares technique.

To model DAC offset errors, we can simply add a digital offset to the ADC result depending on the result of the first stage. To implement this, instead of having the result be simply proportional to the stage one result, $B_{15:12}(k)$, we use a 16-element table of DAC values which is indexed by this integer value. This function is called a look-up table, which we will denote as $LUT[B_{15:12}]$.

Thus:

$$V_{in}(k) = LUT[B_{15:12}(k)] + y_1 B_{11:8}(k) + y_2 B_{7:4}(k) + y_3 B_{3:0}(k) \tag{4.12}$$

To transform Eq. 4.12 into a form suitable for linear least squares, we will invoke the discrete delta function, which is defined as:

$$\delta_{m,n} = \begin{cases} 1 & if\ m = n \\ 0 & if\ m \neq n \end{cases} \tag{4.13}$$

The linear form of the equation set is then:

$$\begin{aligned} V_{in}(k) &= LUT[0]\delta_{0,B_{15:12}(k)} + LUT[1]\delta_{1,B_{15:12}(k)} + \cdots + LUT[15]\delta_{15,B_{15:12}(k)} \\ &\quad + y_1 B_{11:8}(k) + y_2 B_{7:4}(k) + y_3 B_{3:0}(k) \\ Fit(k) &= c_0 + c_1 cos(2\pi f k) + c_2 sin(2\pi f k) \\ Fit(k) &= V_{in}(k) \end{aligned} \tag{4.14}$$

However, while the gain model (Eq. 4.9) was constrained in amplitude on the right-hand side, Eq. 4.14 is unconstrained on both sides. Attempting to use linear least-squares optimization will result in all zero coefficients. One remedy is to enforce $LUT[11] = LUT[4] + 7$. This effectively normalizes the fit to average the middle seven LSB steps of the first stage. We may write this in the form $A\mathbf{x} = \mathbf{b}$, where A has rows of the form:

$$\begin{aligned} A(k) = & \\ & [-\cos(2\pi f k) \quad -\sin(2\pi f k) \quad B_{11:8}(k) \quad B_{7:4}(k) \quad B_{3:0}(k) \\ & \delta_{0,B_{15:12}(k)} \quad \delta_{1,B_{15:12}(k)} \quad \delta_{2,B_{15:12}(k)} \quad \delta_{3,B_{15:12}(k)} \\ & (\delta_{4,B_{15:12}(k)} + \delta_{11,B_{15:12}(k)}) \quad \delta_{5,B_{15:12}(k)} \quad \delta_{6,B_{15:12}(k)} \quad \delta_{7,B_{15:12}(k)} \\ & \delta_{8,B_{15:12}(k)} \quad \delta_{9,B_{15:12}(k)} \quad \delta_{10,B_{15:12}(k)} \\ & \delta_{12,B_{15:12}(k)} \quad \delta_{13,B_{15:12}(k)} \quad \delta_{14,B_{15:12}(k)} \quad \delta_{15,B_{15:12}(k)}] \end{aligned} \tag{4.15}$$

And $\mathbf{x}, \mathbf{b}$ are:

$$\mathbf{x}^T = $$

$$[c_1 \quad c_2 \quad y_1 \quad y_2 \quad y_3$$

$$LUT[0] \quad LUT[1] \quad LUT[2] \quad LUT[3]$$

$$LUT[4] \quad LUT[5] \quad LUT[6] \quad LUT[7]$$

$$LUT[8] \quad LUT[9] \quad LUT[10]$$

$$LUT[12] \quad LUT[13] \quad LUT[14] \quad LUT[15]\ ]$$

$$\mathbf{b}^T = [c_0 - 7\delta_{11,B_{15:12}(0)} \quad c_0 - 7\delta_{11,B_{15:12}(1)} \quad \dots \quad c_0 - 7\delta_{11,B_{15:12}(N-1)}]$$

Again, the linear least-squares solution is $\mathbf{x}_{opt} = (A^T A)^{-1} A^T \mathbf{b}$. Note that the solution remains a one-step procedure, making it computationally efficient. Now $(A^T A)$ has a size of 20x20, which should still be feasible to compute on-chip or on-board. Also, note that while the matrix A is large, it has mostly elements of zero value. Thus, the least-squares solution can be computed without needing to storing the complete A matrix in memory.

4.1.3 Nonlinear Coefficient Fitting

More complex error models may not have a linear dependence on all parameters. In such cases, we can apply the iterative method of non-linear least squares. Consider the optimization problem:

$$min \sum_{k=0}^{N-1} (V_{in}(k, \mathbf{x}) - Fit(k, \mathbf{x}))^2 = min \sum_{k=0}^{N-1} f(k, \mathbf{x})^2 \tag{4.16}$$

where we assume $f(k, \mathbf{x})$ is nonlinear with respect to the fitting parameters, $\mathbf{x}$. Using a first-order Taylor expansion about an initial guess for the parameters of $\mathbf{x}_0$ results in:

$$f(k, \mathbf{x}) \approx f(k, \mathbf{x}_0) + \frac{\partial f(k, \mathbf{x})}{\partial \mathbf{x}} \Delta \mathbf{x} = -\mathbf{b}(k) + A(k) \Delta \mathbf{x} \tag{4.17}$$

Using this Taylor approximation, Eq. 4.16 becomes a linear least squares problem. Experience has shown that after finding initial values for the parameters from the gain model described above, only 3 to 4 iterations are needed to precisely determine the nonlinear parameters.

The vector derivative term in Eq. 4.17 is commonly referred to as the Jacobian. Finding formulas for each column of the Jacobian proved cumbersome in practice. For the purposes of evaluating distortion models on a PC, the nonlinear least-squares optimizer from the SciPy package for Python was used [17, 18, 19, 20]. The SciPy optimizer can estimate the Jacobian, so it was only necessary to specify the fitting equation. If the nonlinear distortion model were selected, one would need to find formulas for the Jacobian in order to make a micro-controller-based solution feasible and efficient.

4.2 Narrow-band Distortion Correction

Narrow-band distortion correction refers to finding an error model for inputs at a single input frequency at an amplitude of full-range. Frequency or slope dependence need not be considered, although the models tended to work best when applied to positively and negatively sloped data separately.

The gain model described above is the simplest model we will consider, involving only three interstage gain coefficients. Fig. 5.12 shows the error with a sinusoidal input, for positive slope only. The figure is colored according to the first stage binary output, which is also the first stage DAC value. The SINAD after applying this correction was 50.3 dB, corresponding to 8 effective bits.

Looking at lower stage one outputs in Fig. 5.12, we see that the error of the residual after stage one looks something like a second order polynomial. So, we developed an inverse model that applied a second-order polynomial to the residual value of stage one. The polynomial coefficients are fit separately for each value of the first stage outputs. This model is named “d1_g1_c1”, because

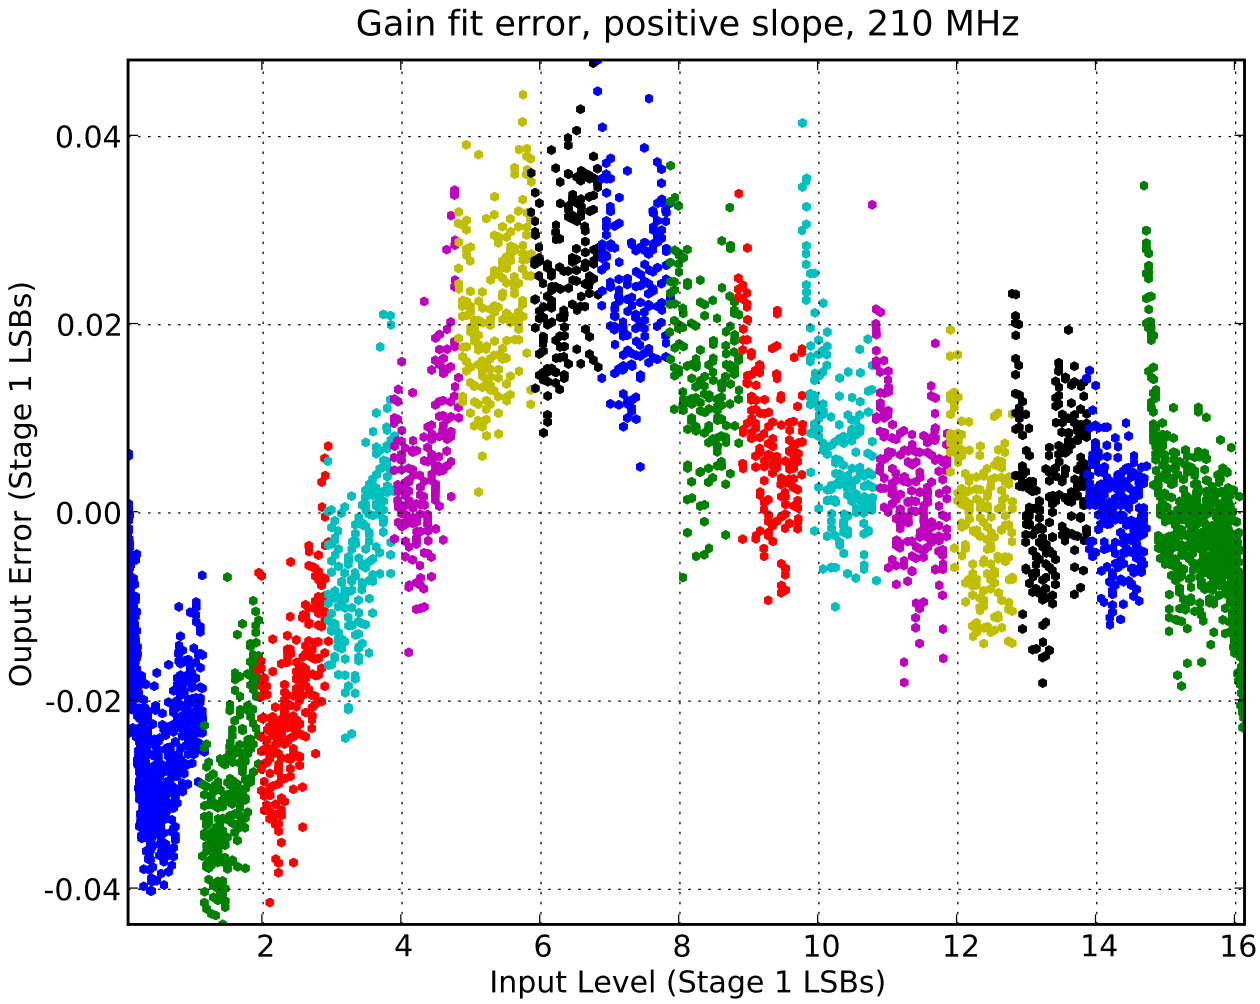

Figure 4.2 Gain fit error, color coded by stage one output ($F_{in} = 210.01\,\mathrm{MHz}$, $F_{clk} = 100\,\mathrm{MHz}$)

