A Continuous Time Sigma-Delta Modulator for Digitizing Carrier Band Measurements

by

Philip Weimin Juang

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2001

Author .....
Department of Electrical Engineering and Computer Science
May 11, 2001

Certified by.....
James K. Roberge
Professor of Electrical Engineering
Thesis Supervisor

Certified by.....
Robert Bousquet
C.S. Draper Laboratory Senior Engineer
Thesis Supervisor

Accepted by...
Arthur C. Smith
Chairman, Department Committee on Graduate Theses

A Continuous Time Sigma-Delta Modulator for Digitizing Carrier Band Measurements

by

Philip Weimin Juang

Submitted to the Department of Electrical Engineering and Computer Science
on May 11, 2001, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

This document discusses the design of a low power, low area, sigma-delta modulator to be built on a .5μm CMOS process. First, the use of a continuous-time topology is investigated and compared with an existing discrete-time design. Next, the design methodology of a 3rd-order continuous-time modulator is explored; the same methodology is then applied to a 4th-order modulator design. Both designs were fabricated and tested in the lab. The final section discusses the test results from the completed IC's and comments on the viability of continuous-time modulators for this application.

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

Thesis Supervisor: Robert Bousquet
Title: C.S. Draper Laboratory Senior Engineer

Acknowledgments

First and foremost, I'd like to thank my thesis supervisors, Professor James K. Roberge and Rob Bousquet, as well as all of my friends and colleagues at Draper Laboratory: Paul Ward, Eric Hildebrant, Bill Kelley, Tom King, Shida Martinez, Chris O'Brien, and Rob Bousquet. They have all made my career at MIT and Draper fun and worthwhile. My understanding of electrical engineering in general and solid-state circuits in particular has benefited from their endless patience and invaluable counsel. I am tremendously grateful for their contributions to my professional growth.

Second, I'd like to acknowledge my friends, who have stood by me throughout my years at MIT and have provided welcome relief after the long days of thesis writing. Thanks go out to Stimpe, Steve, Keith, Jon, Frank, Morris, Spencer, Lahaie, Ryan, Amar, and of course, Christine.

Lastly, I'd like to extend thanks to my parents for their love and support. I am in their debt, as without their patience and wisdom, I might have become a chemical engineer. I hope they enjoy reading the next 100 pages.

This thesis was prepared at The Charles Stark Draper Laboratory, Inc., under Internal Company Sponsored Research No. C312.

Philip Juang

Contents

1 Introduction **17**

1.1 Background 17

1.2 Outline 18

2 Fundamental Operation of Σ-Δ Modulators for Analog to Digital Conversion **21**

2.1 Characteristics of Σ-Δ Converters 21

2.1.1 Oversampling Converters 22

2.1.2 Noise Shaping 26

2.2 A Discrete-Time Third Order Σ-Δ Modulator 31

2.2.1 Basic Modulator Operation 32

2.2.2 Circuit Realization 33

3 Advantages and Disadvantages of Continuous-Time Σ-Δ Modulators **39**

3.1 Benefits 39

3.1.1 Reducing power consumption 40

3.1.2 Decreasing chip area 41

3.1.3 Lowering thermal noise 43

3.2 Drawbacks 46

3.2.1 Absolute component tolerances 46

3.2.2 Sample-and-hold circuit 47

3.2.3 Return-to-zero DAC 47

3.2.4 Clock Jitter 48

4 Design of a third-order continuous-time Σ-Δ modulator **53**

4.1 Design possibilities 53

4.1.1 Direct impulse-invariance transformation 53

4.1.2 Tunable on-chip components 55

4.1.3 Continuous-time front end 55

4.1.4 Lowpass modulator topology 57

4.2 Design Methodology 58

4.2.1 Initial hand design 58

4.2.2 Software simulation design 64

4.2.3 1-bit Return-to-zero Digital-to-Analog Converter 69

4.3 Design Review 70

4.3.1 Simulated Performance 70

4.3.2 Weaknesses 75

5 Design of a fourth-order continuous-time Σ-Δ modulator **77**

5.1 Techniques to decrease quantization noise 77

5.1.1 Increasing sampling frequency 77

5.1.2 Fourth-order topology 78

5.2 Design methodology 79

5.2.1 Initial hand design 79

5.2.2 Minimizing power consumption 81

5.2.3 RTZ DAC redesign 82

5.3 Design Review 83

5.3.1 Simulated Performance 83

5.3.2 Analysis 87

6 Measurement Results of On-chip Continuous Time Σ-Δ Modulators **89**

6.1 Measured performance 89

6.1.1 SD3CT measurements 89

6.1.2 SD4CT measurements 91

6.2 Effects of Clock Jitter 94

6.2.1 Phase noise modulation 95
6.2.2 Clock jitter tolerances for SD3CT and SD4CT 96
6.3 Future Work 98

7 Conclusion **101**

A Circuit Schematics of SD3CT and SD4CT **103**

B SD3CT Simulation and Lab Test Data **109**

C SD4CT Simulation and Lab Test Data **119**

D Clock Jitter Test Results **129**

List of Figures

1-1 Power spectrum of analog input signal 18

2-1 A 1-bit oversampling converter 22

2-2 Spectrum of $x[n]$, sampled from continuous-time signal $x(t)$ 23

2-3 White noise model of quantizer 23

2-4 Spectral Density of $e[n]$ 24

2-5 Spectral Density of r[n] 25

2-6 Brick wall filter used to attenuate the out-of-band noise 25

2-7 Spectral Density of y[n]. All the out-of-band noise has been filtered out. 26

2-8 A 1-bit oversampling converter with noise shaping 27

2-9 Noise shaping loop with the white noise model of the quantizer 27

2-10 Frequency Response of $H(e^{j\omega})$ (single integrator) represented in log scale. 28

2-11 Signal Transfer Function: $\frac{P(e^{j\omega})}{X(e^{j\omega})}$ represented in log scale. 29

2-12 Noise Transfer Function: $\frac{P(e^{j\omega})}{E(e^{j\omega})}$, represented in log scale. 29

2-13 Noise Spectrum: Spectrum of $P(e^{j\omega})$ due to $E(e^{j\omega})$. The dotted line represents the spectral height k_e of the error signal. The solid line depicts the error signal after noise shaping. 30

2-14 Noise Spectrum: Spectrum of $P(e^{j\omega})$ due to $E(e^{j\omega})$ after filtering. The two components are the same as in Figure 2-13. Note that the spectral density of the shaped noise is significantly reduced. 31

2-15 Block diagram of third order $\Sigma\Delta$ circuit. 32

2-16 Frequency Response of the Noise Transfer Function with $f_s = 5.12MHz$. 33

2-17 A fully differential switched-capacitor integrator 34

2-18 1-bit feedback DAC as it is implemented in the existing third order modulator design 37

3-1 A simple continuous-time integrator (fully-differential) 40
3-2 Block diagram of third order ΣΔ circuit. 43
3-3 Two waveforms with 7 "ones." Due to rise and fall times, the energy (area under the curve) in the first waveform is slightly different than the energy in the second. 48
3-4 Waveform of a Return-to-Zero DAC. 49
3-5 Feedback current waveforms for a single sampling period 49
3-6 Clock jitter effects on the current waveforms of two different DAC implementations. 51

4-1 Change in notch frequency with a 25% variance in $c4$ 55
4-2 Block diagram of third order ΣΔ circuit. 56
4-3 Block diagram of a third order lowpass ΣΔ modulator. 59
4-4 Frequency Response of the NTF with all poles located at unity-gain frequency 1.28MHz 60
4-5 Bode plot of hand designed loop gain 62
4-6 Magnitude response of the NTF corresponding to Figure 4-5 63
4-7 Root locus plot of the loop transmission 63
4-8 Circuit simulation of lowpass third order modulator 64
4-9 Output of integrator 2 using initial hand design 65
4-10 Output of integrator 2 with decreased feedback coefficient b_2 66
4-11 Block diagram of third order lowpass ΣΔ modulator with feedforward paths. 67
4-12 Output of integrator 2 with decreased feedback coefficient b_2 and feedforward path from $x[n]$ 68
4-13 Circuit diagram of a monostable multivibrator (one-shot) 69
4-14 Transmission gates that switch the feedback node from NREF, MID, and PREF 70

4-15 Positive and negative return-to-zero pulses generated by the DAC . . 71
4-16 Plot of SD3CT Noise Transfer Function 73
4-17 Power spectrum of SD3CT output, normalized to P_s=1V^2 74