in addition to a lookup table for the DAC offset, it also has lookup tables for first stage residual gain and curvature. The model has the form:

$$V_{in}(k) = LUT_d[B_{15:12}(k)] + LUT_g[B_{15:12}(k)]R_1 + LUT_c[B_{15:12}(k)]R_1^2 \tag{4.18}$$

where LUT_d, LUT_g, LUT_c are each 16-element lookup tables corresponding to the offset, gain and curvature. The stage one residual R_1 is defined by:

$$y_1 B_{11:8}(k) + y_2 B_{7:4}(k) + y_3 B_{3:0}(k) \tag{4.19}$$

The "d1_g1_c1" model involves 50 parameters, so we used records of at least 8192 points to precisely determine the parameters. Although determining the parameters from a nonlinear least-squares fit is computationally intensive, once the parameters are known, this post-compensation model (Eq. 4.18) could be easily implemented on chip.

We studied several other distortion correction models. Brief descriptions of each of the models are listed in Table 4.1.

4.3 Broadband Distortion Correction

The nonlinearity characteristics of an ADC tend to vary with input frequency and amplitude. This is due to the nonlinear settling behavior of the ADC components. To account for the frequency and amplitude dependence, we can make the distortion model depend on the slope of the input in addition to the input level. Equivalently, the inverse distortion model will depend on the binary output code and input slope estimate. The challenge is to add slope dependence to the single-frequency fits without adding too many fitting parameters.

The ideal way to estimate the slope of V_{in} would be to use a separate ADC to measure it within the same chip. The slope would be recorded in parallel with the input signal, simultaneously sampled with the input. Frequency dependence of the distortion tended to be at the level of one part in 8 bits, so a 4-bit flash converter should be sufficient for the slope estimate.

Name	# params.	Description
gain	3	Determines inter-stage gain coefficients only.
d1	19	Gain coefficients and a LUT for DAC offset.
d1_g1	36	Two LUTs for the DAC offsets and residual gain of the first stage, indexed by the first stage binary result.
d1_g1_c1	50	Three LUTs for the DAC offsets, residual gain, and curvature of the first stage, indexed by the first stage binary result.
d1_nl	23	Based on "d1", but also fits a third order polynomial to the first stage THA response.
n33	11	This model uses a third-order two-dimensional polynomial with inputs of the first stage binary output and first stage residual.
d1_n33	21	Like "n33", but also fits offsets for the first stage DAC.
d1_d2	37	A 16-element LUT for the offsets of each of the first two stage DAC outputs, and inter-stage gain coefficients.
d1_d2_g1_g2	65	Four 16-element LUTs for the first and second stage offsets and gains.
d12	258	A 256-element lookup table of offsets indexed by both the first and second stage binary results.
d12_g12	513	Two 256-element lookup table for the combined offset and residual gain of the first two stages taken together.

Table 4.1 Single Frequency Model Descriptions

Unfortunately, our first revision of the BDMS-ADC does not contain circuitry to estimate the input slope. Instead, we estimate the input slope by fitting the input to a sinusoid. This of course would not work with a truly broadband input, but it allows us to easily compare the performance of various distortion correction algorithms.

It is also possible to construct a slope estimate by applying digital filtering to the ADC binary outputs. After post-processing the data with a non-slope dependent correction, the data is fed into a slope-estimation filter. Hummels [21] used a 23-point FIR filter to adequately estimate the slope up to 80% of the Nyquist frequency. However, such a technique is prone to aliasing errors, which would be avoided by measuring the slope directly.

4.3.1 Broadband Model Descriptions

In order to add slope dependence to our existing single-frequency models, we simply allowed each parameter to vary as a second-order polynomial. Such a polynomial may be described by three coefficients, so we end up with three times the number of fitting parameters for each inverse distortion models. The use of this type of slope dependence is indicated by the suffix "_s" appended to a model name, as in "d1_g1_c1_s".

As suggested by Larrabee [16], we found that the optimal parameters in our single-frequency distortion models were different for positive and negative slopes of V_{in}. This is born out in Table 5.1 and Table 5.2, by the fact that the models work better when applied to only positive or negatively sloped data as compared to when applied to the entire data set. Also, note that the performance of each model tends to favor the positive slope input set over the negative slope set. This suggests that variance in the model parameters is not symmetrical with slope. Since the ADC uses a differential design, we suspect this asymmetry is due mostly to the balun at the ADC input.

In order to account for the asymmetry between positive and negative slopes, we allowed the slope dependence to use a different second-order polynomial for positive and negative input slopes.

This way, the slope dependence may be discontinuous around zero input slope. The use of this type of slope dependence is indicated by the suffix "_sd" appended to a model name, i.e. "d1_g1_c1_sd".

4.3.2 Broadband Model Parameter Optimization

In order to make sure we have adequate coverage of the input state space when using a sinusoidal input, we must take many records at unique amplitudes and frequencies. Fig. 4.3 shows the state coverage for input frequencies of 10, 20, 50, 100, 130, 170, 210 MHz and amplitudes of 40, 60, 80 and 100 percent of full scale. The input slope has been normalized to the maximum slope of a full range signal at the Nyquist frequency.

To determine the model parameters, the non-linear least squares technique was applied to the full set of data records simultaneously. The number of rows in the fitting equation is increased by the number of input frequencies multiplied with the number of input amplitudes, in this case a factor of 28. However, the main increase in computational burden is due to the increase in number of parameters. Additionally, the amplitude, frequency, phase and offset must be determined for each data set, greatly adding complexity. To illustrate the added complexity, the Python implementation of the single-frequency model "d1_g1_c1" took 2 seconds to complete, while the implementation of "d1_g1_c1_sd" took 1670 seconds.

In order to find the model parameters with a micro-controller for these slope-dependent models, it would be necessary to reduce the number of slope dependent parameters and determine formulas for the Jacobian matrix.