5-1 Bode plot of the hand designed fourth-order loop gain 80
5-2 Frequency Response of the NTF corresponding to Figure 5-1 80
5-3 Root locus plot of the fourth order loop transmission 81
5-4 Circuit implementation of the redesigned DAC 83
5-5 Return-to-zero waveforms for the redesigned DAC 84
5-6 Magnitude response of the SD4CT noise transfer function 85
5-7 Power spectrum of the SD4CT output, normalized to P_s=1V^2 86

6-1 Nominal output spectrum of SD4CT 92
6-2 Nominal output spectrum of SD4CT (zoomed out) 93
6-3 Output spectrum of SD4CT using a clock with white jitter (bandlimited to 3kHz) 95

List of Tables

1.1 Specifications for A/D conversion of input signal 19

6.1 Measured performance for A/D conversion using SD3CT 90
6.2 Measured performance for A/D conversion using SD4CT 91
6.3 Relationship between clock jitter and output noise floor for SD3CT . 97
6.4 Relationship between clock jitter and output noise floor for SD4CT . 97

Chapter 1

Introduction

1.1 Background

The Microelectronics Group at Charles Stark Draper Laboratories has a need to digitize an analog waveform for digital processing. The signal of interest is bandlimited to 100Hz and modulated on a sinusoid whose frequency is $2f_c$ (nominally 20kHz). Figure 1.1 approximates the power spectrum of the analog input; the signal of interest is located about the second harmonic. While this signal can be extracted and processed using analog circuitry, significant error can be introduced by DC bias drift. It has been shown that digital processing of the signal results in the elimination of this problem.

A sigma-delta (Σ-Δ) A/D converter is suited to the task of digitizing the signal, due to its ability to achieve high-accuracy data conversions in a relatively short period of time. Furthermore, since the linearity of the conversion is important to this application, Σ-Δ converters are well suited to this job A 3rd-order discrete-time (DT) Σ-Δ modulator has been designed and fabricated on a CMOS process to perform this task. The device is powered by a single 5V supply—all other voltages needed are generated on chip. The input signal is centered about a virtual ground of 2.5V (which we designate MID). Table 1.1 displays the major specifications of a Σ-Δ modulator, as well as the measured results obtained from a discrete-time design [10].

The first thing we notice about this design is that the SNR is too low. While

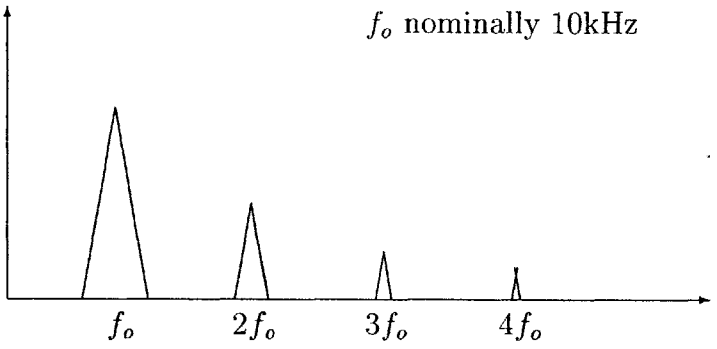

Figure 1-1: Power spectrum of analog input signal

this is somewhat problematic, it has been found that this Σ-Δ modulator functions adequately; the specification is somewhat conservative. Since all other specifications appear to have been met, it will be used as intended. However, we foresee a need for a smaller, lower power A/D converter that can achieve the specified noise level in future applications. For this reason, we are interested in investigating new Σ-Δ designs; of particular interest to us is a continuous-time topology. It is believed that such a design will be useful for future ASICs.

1.2 Outline

Usually, Σ-Δ modulators are designed using a DT topology, employing the use of switched-capacitor filters when implemented. The goal of this thesis is to explore the use of a continuous-time (CT) implementation for our specific need. We begin by discussing the operation of Σ-Δ converters, the principle of oversampling, and the techniques involved in noise shaping. We will also highlight the major advantages and limitations of CT Σ-Δ modulators and explore its use for digitizing our carrier band measurement. As we shall see, using a CT Σ-Δ modulator is extremely beneficial in terms of reducing power consumption and minimizing circuit noise.

The next part of this document will focus on the design of a 3rd-order CT modulator. We will see that the most difficult tasks are achieving sufficient quantization

Characteristic	Specification	Result
Clock Frequency	$256^* f_o$	—
Full Scale SNR within 20kHz$\pm$100Hz	117 dB	111 dB
Input Range	$> \pm$ 1V Full scale	$\pm$ 2.5V FS
Second Harmonic Distortion (from 10kHz input)	< -80dBcFS	-82dBcFS
AC Gain Stability at 20kHz	1% over temperature	0.4% over temperature
Spurious Tones	< -80dBcFS at 20kHz$\pm$4kHz	No tones in output spectrum
Power Consumption	—	40mW
Chip area	—	$9mm^2$

Table 1.1: Specifications for A/D conversion of input signal

noise shaping, eliminating the factors changing the scale factor of our modulator (i.e. maintaining AC gain stability), and minimizing the effects of clock jitter, which injects undesired noise into the system. A 3rd-order modulator is found to be insufficient for our needs; hence, we then discuss a 4th-order design while addressing issues uncovered in the 3rd-order modulator.

Finally, we compare our expected results to those obtained from lab testing of the manufactured silicon test chips. The last section discusses the discrepancies discovered, studies and documents the effects of clock jitter on our modulator, and suggests future work to be done on this topic.

Chapter 2

Fundamental Operation of Σ-Δ Modulators for Analog to Digital Conversion

2.1 Characteristics of Σ-Δ Converters

Σ-Δ converters are generally defined to be the class of converters which use shaped quantization noise with oversampling to achieve high resolution. Σ-Δ conversion has become popular for many reasons [2]:

- It yields high accuracy results fairly quickly; thus, it is useful for high-resolution, medium speed applications.
- It is easily realized on a CMOS process. Often, Σ-Δ converters rely on well-matched components, which are provided on an integrated circuit.
- A/D converters are usually preceded by an anti-aliasing filter. Σ-Δ converters reduce the required performance of these filters.
- The static power dissipation can be made very low without significantly degrading the performance of the converter.

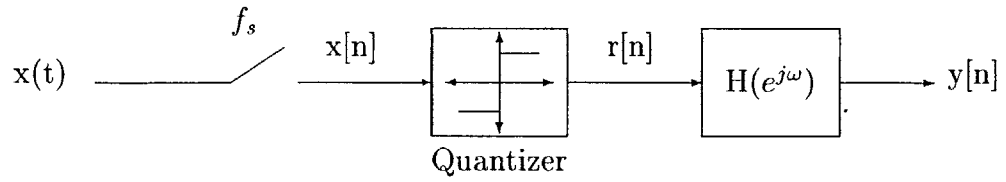

Figure 2-1: A 1-bit oversampling converter

Σ-Δ conversion is done in two stages. The first stage samples the analog signal and shapes the quantization noise. The second stage downsamples the bits obtained from the first stage and digitally filters the out-of-band noise. This decimator is realized using digital circuitry. This document is mainly concerned with the design of the first stage, often called the *modulator*. Design of the second stage, called the *decimator*, is beyond the scope of this thesis.

To fully understand the issues involved in designing Σ-Δ modulators, one must be familiar with the concepts of oversampling and noise shaping. This section will discuss these two techniques and illustrate them in a simple 1st-order modulator design.

2.1.1 Oversampling Converters

Oversampling is the method of sampling a signal above the Nyquist rate. It is a useful technique in A/D conversion, because it reduces the power of the quantization noise in the band of interest. To illustrate this effect, consider a 1-bit oversampling converter, whose block diagram is shown in Figure 2-1.