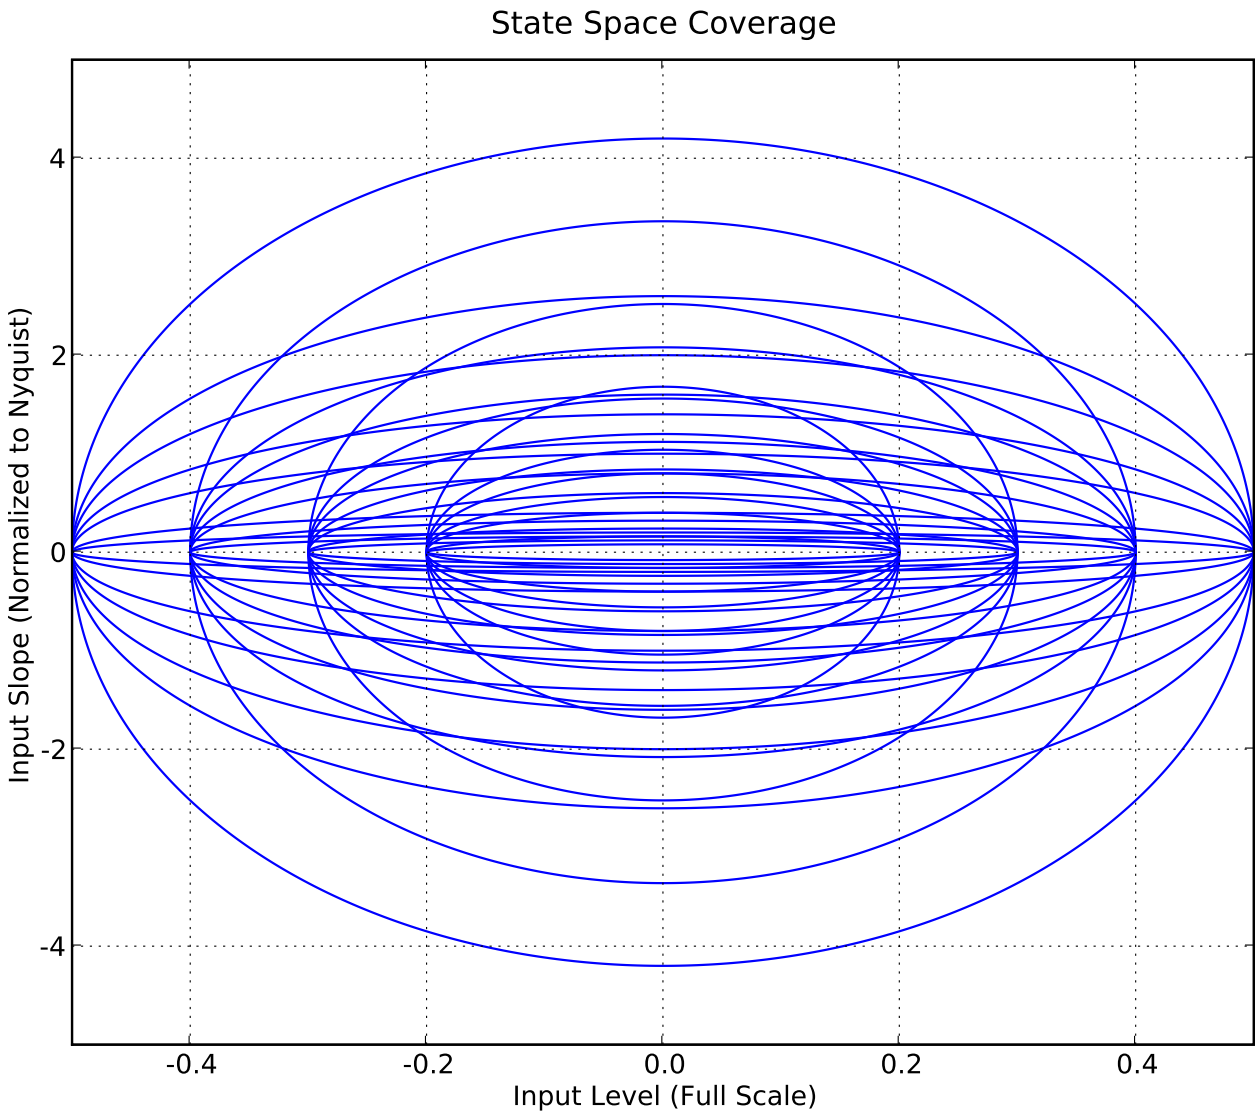

Figure 4.3 Input State Coverage Map

Chapter 5

Measurement and Results

The BDMS-ADC was developed in the Jazz Semiconductor 0.18 μm SiGe BiCMOS process. It was then professionally wire-bonded to a 64-pin quad-flat pack (QFP) package. A micrograph of the die is shown in Fig. 5.1. The die size is 1.8 mm by 1.2 mm.

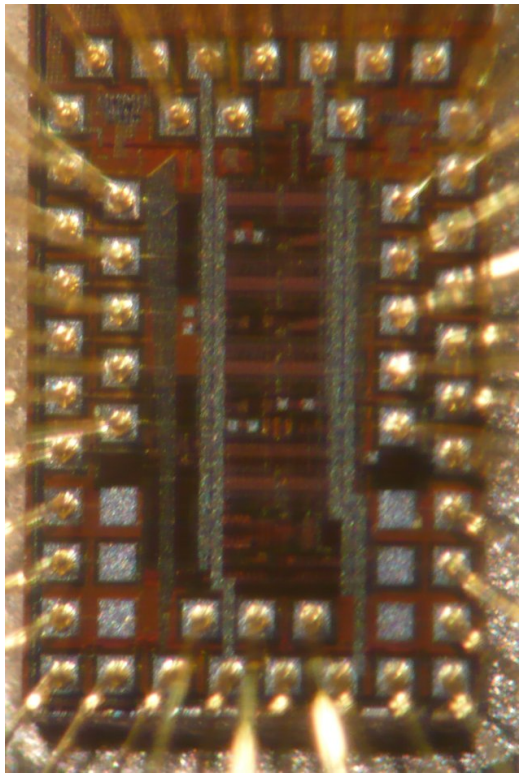

Figure 5.1 BDMS-ADC Die Photo

Due to sub-optimal routing of the digital outputs of the first two pipeline stages, the binary outputs begin to show glitches at certain transitions when the clock frequency approaches 200 MHz, as shown in Fig. 5.2. This is well short of our design goal of 500 MHz. However, we can still test the high-frequency capabilities of the input THA by operating in sub-sampling mode ($f_{in} > f_{clk/2}$). Even at a clock frequency of 100 MHz, where no glitches occur, the ADC has competitive performance.

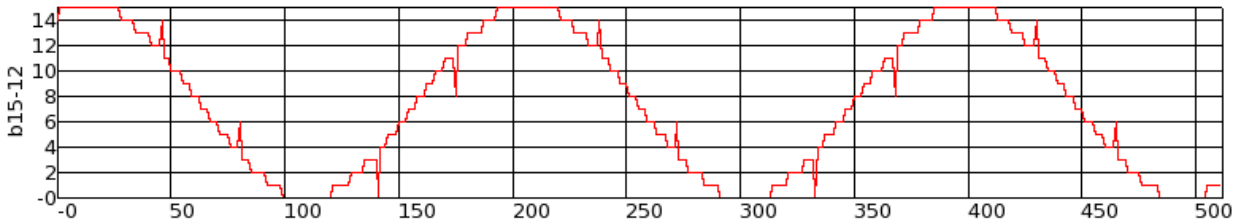

Figure 5.2 Glitches in Stage 1 Output (F_{clk} = 200 MSps)

5.1 Test Setup

A diagram and photograph of the test arrangement are shown in Fig. 5.3 and Fig. 5.4. The ADC input and clock are generated by synthesized signal generators. The two synthesized sources were phase locked by a 10 MHz link. We found that the distortion level of these sources was between 45 and 50 dB. This corresponds to only about 7 effective bits. Consequently, we added fourth-order bandpass filters to the analog input to ensure significant spectral purity and to limit the integrated input noise. We used filters at 10, 20, 50, 70, 100, 130, 210 and 374 MHz, which had to be manually swapped between data recordings.

The outputs of the ADC were recorded by a Xilinx Virtex-4 FPGA, which is part of the Digilent FX-12 development board. The interface supports up to 500 MSps by using Low-Voltage Differential Signaling (LVDS). A continuous data recording of 32 k-samples is stored in the FPGA memory and later read out by a computer over a 57.6 kBps serial link.