Let us suppose that x(t) is bandlimited to some frequency f_o. The input signal $x(t)$ is sampled at a frequency f_s to produce a digitized signal $x[n]$, whose power spectrum is shown in Figure 2-2. It is then put through a 1-bit quantizer to obtain an output of either $\frac{\Delta}{2}$ or $-\frac{\Delta}{2}$ (where Δ is equal to the difference in quantization levels). The filter $H(e^{j\omega})$ serves to bandlimit the resulting output signal $r[n]$ to some frequency bandwidth f_o. The behavior of the quantizer causes an error (a significantly large one, in the case of a 1-bit quantizer). It has been shown that this quantizer noise can

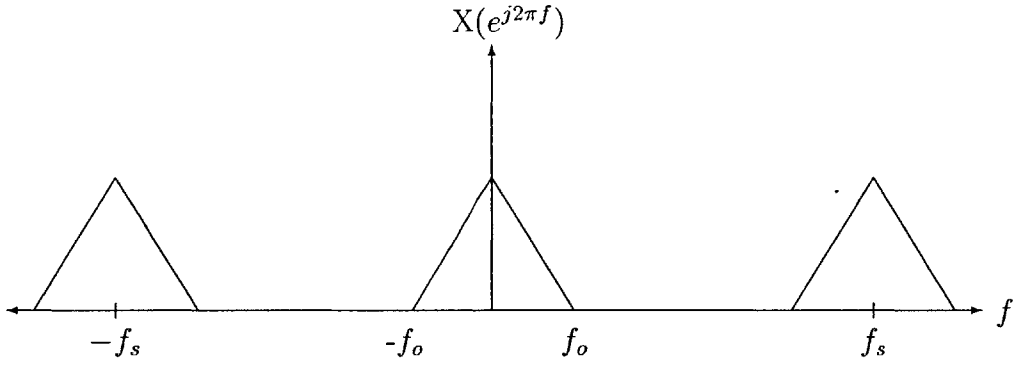

Figure 2-2: Spectrum of $x[n]$, sampled from continuous-time signal $x(t)$

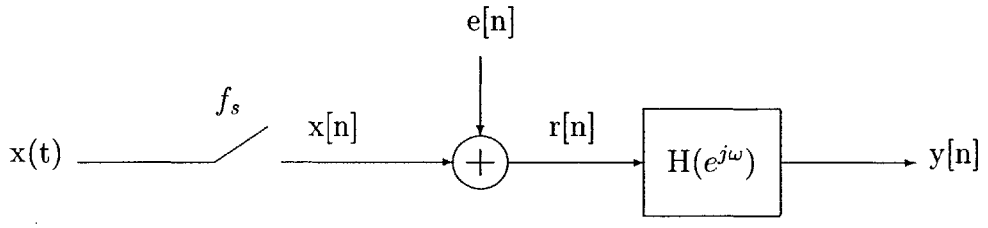

Figure 2-3: White noise model of quantizer

be modeled as an additive white noise source ($e[n]$), as shown in Figure 2-3. [7]

Intuitively, one can think of the quantization error sequence $e[n]$ as being an independent random variable uniformly distributed between $\pm\frac{\Delta}{2}$–since it is uniformly distributed, the spectrum will be flat. However, this approximation assumes that the quantization noise $e[n]$ is uncorrelated with the input signal $x[n]$. In reality, these two signals have a small correlation; however, for hand calculations we can ignore this correlation and assume white noise. Thus, this model provides a good basis for estimating the signal-to-quantization-noise ratio (SQNR) of the A/D converter. This ratio is defined as follows:

$$SQNR = \frac{P_s}{P_e} \tag{2.1}$$

where P_s is the signal power at the output and P_e is the quantization noise power at the output. We can easily hand calculate these power values. First, if the input is a

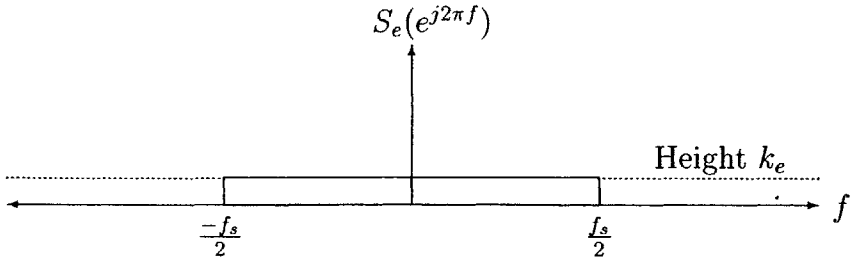

Figure 2-4: Spectral Density of $e[n]$

sinusoid, the signal power at the output of an N-bit quantizer is given by [2]:

$$P_s = \frac{\Delta^2 2^{2N}}{8} \tag{2.2}$$

For our example of a one-bit quantizer, N=1:

$$P_s = \frac{\Delta^2}{2} \tag{2.3}$$

P_e can be easily calculated as well. As previously mentioned, we can model the spectral density of $e[n]$ (which we designate $S_e(f)$ as white noise (Figure 2-4). The total quantization noise power is known to be $\frac{\Delta^2}{12}$ [2], so the spectral density height, k_e, of the error signal can be calculated to be:

$$\int_{-f_s/2}^{f_s/2} \left|S_e(f)^2\right| df = k_e^2 f_s = \frac{\Delta^2}{12} \tag{2.4}$$

$$k_e^2 = \frac{\Delta^2}{12}\frac{1}{f_s} \tag{2.5}$$

We have examined the spectrum of both $x[n]$ and $e[n]$, so we can determine what the spectral density of $r[n]$ looks like; we denote this spectrum to be $S_r(f)$, which is shown in Figure 2-5. The triangle is the component due to the signal $x[n]$ (which represents the signal power P_s). The flat box is the component due to the quantization noise $e[n]$ (which represents the signal power P_e and has spectral density height k_e). The lowpass filter $H(e^{j\omega})$ (shown in Figure 2-6) serves to attenuate the out-of-band

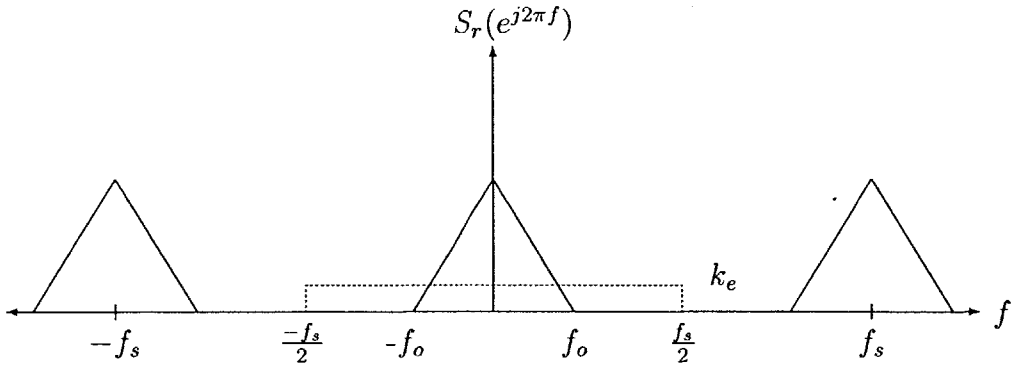

Figure 2-5: Spectral Density of r[n]

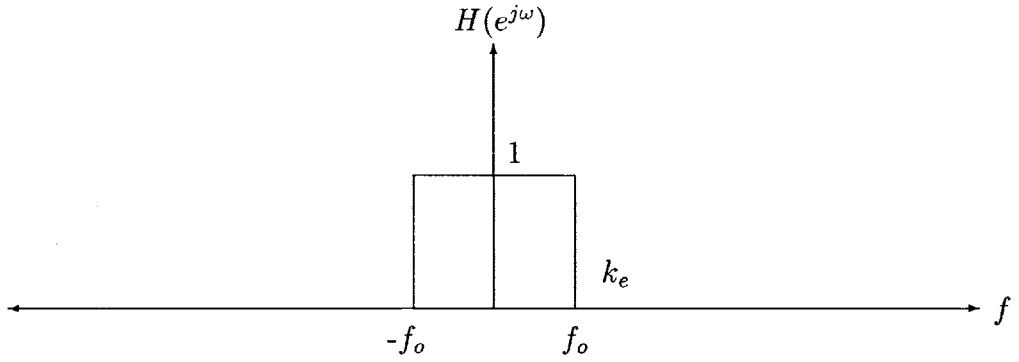

Figure 2-6: Brick wall filter used to attenuate the out-of-band noise

noise. Since the $x[n]$ is already bandlimited to f_o, the signal power remains unchanged. However, the out-of-band quantization noise is filtered out, so the quantization noise power in $y[n]$ is:

$$P_e = \int_{-f_s/2}^{f_s/2} S_e(f)^2 \left|H(e^{j\omega})\right|^2 df = \int_{-f_o}^{f_o} S_e(f)^2 df = \int_{-f_o}^{f_o} k_e^2 df = \frac{\Delta^2}{12}\frac{2f_o}{f_s} \tag{2.6}$$

We define the term $f_s/2f_o$ as the oversampling ratio (OSR). Equation 2.6 shows that the quantization noise power of $y[n]$ is inversely proportional to the oversampling ratio. Hence, the OSR and the SQNR have the following relationship:

$$SQNR = \frac{P_s}{P_e} = \frac{\Delta^2}{2} \times \frac{12}{\Delta^2}\frac{f_s}{2f_o} = 6 \times OSR \tag{2.7}$$

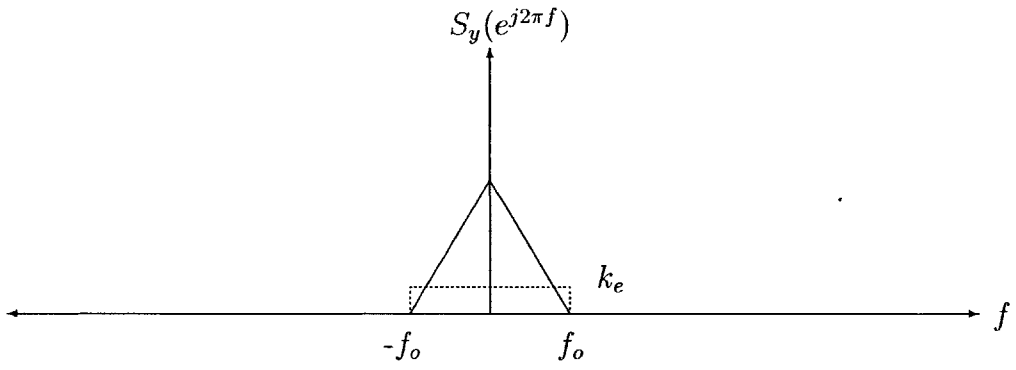

Figure 2-7: Spectral Density of y[n]. All the out-of-band noise has been filtered out.

$$SQNR(dB) = 10log(6) + 10log(OSR) \tag{2.8}$$

So for every doubling of the oversampling ratio, the SQNR improves by approximately 3dB.

The spectral density of $y[n]$ is depicted in Figure 2-7. It shows that through the use of oversampling, much of the noise power has been spread out across a wider band and then filtered out, thereby reducing the noise content in the output.

2.1.2 Noise Shaping

While oversampling is advantageous, it is usually not sufficient to obtain high-resolution A/D conversion. For example, in order for a 1-bit oversampling converter to achieve 70dB SQNR for a 20kHz bandlimited signal, a sampling frequency of 40GHz is required! One solution is to use multi-bit quantizers, which reduces the quantization error and thus the noise power. While achieving higher resolution with multi-bit quantizer architectures is possible, there are several drawbacks. First, the output of such an architecture would be a multi-bit word, instead of a single bit. This topology increases the complexity of the modulator; most notably, the D/A converter in the feedback path would require much more design work. Secondly, the FIR filter in the decimator would need to be able to handle multi-bit words, significantly adding to its complexity. Lastly, in realizing a multi-bit quantizer, there will be a higher degree

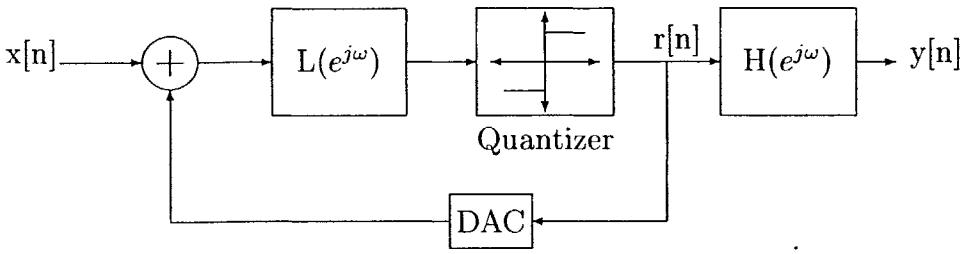

Figure 2-8: A 1-bit oversampling converter with noise shaping

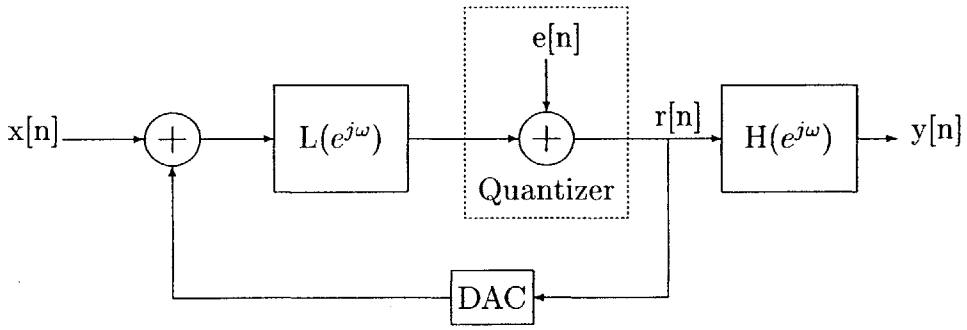

Figure 2-9: Noise shaping loop with the white noise model of the quantizer

of nonlinearity in the A/D conversion. The advantage of using a 1-bit converter is that it is inherently linear.

An alternative technique that is used to achieve a higher SQNR ratio is noise shaping. Noise shaping pushes the quantization noise out of band, so that it can be filtered out. The most common way of achieving this noise shaping is by feeding back the output of the quantizer to the sampled input, forming a feedback loop. The block diagram of this loop is shown in Figure 2-8. (Note: The D/A converter in the feedback path serves mainly to buffer the feedback signal and avoid loading the modulator output. In this case, it can be considered to be a unity gain block).

Using the white noise approximation for quantization noise (shown in Figure 2-9, it can be seen that the Signal Transfer Function (STF) and the Noise Transfer Function (NTF) are:

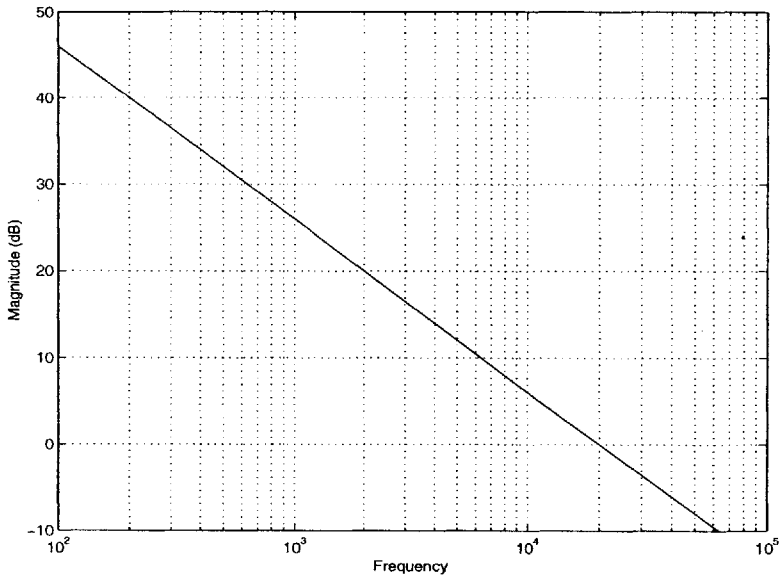

Figure 2-10: Frequency Response of $H(e^{j\omega})$ (single integrator) represented in log scale.

$$STF = \frac{R(e^{j\omega})}{X(e^{j\omega})} = \frac{L(e^{j\omega})}{1 + L(e^{j\omega})} \approx 1 \quad for\ L(e^{j\omega}) > 1 \tag{2.9}$$

$$NTF = \frac{R(e^{j\omega})}{E(e^{j\omega})} = \frac{1}{1 + L(e^{j\omega})} \tag{2.10}$$

By choosing the magnitude of $L(e^{j\omega})$ to be significantly greater than 1 in the band of interest, the input signal $x[n]$ will remain unaffected while the spectrum of the noise signal $e[n]$ will be shaped so that the inband noise is attenuated. For example, choosing $L(e^{j\omega})$ to be an integrator will have the desired effect—$L(e^{j\omega})$ has a large gain for low frequencies and small gain (less than unity) for high frequencies. This feedback structure, which provides the noise shaping, makes up a large part of Σ-Δ modulator design.

To illustrate this technique, let us assume that the input signal $x[n]$ is bandlimited to 10kHz, and $L(e^{j\omega})$ is designed to be a single integrator with a unity-gain frequency of 20kHz. The STF is approximately unity at frequencies below 10kHz. The NTF attenuates the quantization noise error at low frequencies. Figures 2-10 through 2-12 show the magnitude responses in log-log scale.