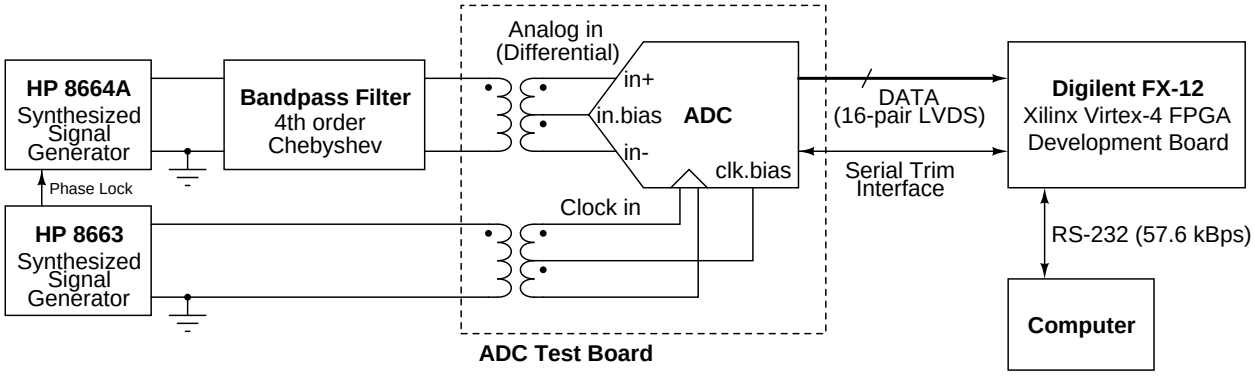

Figure 5.3 Test Setup Schematic

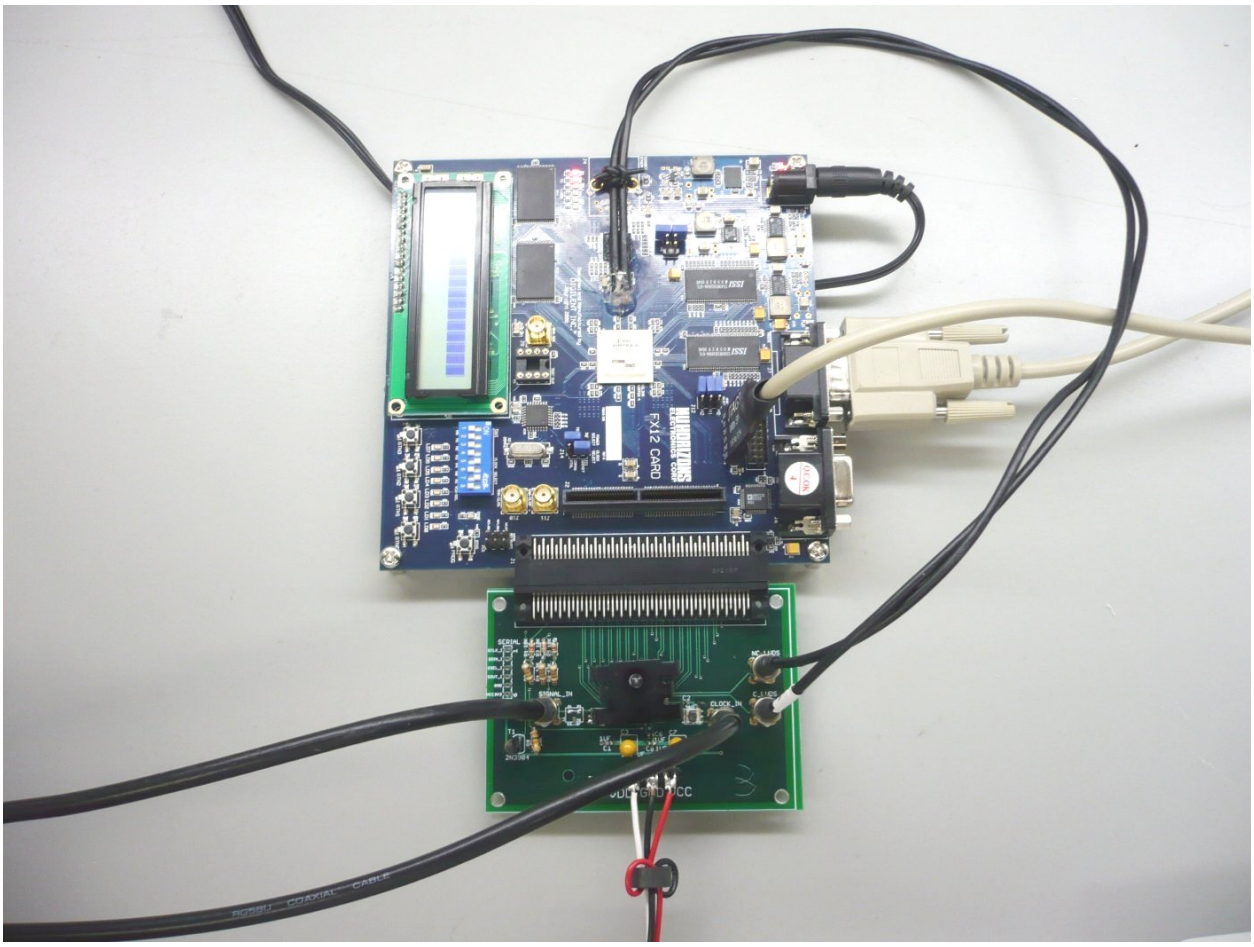

Figure 5.4 Test Setup Photo

5.1.1 Test Board

A photograph of the test board is shown in Fig. 5.5. This PCB contains a single-ended to differential balun for both the clock and analog inputs. These baluns also perform a DC level shift. The ADC chip itself generates the bias voltages necessary for the level shift in each case.

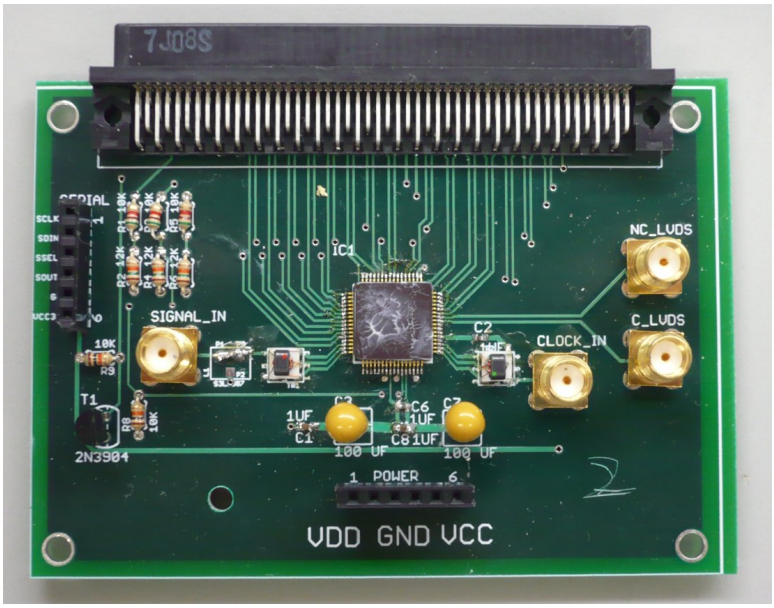

Figure 5.5 Test PCB Photo

The balun used at the ADC analog input was Coilcraft part number WBC2-1T. This balun has a 2:1 impedance ratio, which is ideal for transforming from a 50 Ω source to a differential pair of signals which are each terminated at 50 Ω. The 3 dB bandwidth of this balun extends from 0.2 MHz to 500 MHz.

The ADC has a 2:1 divider in its clock buffer. Since the ADC was intended to be used up to 500 MSps, the balun for the clock buffer was required to support up to a 1 GHz frequency. We used the Coilcraft WBC4-14, which has a 4:1 impedance ratio and a bandwidth extending from 1.5 MHz to 1.2 GHz

The ADC outputs connect to the FPGA by means of an 100-pin connector. This connector is the counterpart to one on the Digilent FX-12 prototyping board. This greatly reduced the clutter of the test setup.

5.1.2 FPGA Firmware

Fig. 5.6 shows a schematic of the FPGA firmware we developed. The FPGA is responsible for storing data records at clock rates of up to 500 MSps. To capture data at such a high rate, it is necessary to use input deserializers, which increase the data width by a factor of 8. Next, the data is written to a 32 k-word first-in-first-out buffer (FIFO). On the Xilinx FPGA, the FIFO blocks allow use of separate clocks on the write and read sides. We take advantage of this feature to use a constant clock for the control logic on the FPGA, independent of the ADC sampling rate. The three clock domains are indicated in Fig. 5.6 by dashed lines.

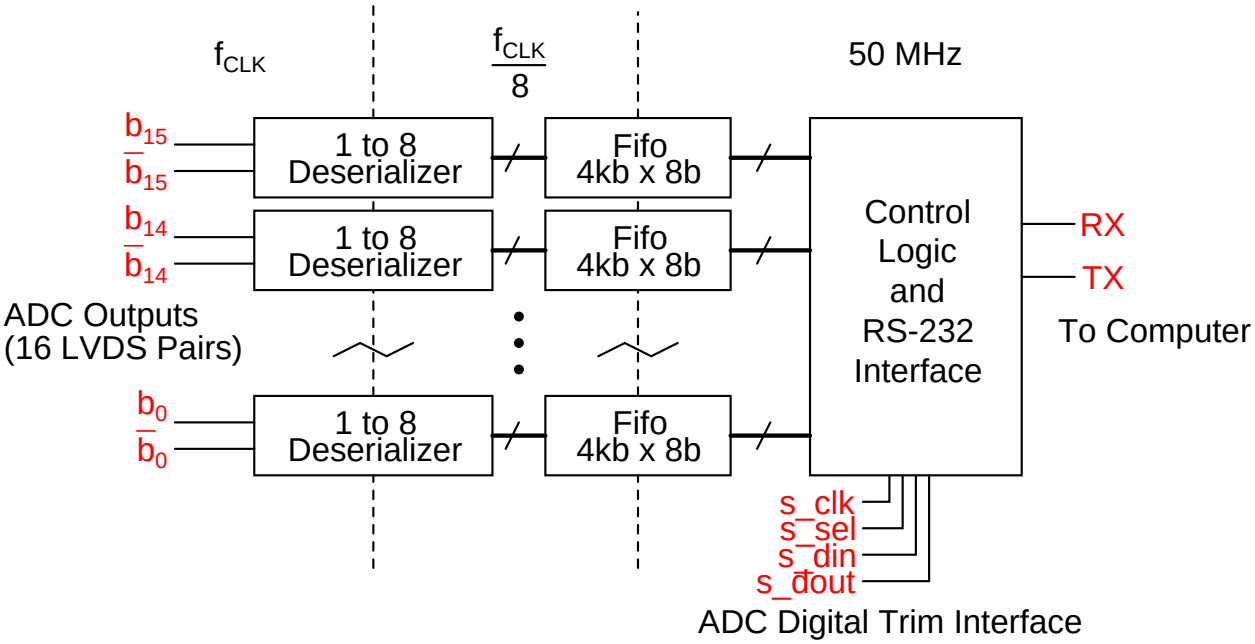

Figure 5.6 FPGA Data Recorder Schematic

The FPGA is also responsible for relaying digital trim commands from the computer to the ADC. The trim controller on the ADC uses simple serial interface, in which the "s_sel" signal designates sending an address or data.

5.1.3 Software for Acquisition and Analysis

The BDMS-THA requires calibration of the ADC range and offset in each stage in order to prevent signal clipping. To facilitate this, GUI was developed to show the ADC response in near real-time. This interface is shown in Fig. 5.7 with a typical low frequency sinusoidal input.

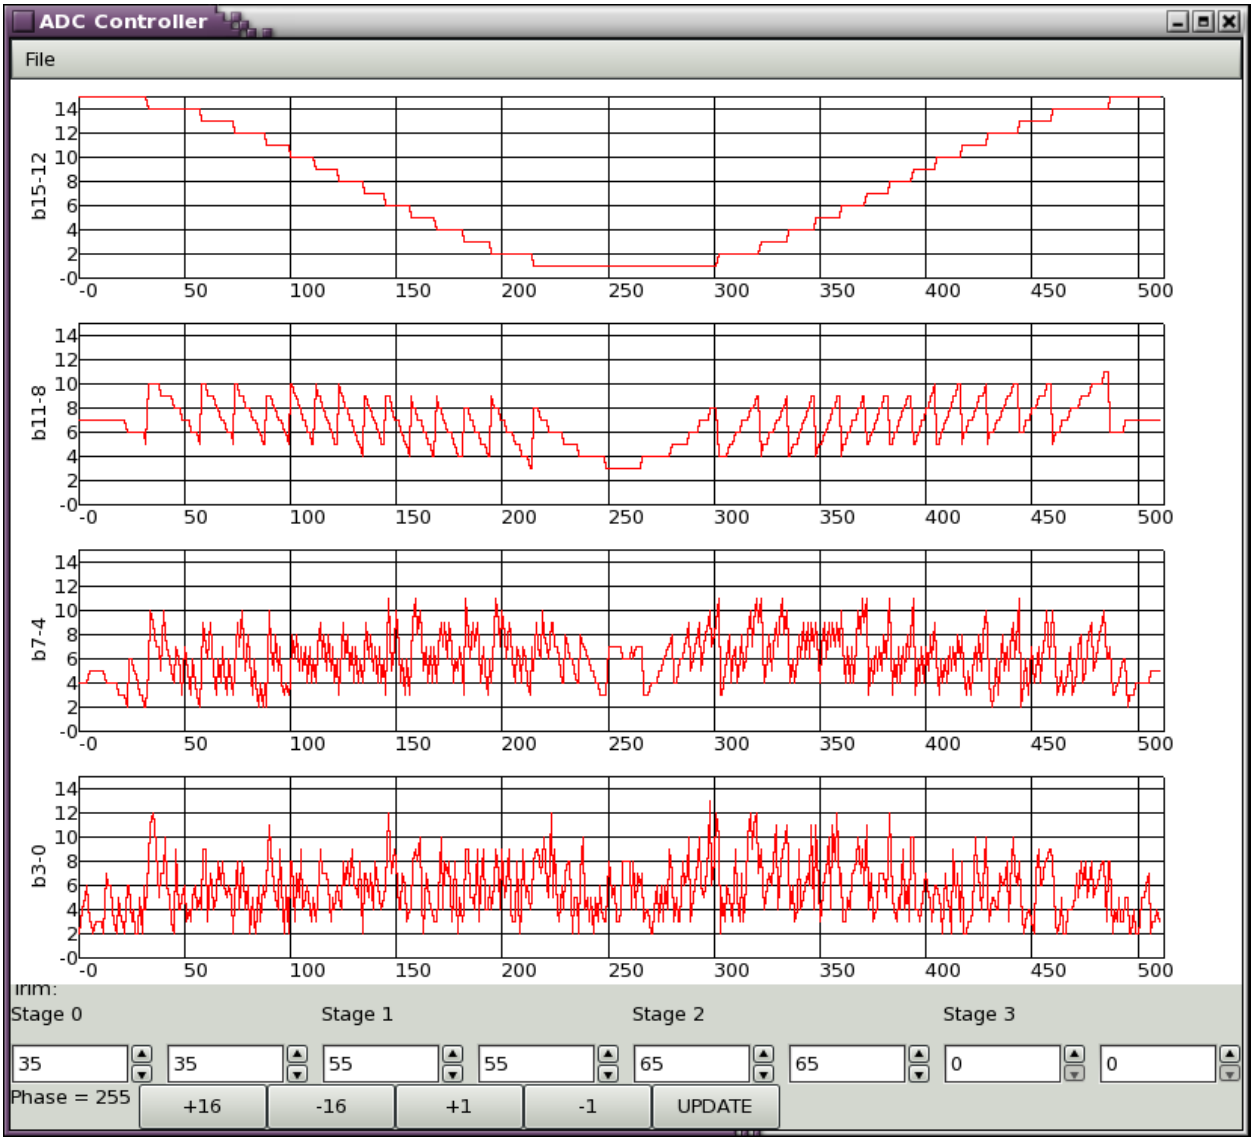

Figure 5.7 ADC Controlling Software GUI

5.2 Results

5.2.1 DC Noise

We obtained a measure of the thermal noise response at DC by terminating the ADC input to 50 Ω, after the balun. This produced a data recording, in which the first three pipeline stages had constant outputs. Only the output of the final stage ADC varied. Thus, in order to obtain the noise power, we need only know the gain coefficient of stage 4, which is readily obtained from fitting the gain model to a low frequency input. The resulting ADC output is shown in Fig. 5.8. When we compare the noise level from this recording to a full scale sinusoidal input, the result is a signal to noise ratio of 67.9 dB or 11.0 effective bits.

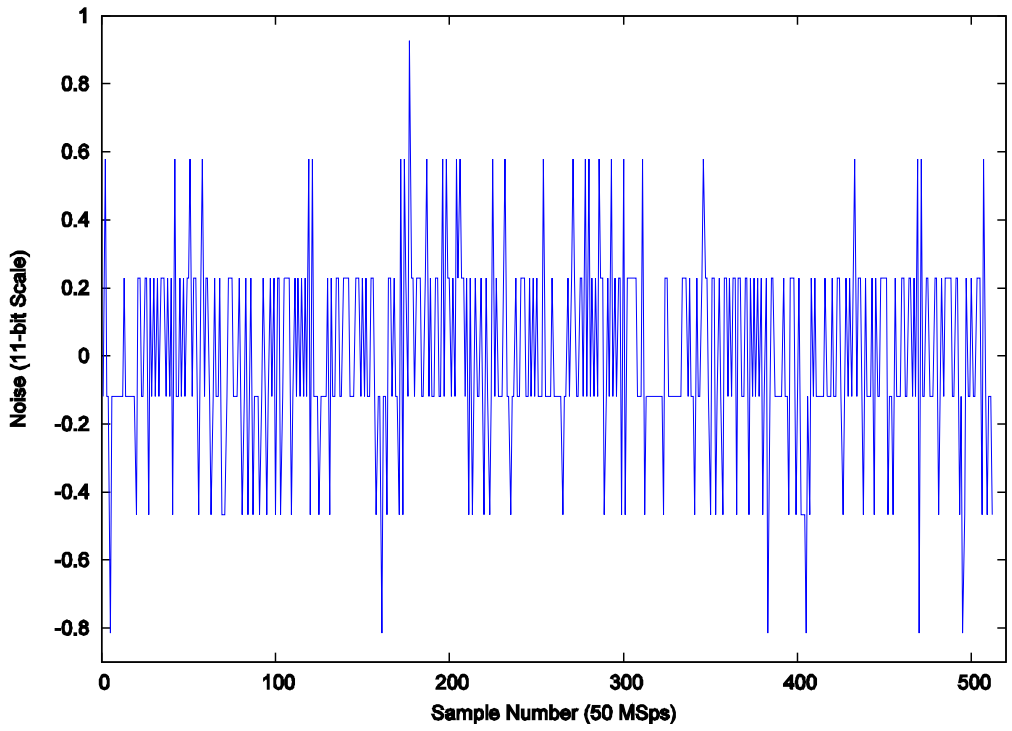

Figure 5.8 DC Noise, input terminated to 50 Ω

5.2.2 Jitter

The sampling time jitter of the BDMS-ADC was measured by applying the same signal to the ADC input and clock. In this manner, the jitter of the signal generator is removed from the measurement. This results in an ADC output similar to that of the thermal noise measurement. Using frequencies of 210 and 374 MHz, we obtained a jitter of 300-400 fs. This results in a jitter-limited performance of about 10 bits at the Nyquist frequency of 500 MHz.

5.2.3 Single Frequency Digital Compensation

In this section, the data correspond to models with parameters which are fit separately for each input frequency. The models here do not allow for frequency dependent distortion.

5.2.3.1 SINAD

Fig. 5.9 shows the SINAD performance under the gain and d1_g1_c1 distortion compensation models. The low frequency performance of the d1_g1_c1 model is about 10 bits. For comparison to the other single frequency models we investigated, the low frequency performance of each model is given in Table 5.1 and the high frequency performance in Table 5.2.

5.2.3.2 SFDR

The SFDR performance of the ADC using single-frequency distortion compensation is shown in Fig. 5.10. The SFDR is limited by the second harmonic of the input frequency. This can be seen in Fig. 5.11, which shows the output spectrum of a full scale input at 10 MHz. The ADC uses a fully-differential design which should greatly reduce even-order harmonics. The balun at the ADC input is a possible source of the second harmonic distortion.

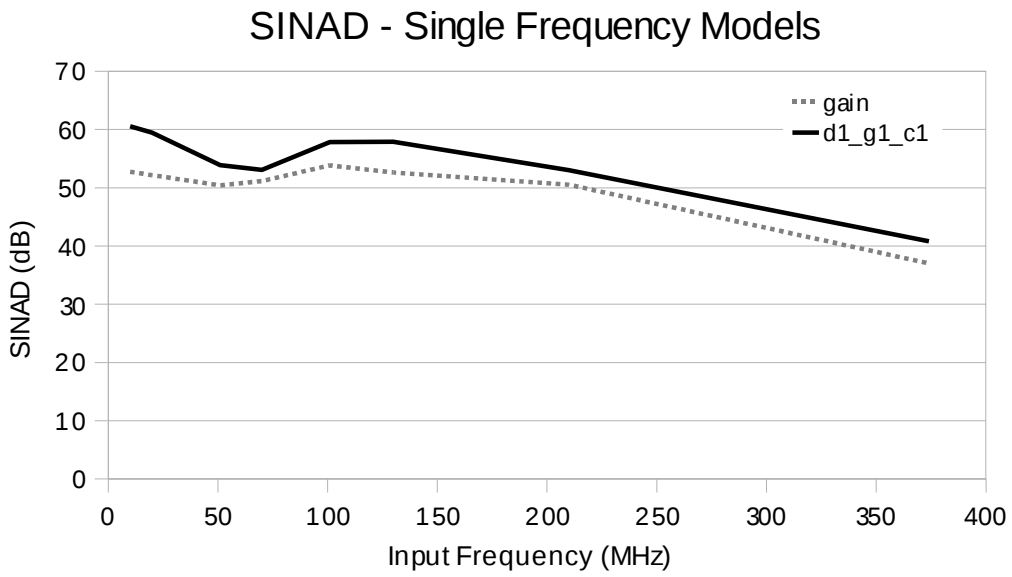

Figure 5.9 SINAD for single frequency models fit individually. $F_{clk} = 100\,\mathrm{MHz}$