The power spectrum of the output signal $R(e^{j\omega})$ can be determined by multiplying

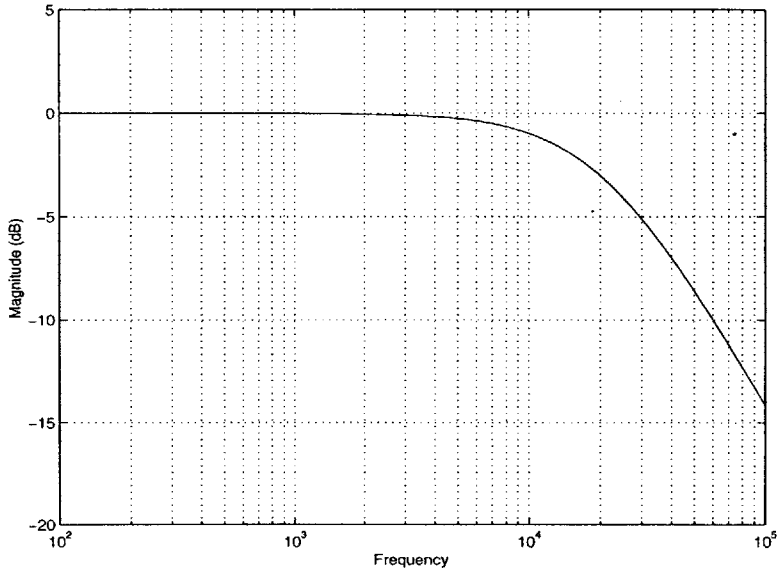

Figure 2-11: Signal Transfer Function: $\frac{P(e^{j\omega})}{X(e^{j\omega})}$ represented in log scale.

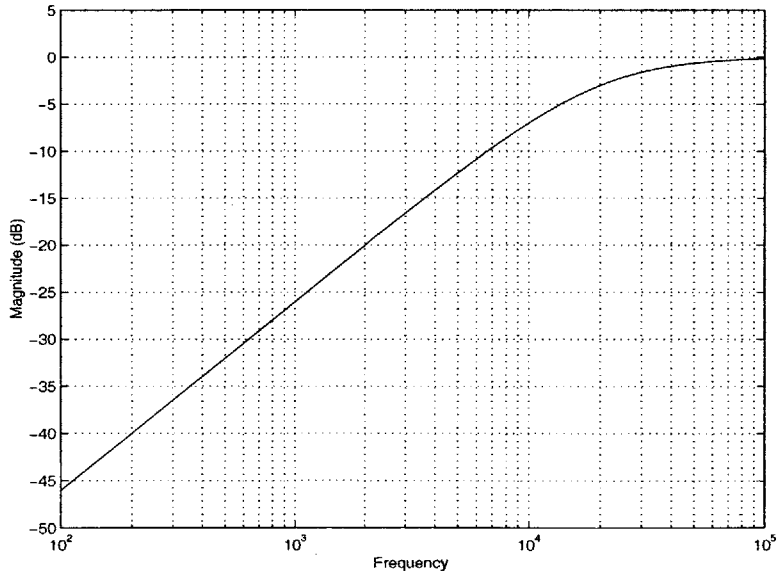

Figure 2-12: Noise Transfer Function: $\frac{P(e^{j\omega})}{E(e^{j\omega})}$, represented in log scale.

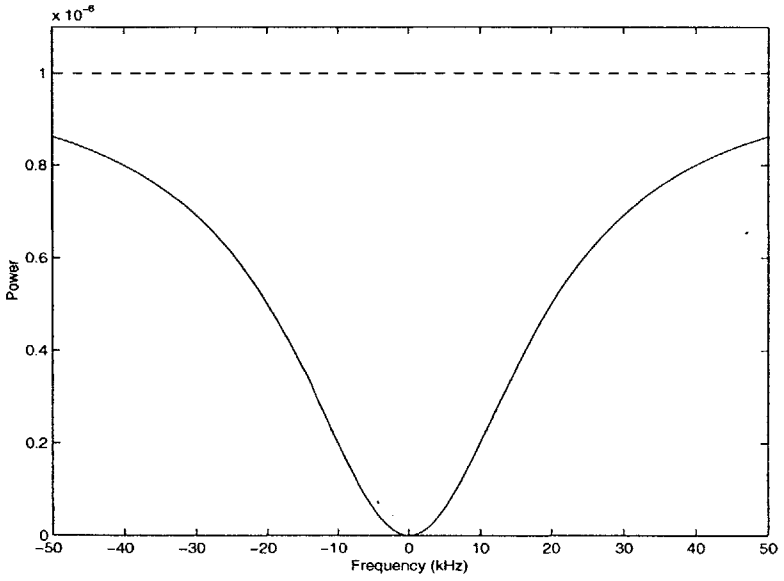

Figure 2-13: Noise Spectrum: Spectrum of $P(e^{j\omega})$ due to $E(e^{j\omega})$. The dotted line represents the spectral height k_e of the error signal. The solid line depicts the error signal after noise shaping.

the input signal's spectrum by the square of the magnitude of the NTF [7]. Hence, Equation 2.6 can be rewritten to include the effects of the noise shaper:

$$P_e = \int_{-f_s/2}^{f_s/2} S_e(f)^2 \left|NTF(f)H(e^{j2\pi f})\right|^2 df = \int_{-f_o}^{f_o} |NTF(f)|^2 k_e^2 df \qquad (2.11)$$

Figures 2-13 and 2-14 show the two-sided spectrum of the modulator output $P(e^{j\omega})$ due to the quantization noise error $E(e^{j\omega})$ before and after filtering with a 10kHz brick wall filter. (For these figures, the spectral density height k_e is arbitrarily chosen to be 1μV after oversampling.) It is easy to see that noise shaping has greatly reduced the noise power in the output; even for a first-order example, the noise power has been roughly reduced by a factor of 10 (20dB).

The characteristics of the noise shaper has a significant effect on the SQNR. While using a single integrator for $L(e^{j\omega})$ provides a great deal of noise shaping, using two integrators can often provide more noise attenuation, especially at lower frequencies. Designing the filter $L(e^{j\omega})$ to be second order usually results in more noise shaping and better SQNR. However, second order feedback loops have the possibility of becoming unstable; hence, they must be designed to assure stability. Using three or more

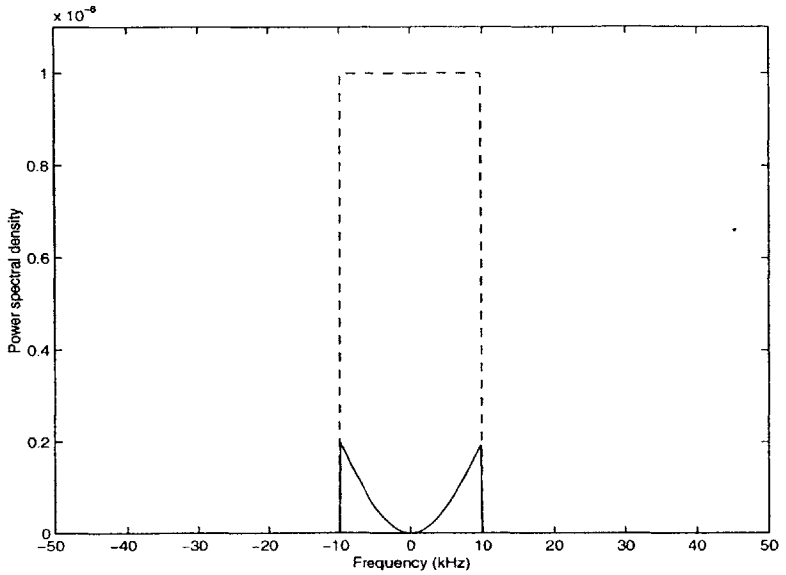

Figure 2-14: Noise Spectrum: Spectrum of $P(e^{j\omega})$ due to $E(e^{j\omega})$ after filtering. The two components are the same as in Figure 2-13. Note that the spectral density of the shaped noise is significantly reduced.

integrators in the loop results in increasing improvements in SNR, but the tradeoff is that such high order modulators become increasingly difficult to stabilize.

This section has mainly discussed a 1-bit, single-stage architecture, as it is the focus of the Σ-Δ modulator designs to be presented in this paper. However, there are many other Σ-Δ modulator design topologies, including those which utilize multi-bit quantizers, multi-stage noise shaping (MASH), and interpolative structures to name a few. These designs are not discussed in this thesis.

2.2 A Discrete-Time Third Order Σ-Δ Modulator

The previous section discussed the techniques of oversampling and noise shaping as they apply to a baseband signal. Since we are designing a digitizer for carrier band measurements, our modulator should be designed for the passband. Recall that the signal to be digitized lies in the spectrum of $f_c \pm 100$Hz (f_c is nominally 20kHz). This section will discuss the topology of a previously designed third-order modulator for this application.

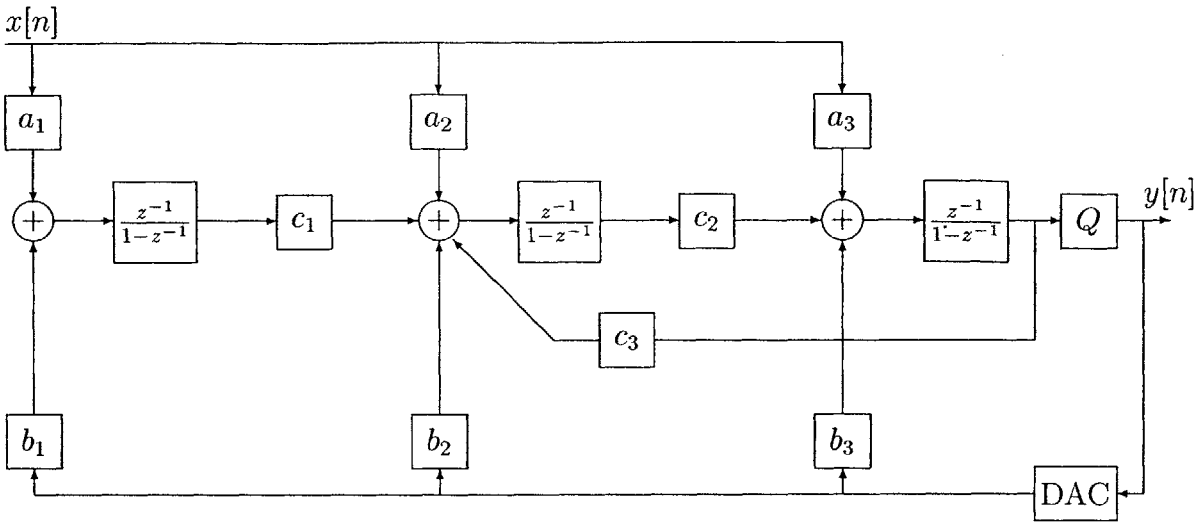

Figure 2-15: Block diagram of third order $\Sigma\Delta$ circuit.

2.2.1 Basic Modulator Operation

The block diagram for the modulator is shown in Figure 2-15. Note that it is implemented in discrete time—the sampling frequency f_s is set to be $256\times f_c$ (f_s is nominally 5.12MHz). The STF and NTF are:

$$STF = \frac{(a_1c_1c_2 - a_2c_2 + a_3)z^{-3} + (a_2c_2 - 2a_3)z^{-2} + a_3z^{-1}}{(-1 - b_3 + (b_2 + c_3)c_2 - b_1c_1c_2)z^{-3} + (3 - c_2(c_3 + b_2) + 2b_3)z^{-2} + (-3 - b_3)z^{-1} + 1} \tag{2.12}$$

$$NTF = \frac{(1 - z^{-1})(1 - 2z^{-1} + (1 - c_2c_3)z^{-2})}{(-1 - b_3 + (b_2 + c_3)c_2 - b_1c_1c_2)z^{-3} + (3 - c_2(c_3 + b_2) + 2b_3)z^{-2} + (-3 - b_3)z^{-1} + 1} \tag{2.13}$$

The frequency response is shown in Figure 2-16 for the values of a_n, b_n, and c_n in this design. The poles and zeroes have been chosen so that the NTF contains a notch at $f_s/256$, which is f_c. This notch serves to further attenuate the noise power at the frequency of interest, thereby increasing the SQNR.

We can calculate the noise power at the output using Equation 2.11. Instead of integrating over the interval $-f_o$ to f_o, we integrate over the passband interval $f_c - f_o$ to $f_c + f_o$.

$$OSR = \frac{f_s}{2 \times 100Hz} = 25600 \tag{2.14}$$

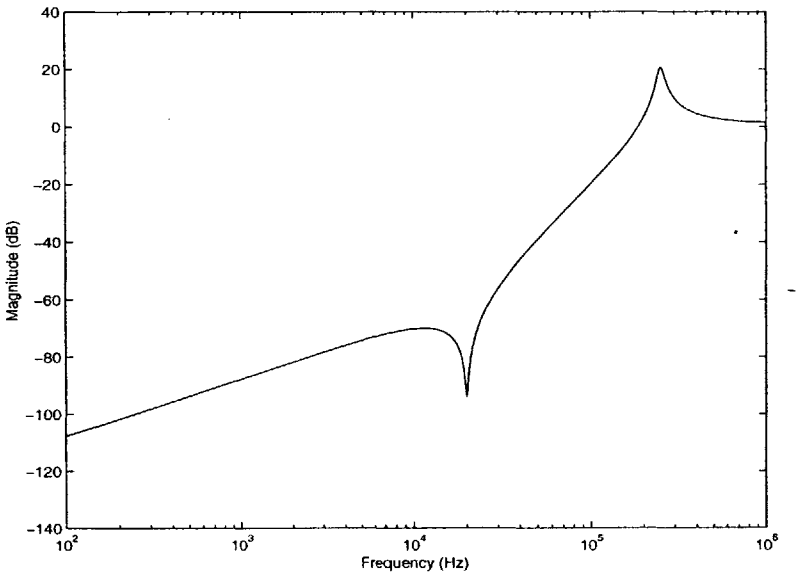

Figure 2-16: Frequency Response of the Noise Transfer Function with $f_s = 5.12MHz$.

$$k_e^2 = \frac{(\Delta = 5V)^2}{12 f_s} = 4.07 \times 10^{-7}\ V^2/Hz \tag{2.15}$$

$$P_e = \int_{f_c - 100Hz}^{f_c + 100Hz} k_e^2 \left| NTF(f) \right|^2 df = 9.27 \times 10^{-20} V^2 \tag{2.16}$$

$$SQNR = 10log\frac{P_s}{P_e} = 10log\frac{12.5V^2}{6.03 \times 10^{-13}} = 201dB \tag{2.17}$$

A SQNR of 201dB is extremely high. This resolution is much better than needed; however, we will see in the following section that implementing the system introduces added noise sources which limit the performance of this modulator.

2.2.2 Circuit Realization

When we examine the distinctions between CT and DT topologies in the next chapter, we will find that most of the differences lie in the circuit realization of the modulator. Therefore, it is important that we understand the implementation of this particular modulator. Of particular interest is the switched-capacitor integrator circuit and the 1-bit DAC in the feedback path.

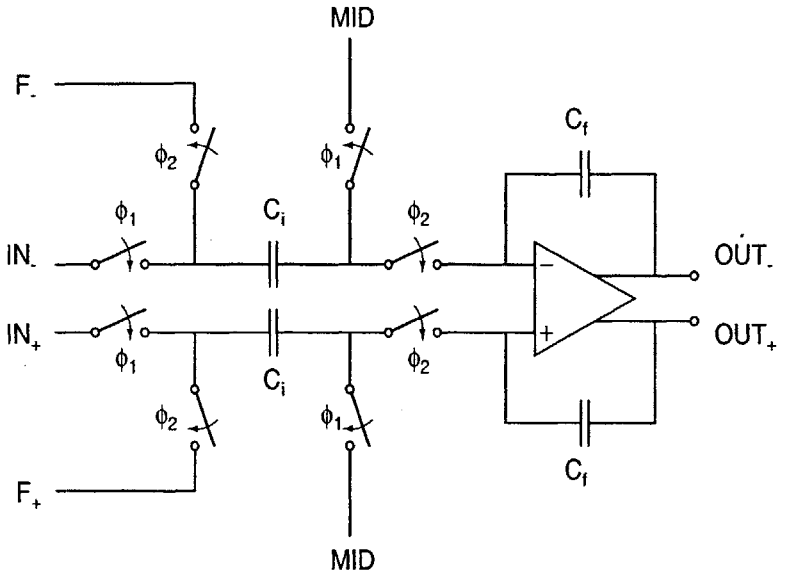

Figure 2-17: A fully differential switched-capacitor integrator

Switched-Capacitor Filters

The integrator blocks in the block diagram are realized using fully differential switched-capacitor integrators, shown in Figure 2-17. There are several reasons for using a fully differential structure. First, the input dynamic range is essentially doubled; since each input can span from 0V to 5V, two inputs can span from 0V to 10V when added together. Second, the differential output gain of the opamps can be made to be stable over process and temperature variations; therefore, any common mode DC offset in the input signal due to these changes will have little effect on the modulator performance, especially since the amplifiers and quantizer will effectively reject the common mode. The drawback of fully differential opamps is that they are usually larger in size and consume more power than single-ended designs (not to mention they are a little trickier to design).