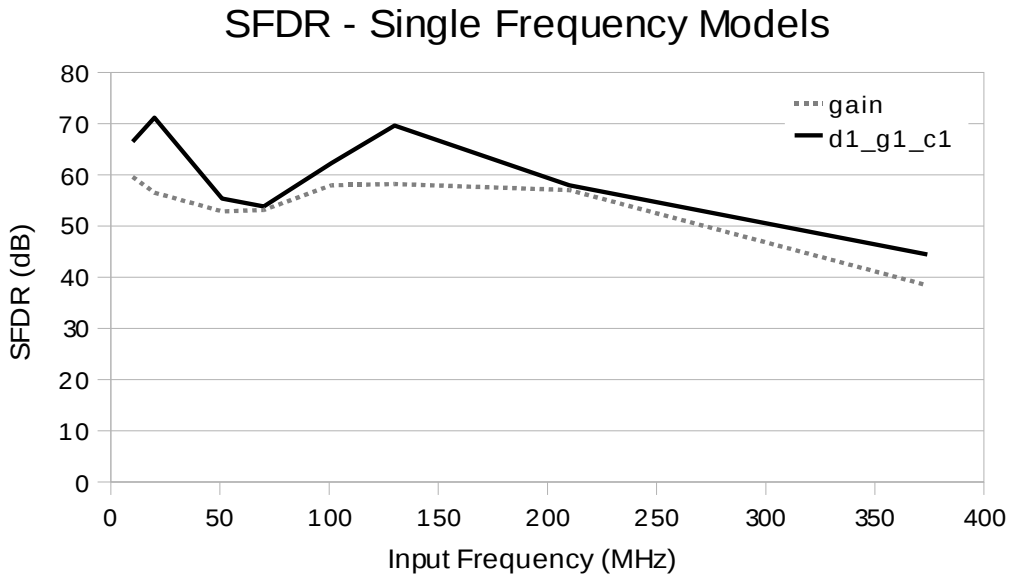

Figure 5.10 SFDR for single frequency models fit individually. $F_{clk} = 100\,\mathrm{MHz}$

Name	# params.	SINAD for slope:		
		both	pos.	neg.
gain	3	54.54	54.24	54.85
d1	19	57.45	57.61	57.70
d1_g1	36	60.37	60.95	60.89
d1_g1_c1	50	61.27	62.21	61.77
d1_nl	23	60.38	61.30	60.43
n33	11	57.92	57.63	58.68
d1_n33	21	57.47	57.67	57.72
d1_d2	37	58.45	59.51	60.11
d1_d2_g1_g2	65	60.82	61.44	61.53
d12	258	61.77	62.48	62.65
d12_g12	513	62.39	63.21	63.53

Table 5.1 Comparison of single-frequency models for a low frequency input, ($F_{in} = 10.01\,\mathrm{MHz}, F_{clk} = 100\,\mathrm{MHz}$)

Name	# params.	SINAD for slope:		
		both	pos.	neg.
gain	3	50.52	50.25	50.80
d1	19	51.75	56.90	56.21
d1_g1	36	52.78	58.88	57.06
d1_g1_c1	50	52.90	59.33	57.28
d1_nl	23	52.70	58.33	56.99
n33	11	51.92	57.66	55.50
d1_n33	21	52.15	57.56	56.46
d1_d2	37	52.27	58.72	56.58
d1_d2_g1_g2	65	52.90	59.52	57.36
d12	258	53.55	59.92	57.89
d12_g12	513	54.44	60.70	58.36

Table 5.2 Comparison of single-frequency models for a high frequency input, ($F_{in} = 210.01\,\mathrm{MHz}, F_{clk} = 100\,\mathrm{MHz}$)

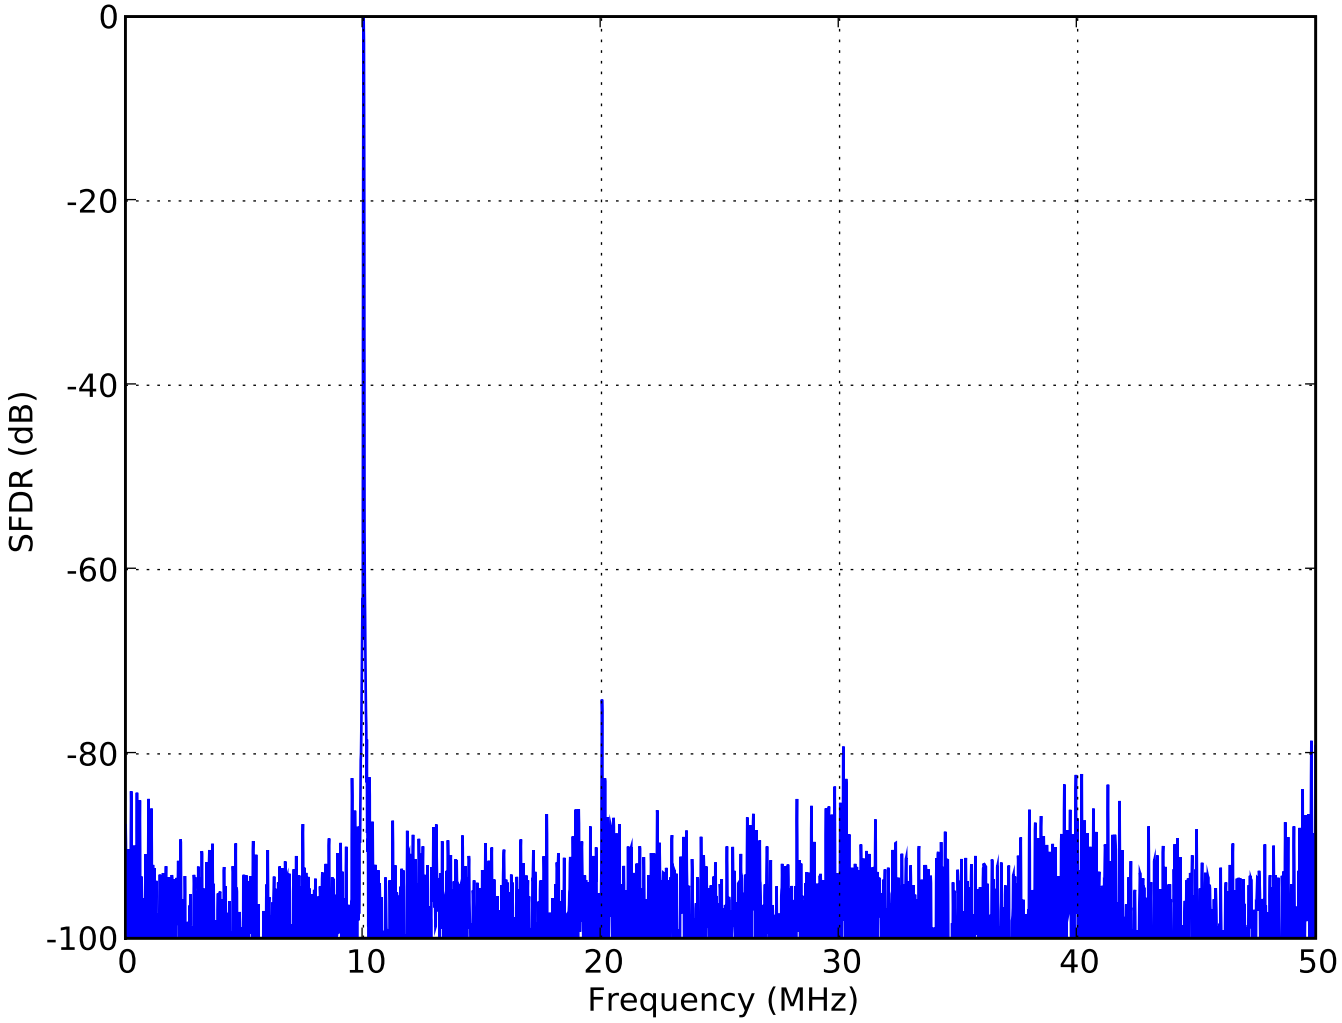

Figure 5.11 Output Spectrum for Full-Scale 10 MHz Input

5.2.4 Broadband Compensation

In this section, the data correspond to models which are intended to compensate distortion over a wide input bandwidth. The model parameters are obtained from a simultaneous fit to data records of varied frequency and amplitude.

The fact that distortion is frequency dependent is made plain by Fig. 5.12. This plot shows the ADC output error under the gain model for sinusoids inputs of 10, 20, 50, 70, 100, 130, 210 and 374 MHz and 40, 60, 80 and 100 percent of the full range input amplitude. The error waveforms are intentionally offset for visibility.

Note that the error records at 100% amplitude and low frequency should approach the ADC's integrated nonlinearity (INL) at DC. Thus, we need not try to measure the INL directly by applying a DC input, which would require bypassing the input balun.

The ADC broadband errors for the d1_g1_c1_ds model are shown in Fig. 5.13. This is the most successful distortion model, but also the most complex. The data recordings at 374 MHz, tend to compromise the other high frequency points, so these are removed from the parameter fitting from this point on.

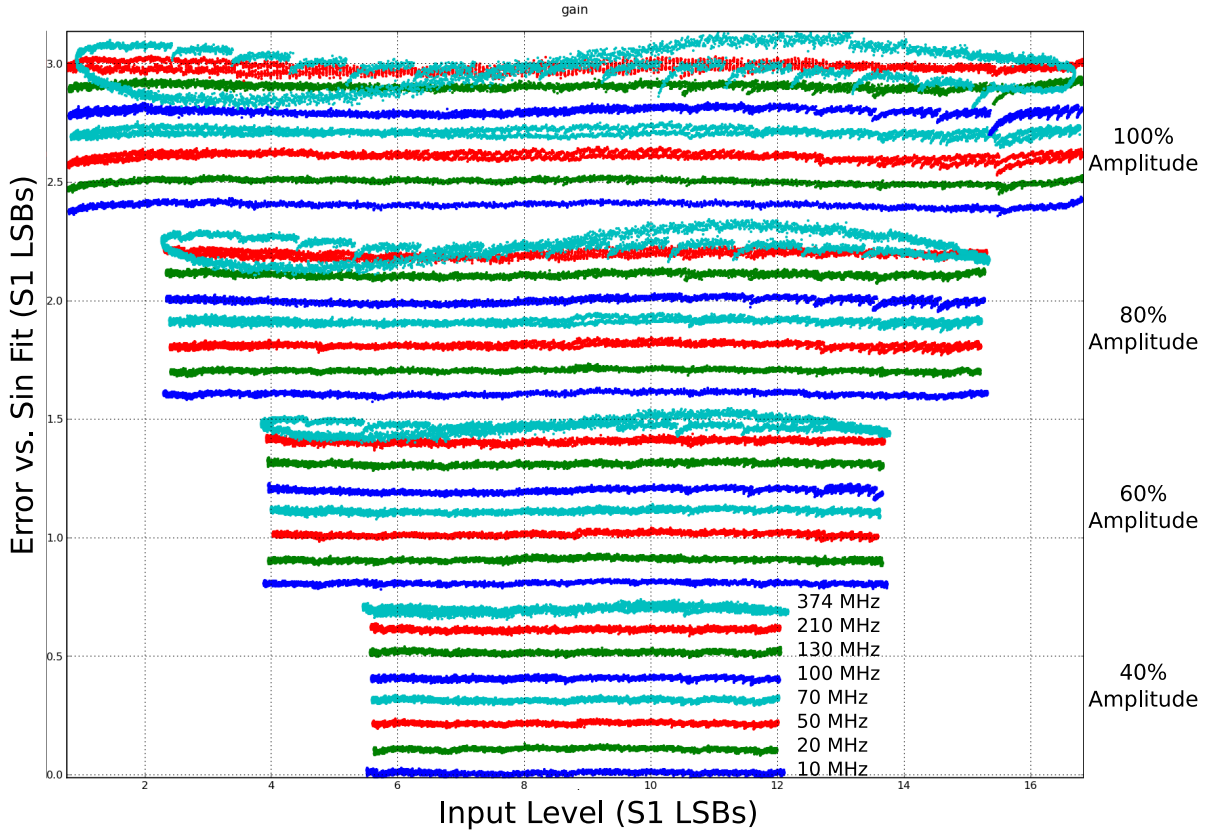

Figure 5.12 Broadband Fit Errors: Gain Fit

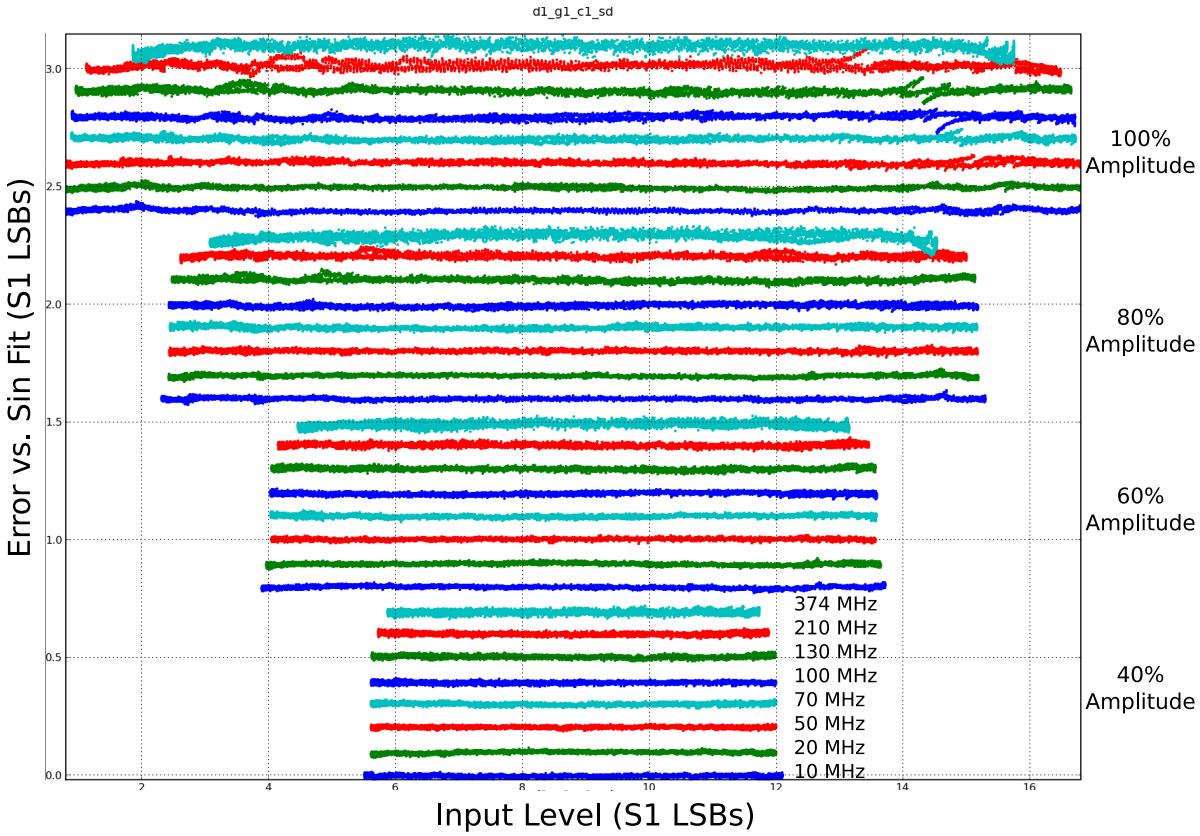

Figure 5.13 Broadband Fit Errors: "d1_g1_c1_sd" Discontinuous Slope Compensation

5.2.4.1 SINAD

Table 5.3 lists the SINAD levels of the broadband distortion compensation models. The SINAD is listed for various input amplitudes and frequencies. When calculating SINAD here, we use the amplitude of a full range sinusoid as the signal level, instead of the actual signal amplitude. This makes it easier to compare signals at different amplitudes.

Model	Amp.	70 MHz	101 MHz	130 MHz	210 MHz
gain	40%	54.89	58.65	53.63	56.10
	60%	54.51	55.79	54.59	55.68
	80%	52.21	51.98	54.26	50.37
	100%	50.91	51.66	52.47	46.72
d1_g1_c1	40%	57.47	58.58	55.80	58.24
	60%	56.31	57.88	56.51	57.56
	80%	54.46	55.55	56.16	52.55
	100%	51.99	53.87	54.57	47.44
d1_g1_c1_s	40%	61.17	55.81	58.57	59.21
	60%	58.72	55.73	57.97	57.96
	80%	56.45	54.88	56.23	57.43
	100%	54.38	54.36	53.81	56.05
d1_g1_c1_sd	40%	60.85	56.08	58.49	59.58
	60%	59.52	55.77	58.14	58.27
	80%	58.67	54.89	56.87	58.07
	100%	56.13	54.32	55.25	57.04

Table 5.3 SINAD for Broadband Models

Note that SINAD levels become more consistent across amplitude and frequency as more slope dependence is added. This shows that distortion models must allow frequency dependence in order to work with broadband inputs.

5.2.4.2 SFDR

The SFDR performance of the broadband models is shown in Table 5.4. Again, the distortion is improved and made more consistent with frequency and amplitude by the models with higher-order slope dependence.

Model	Amp.	70 MHz	101 MHz	130 MHz	210 MHz
gain	40%	61.73	56.59	61.39	62.77
	60%	54.57	54.81	57.90	58.90
	80%	55.81	56.93	58.15	55.14
	100%	53.07	57.95	58.27	56.93
d1_g1_c1	40%	65.88	56.58	66.10	65.51
	60%	57.72	57.30	65.08	61.40
	80%	55.76	61.99	63.82	58.21
	100%	53.16	61.30	60.58	55.72
d1_g1_c1_s	40%	68.28	56.82	67.05	63.26
	60%	59.33	57.60	62.34	65.64
	80%	57.93	63.63	62.84	65.71
	100%	56.03	64.28	59.27	67.54
d1_g1_c1_sd	40%	70.06	56.82	65.88	63.62
	60%	65.48	57.56	63.66	64.12
	80%	67.00	62.17	64.31	66.42
	100%	61.04	63.35	65.40	73.56

Table 5.4 SFDR for Broadband Models

5.3 Performance Comparison

Walden [10] plotted the performance of several commercial and prototype ADCs, which is reproduced in Fig. 5.14. The performance of the BDMS-ADC is plotted as an **x**. With minor layout modification, we should be able to increase the clock frequency of this ADC up to 500 MSps, with similar noise and distortion performance. This is indicated by an **o** added to the plot.