The clock phases ϕ_1 and ϕ_2 represent alternating control signals; when ϕ_1 is ON, ϕ_2 is OFF and vice versa. Assuming a 50% duty cycle, this circuit's input to output transfer function is that of a DT integrator (such as the ones used in Figure 2-15). Taking $v_o[n]$ to be the voltage across OUT_+ and OUT_-, $v_i[n]$ to be across IN_+ and IN_-, and $f[n]$ to be across F_+ and F_-, the transfer function is:

$$V_o(z) = \frac{C_i}{C_f} \frac{z^{-1}}{1 - z^{-1}} (V_i(z) - F(z)) \tag{2.18}$$

This circuit provides not only the integration, but also the summing node needed in the modulator. Since the feedback coefficients a_i, b_i, and c_i are determined by the ratios of the input capacitors to the feedback capacitors of the integrators, the precise component matching of the CMOS process used in ASIC fabrication is a tremendous advantage when realizing these blocks. In fact, no part of this third order system is dependent on the absolute tolerances of any component; this characteristic is advantageous, because the on-chip capacitors in this fabrication process (.5μm CMOS) have an absolute tolerance of 8% in addition to a temperature coefficient of -25 ppm/°C. The matching tolerance, however, is accurate to better than 1%. The bandwidth of the notch filter is wide enough to accomodate this matching error; hence, we are able to place the poles and zeroes of our system to reasonable degrees of accuracy.

This integrator circuit is also the key to understanding where thermal noise is generated. Examining the circuit, it can be determined that the sampling noise from the input capacitors (C_i) is the major source (the noise power density is given by $\frac{kT}{C}$). Several capacitors on the input provide additive noise sources; hence, the input-referred noise is high enough to be troublesome. Increasing the capacitor sizes will diminish the noise level—however, this translates into quite a bit more area needed on chip.

Experimentally, the thermal noise floor is measured to be approximately -115dBc. The shaped quantization noise at 20kHz has been calculated to be much less than this value, so we can conclude that the limiting factor in the signal-to-noise ratio of this DT third-order modulator is thermal noise.

1-bit Digital to Analog Converter

At first glance, a 1-bit DAC does not appear to be useful at all. After all, a digital bit is represented by an analog voltage, so what is there to convert? Quite the

contrary, this block is a vital part of Σ-Δ modulator design.

The DAC is useful for several reasons. The first is that the DAC serves to buffer the output of the modulator from the summing nodes of the integrators. Since the output is fed back to multiple nodes in the system, a voltage buffer is necessary to avoid loading the output. Second, the DAC can scale the output voltage to the appropriate levels. This ability is not particularly useful, since our capacitor ratios can be altered to provide the correct scaling. Lastly, proper design of the DAC helps preserve the AC gain stability of the modulator. Even if the output was able to drive large capacitive loads, it would still not be prudent to simply connect the output directly to the summing nodes, because doing so results in large variations in AC gain.

In order to maintain AC gain stability, the charge fed back to the integrators must be consistent for a given quantization error signal; that is to say, $f[n]$ must be stable (either 0V or 5V for a given "n"). Since the quantizer is simply a comparator with a D flip-flop on the output, the analog output voltage level is dependent on the supply (VCC) that powers it. Over time, a digital "1" can change from 4.75V to 5.25V, depending on temperature and other external factors. So while a digital "0" would correspond to 0V, a digital "1" can vary over a large range of analog voltages. Such inconsistencies in the feedback signal will cause variations in the AC gain of the modulator.

In this system, the 1-bit DAC is implemented as shown in Figure 2-18. The digital output of the modulator is used as a control signal to switch F_+ and F_- between the stable reference voltages PREF and NREF (for this ASIC, PREF=4.25V and NREF=0V). These reference voltages source enough current to drive the integrators' input capacitors and vary less than .35% over the full range of temperature and process variations. This design is necessary to stabilize the AC gain of the modulator.

Understanding the operation of the switched-capacitor integrator (as well as its limitations) is crucial to understanding the advantage of a CT implementation. We have already seen that the thermal noise generated by the input capacitors is the limiting factor in the achievable resolution of this modulator. The next chapter will

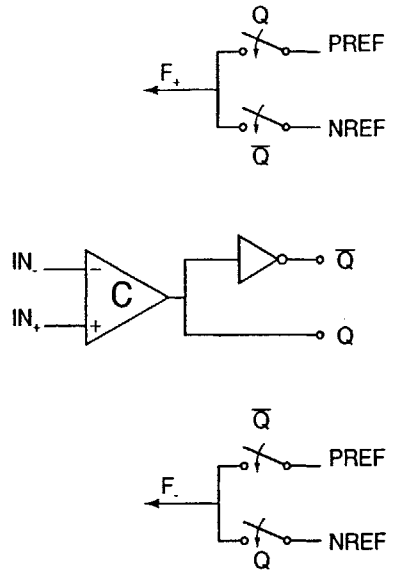

Figure 2-18: 1-bit feedback DAC as it is implemented in the existing third order modulator design

reveal further weaknesses and suggest a CT topology as a solution. Furthermore, we shall see that when it comes time for us to design a CT modulator, the 1-bit DAC will require extensive redesign.

Chapter 3

Advantages and Disadvantages of Continuous-Time Σ-Δ Modulators

In the previous chapter, a third-order discrete-time (DT) modulator was presented. The major performance limitation is due to the high level of thermal noise at the input of the modulator. It is believed that a continuous-time (CT) modulator topology may eliminate this problem. A fair amount of research has been done on CT topologies, as they are potentially useful in minimizing power dissipation, silicon area, and thermal noise.

While there are many benefits to be had from this implementation, there are several significant drawbacks which make such a topology unpopular, especially for design on silicon. The following chapter presents research done on CT modulator topologies and investigates its viability for our specific application.

3.1 Benefits

In redesigning the third-order modulator from the previous chapter, we are definitely interested in reducing the thermal noise floor at the input, but other areas of interest include reducing power dissipation and area on chip. Since ultimately, several of these A/D converters may be needed for "system-on-chip" designs, it is worthwhile to target these areas when redesigning.

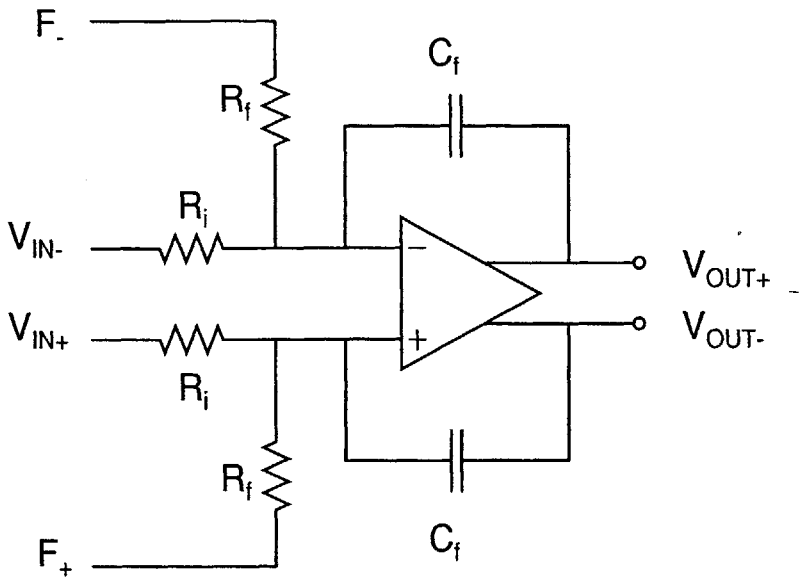

Figure 3-1: A simple continuous-time integrator (fully-differential)

One of the major building blocks of our Σ-Δ modulator is the integrator, which we have seen implemented in discrete time using switched-capacitor techniques (Figure 2-17). A simple CT integrator can be implemented as shown in Figure 3-1.

3.1.1 Reducing power consumption

The desire for increasingly better digitizers continues to grow, and more often than not, the solution is to design higher order Σ-Δ modulators. Much work has been done in obtaining methods to stabilize such designs; however, another drawback of high order modulators exists: increased power consumption. Adding another switched-capacitor integrator (and thus another opamp) to a modulator design requires the use of a significant amount of additional power.