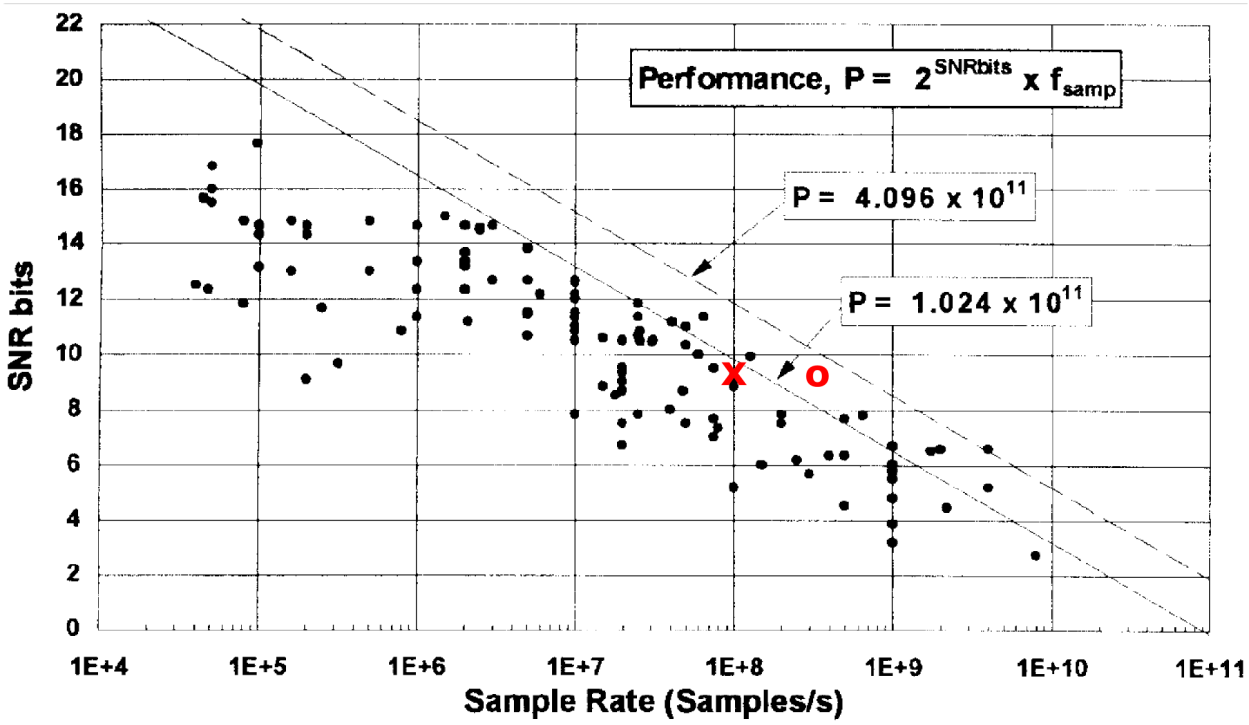

Figure 5.14 Comparison to Other High-Performance ADCs [10]

Chapter 6

Conclusion

The Bipolar-Driven MOS Switch THA developed in this research has important benefits to high-speed ADCs. First, it allows operation with relatively low power supply voltage range. Second, since the main switch acts like a resistor, the added noise is minimal. Also, the switches act as part of a low-pass filter to further limit noise. Finally, since the sampling time jitter is limited by bipolar switches, the jitter performance is much improved over what can be achieved using only MOS devices.

In the design of an ADC using the BDMS-THA, we did only minimal distortion compensation with analog circuitry. The only remaining distortion compensation techniques in the analog-domain are signal-tracking at the MOS gates of the THA in track mode and digital trim for the 4-bit flash ADC in each pipeline stage. As much compensation as possible was moved into the digital domain. The compensation was applied by computer processing after the records were acquired. This allowed us to fabricate the chip before developing the compensation algorithms, and allowed us to evaluate the performance of several compensation schemes.

Using the distortion compensation models we developed, the BDMS-ADC shows a SINAD performance 55 to 58 dB over an input bandwidth of 200 MHz, corresponding to around 9.5 effective bits. This makes the ADC competitive with ADC designs using GaAs technology. However, since the BDMS-ADC was designed with a SiGe process, it has a much lower fabrication cost.

6.1 Future Work

For the best distortion correction performance, we found it necessary to allow a slope dependence to the fitting parameters. For the work presented, the slope information was found by assuming the input to be sinusoidal in shape. In order for the ADC to cover a wide bandwidth, a more local slope estimate is required. Hummels [21] implemented a slope estimation using a 23-point FIR filter that adequately estimated the slope up to 80% of the Nyquist frequency. However, it would be more efficient to have a separate ADC for recording slope information. Having slope information recorded separately would also allow a more broadband compensation and allow for sub-sampling.

List of References

[1] R. Yu et al. A 1 GS/s 11-b track-and-hold amplifier with < 0.1 dB Gain Loss. In *GaAs IC Symposium*, pages 87–90, 1997.

[2] Thorsten Baumheinrich, Bernd Pregardier, and Ulrich Langmann. A 1-GSample/s 10-b full nyquist silicon bipolar track&hold IC. *IEEE J. Solid-State Circuits*, 32(12):1951–60, December 1997.

[3] C. D. Motchenbacher and J. A. Connelly. *Low noise electronic system design*. J. Wiley & Sons, 1993.

[4] Dirk J. Robinson and George S. La Rue. 1 GSps 11-bit track-and-hold in SiGe BiCMOS. In *Workshop on Microelectronics and Electron Devices*, pages 67–70. IEEE, 2005.

[5] Klass Bult. Analog design in deep sub-micron CMOS. *Solid-State Circuits Conference, 2000. ESSCIRC '00. Proceedings of the 26th European*, pages 126–32, 2000.

[6] Mayank Garg, Sushant S. Suryagandh, and Jason C.S. Woo. Scaling impact on analog performance of sub-100nm MOSFETs for mixed mode applications. *European Solid-State Device Research, 2003. ESSDERC '03. 33rd Conference on*, pages 371–74, 2003.

[7] Georges Gielen et al. Analog and digital circuit design in 65nm CMOS: end of the road? *Design, Automation and Test in Europe, 2005. Proceedings*, pages 37–42, 2005.

[8] Erix A. Vittoz. Future of analog in the VLSI environment. *Circuits and Systems, 1990., IEEE International Symposium on*, pages 1372–75, 1990.

[9] David A. Johns and Ken Martin. *Analog Integrated Circuit Design*. Wiley, 1997.

[10] Robert H. Walden. Analog-to-digital converter survey and analysis. *IEEE J. Select. Areas Commun.*, 17(4):539–50, April 1999.

[11] P. Vorenkamp and R. Roobers. A 12b 60 MSample/s cascaded folding and interpolating ADC. *IEEE J. Solid-State Circuits*, 32(12):1876–95, December 1997.

[12] Alex R. Bugeja and Sung-Ung Kwak. Design of a 14b 100MS/s switched-capacitor pipelined ADC in RFSiGe BiCMOS. In *ISCAS*, pages 428–31, 2001.

[13] Sang-Min Yoo et al. A 2.0V 12b 120 MSample/s CMOS pipelined ADC. In *ISCAS, 2006*, page CDROM. IEEE, 2006.

[14] Eulalia Balestrieri, Pasquale Daponte, and Sergio Rapuano. A state of the art on adc error compensation methods. *IEEE Trans. Instrum. Meas.*, 54(4):1388–94, August 2005.

[15] J. Tsimbinos and K. V. Lever. Improved error-table compensation of A/D converters. *IEE Proc.-Circuits Devices Syst.*, 144(6):343–49, December 1997.

[16] J. Larrabee, F. H. Irons, and D. M. Hummels. Using sine wave histograms to estimate analog-to-digital converter dynamic error functions. *IEEE Trans. Instrum. Meas.*, 47(6):1448–56, December 1998.

[17] Adam Ginsburg. SciPy cookbook / FittingData. http://www.scipy.org/Cookbook/FittingData.

[18] Alex Martelli, Anna Martelli Ravenscroft, and David Ascher. *Python Cookbook*. O'Reilly, 2 edition, 2005.

[19] Noel Rappin and Robin Dunn. *wxPython in Action*. Manning, 2006.

[20] Alex Martelli. *Python in a Nutshell: A Desktop Quick Reference*. O'Reilly, 2 edition, 2006.

[21] D.M. Hummels, R.W. Cook, and F.H. Irons. Discrete-time dynamic compensation of analog-to-digital converters. *Circuits and Systems, 1993., ISCAS '93, 1993 IEEE International Symposium on*, pages 1144–47 vol.2, May 1993.

[22] C. Fiocchi, U. Gatti, and F. Maloberti. Design issues on high-speed high-resolution track-and-holds in BiCMOS technology. *IEE Proc.-Circuits and Devices Syst.*, 147(2):100–06, April 2000.

[23] Philip W. Yee. Noise considerations in high-accuracy A/D converters. *IEEE J. Solid-State Circuits*, 21(6):1011–15, December 1986.

[24] Peter R. Kinget. Device mismatch and tradeoffs in the design of analog circuits. *IEEE Journal of Solid-State Circuits*, 40(6):1212–24, June 2005.

[25] Koen Uyttenhove and Michel S. J. Steyaert. Speed-power-accuracy tradeoff in high-speed CMOS ADCs. *IEEE Transactions on Circuits and Systems - II*, 49(4):280–87, April 2002.

[26] Jan M. Rabaey, Anantha Chandrakasan, and Borivoje Nikolic. *Digital Integrated Circuits*. Prentice Hall, 2nd edition, 2003.