Switched-capacitor integrators require a large amount of current in order to function normally. This trait is due to the fact that the output must be able to settle quickly when driving capacitive loads: the feedback capacitor (C_f) and the input capacitance of the successive integrator. It is imperative that the outputs not be slew-rate limited, since studies have shown that it can cause an increase in harmonic distortion and quantization noise [1]. While DT designs are usually sensitive to slew-

rate limiting at the input, oversampled system like Σ-Δ modulators are not, because the input changes slowly with respect to the sampling clock. However, Σ-Δ modulators can often have slew-rate limiting problems inside the loop at the integrator outputs.

Furthermore, the integrator output must settle to its final value prior to the next switching cycle; otherwise, a gain error will be introduced into the modulator, which can further degrade SNR. At low sampling frequencies, this error may be easily avoided, but for high-resolution Σ-Δ modulators, it is most likely that the oversampling ratio will be as large as possible, resulting in a short amount of time between clock phases. With such a small amount of time to charge these load capacitances, the switched-capacitor opamps must be designed so that the output stage is capable of supplying large currents, thus avoiding long settling times or slew-rate limiting. The use of large currents translates into using high-power opamps for Σ-Δ modulator design.

CT modulators can more easily avoid the slew-rate problem. The main reason is that the load capacitance seen at the integrator outputs can more easily be minimized when using CT integrators. Therefore, the integrator opamps can be designed to supply less current to the output in order to avoid slew-rate limiting. Moreover, the bandwidth requirement of the opamps is reduced; designing opamps with less bandwidth can also reduce power consumption. The extent to which power can be saved is still dependent on the capacitances seen at the output node, but the opamp design constraints are now more relaxed. Qualitatively, we can see that the power consumption can be reduced from our original DT design—we will hopefully get a better feel of the quantitative reduction once we design a comparable CT Σ-Δ modulator.

3.1.2 Decreasing chip area

There are two major characteristics of CT modulators that decrease the amount of real estate required on chip. The first involves the need for an anti-aliasing filter in DT systems. In order to avoid high frequency noise from aliasing into the band of interest

(and reducing the SNR), an anti-aliasing filter must be placed before the analog signal is sampled. Since DT modulators inherently sample (and hold) when using switched-capacitors, the filter must be placed before the modulator input. However, CT modulators sample the signal inside the modulator loop, prior to the quantizer. In this type of design, the modulator loop inherently provides anti-aliasing!

The Nyquist criterion tells us that if we sample a signal at a frequency of f_s, any two tones that differ in frequency by a multiple of f_s overlap in the spectral plot; that is, they are indistinguishable. Our band of interest is nominally 20kHz; hence, any spectral component (more specifically, noise) at $f_s + 20kHz$ will be appear at 20kHz, thus adding more noise power in the signal band! Anti-aliasing filters remedy this problem by lowpass filtering the signal before sampling, thus attenuating any high-frequency noise which may alias into the signal band.

Examining the block diagram of Figure 3-2, we can see that *where* the signal is sampled has some importance with respect to aliasing. If the signal is sampled at point "A" (prior to the input), then an anti-aliasing filter is needed to attenuate the high frequency noise. However, if the sampling is done at point "B" (as in CT designs), then the first three integrator stages act as an anti-aliasing filter! The transfer function from point A to point B should resemble a third-order low-pass filter, giving us roughly 60dB of attenuation at the sampling frequency (assuming each integrator has a crossover frequency one decade below f_s in order to stabilize the loop). Thus, higher-order CT modulator designs provide increasingly better anti-aliasing.

The second characteristic which allows us to reduce chip area is evident in the circuit requirements of the switched-capacitor and CT integrator. Making resistors in silicon requires much less space than making capacitors. Hence, we can realize integrators with larger gains and bandwidths using less silicon area. This fact is easy to see if we examine the transfer function for the CT integrator shown in Figure 3-1.

$$V_{out}(s) = \frac{1}{sR_iC_f}(V_{in}(s) - \frac{R_i}{R_f}F(s)) \tag{3.1}$$

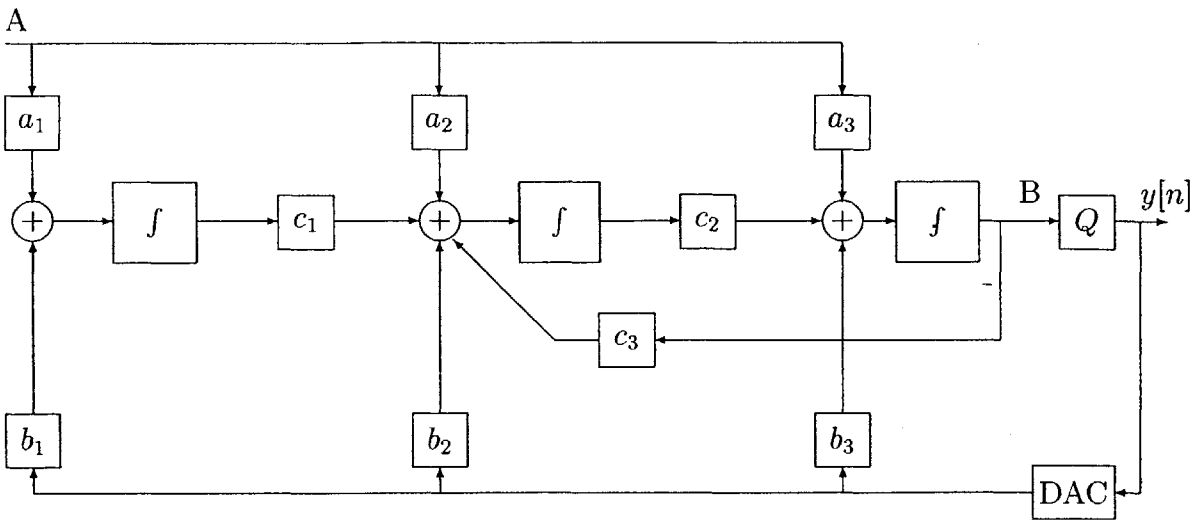

Figure 3-2: Block diagram of third order ΣΔ circuit.

Smaller resistors and capacitors result in larger gains, whereas in the switched-capacitor case, the integrator characteristics are determined mainly by the ratio between two capacitors (see Equation 2.18). It should be obvious that a CT integrator can use the minimum sized capacitor for that particular process and vary the resistor size to realize a wide range of gains (minimum and maximum resistor sizes differ by about three orders of magnitude for this particular process). A DT integrator's gain has a much greater effect on its size (assuming the sampling frequency cannot be varied). This effect relaxes the area constraints on our modulator designs, allowing smaller designs to achieve similar resolutions.

Another way CT modulators can become more compact is that certain physical devices behave as CT integrators (some accelerometers and fluxgate magnetic sensors are examples [1]). There exist certain circuit designs which use the physical characteristics of these devices to implement the first stage of a Σ-Δ modulator, thus eliminating the space needed for a circuit implementation of the first stage. While this technique is not relevant to this design, it is an interesting one nonetheless.

3.1.3 Lowering thermal noise

It is has been stated that the first stage to any Σ-Δ modulator design is the most

important. The major reason for this assertion is that any circuit non-linearities or noise sources in the later stages of the modulator are divided by the gain of the first stage when those errors are input-referred. Thus, any precise circuit requirements needed to minimize such non-linearities can be relaxed in the second and third stages of the modulator, even though the first stage necessitates excruciating attention to detail.

This feature is especially relevant when discussing circuit noise. The main circuit noise sources are those at the first stage. The main noise sources include those from opamps and from components such as switches, capacitors, and resistors. Opamp noise is generated by the MOSFET transistor pair in the opamp's differential input stage. There are two types of noise generated by MOSFETs: flicker noise (or 1/f noise) and thermal noise. Flicker noise is small at higher frequencies, but can have a substantial effect at low frequencies. The voltage noise appears at the gate, and it's noise power is given by [2]:

$$V_{g1}^2(f) = \frac{K}{WLC_{ox}f} \tag{3.2}$$

The term K (not to be confused with Boltzmann's constant) is a value dependent on device characteristics. Since flicker noise is inversely proportional to the area of the MOSFET, we need only to increase the device sizes of the appropriate transistors. The MOSFET's thermal noise is generated by the resistive channel of the MOSFET in the active region. Although current noise is generated in the channel, it translates to voltage noise at the gate. The noise power at the gate due to thermal noise is calculated to be [2]:

$$V_{g2}^2(f) = \frac{8kT}{3g_m} \tag{3.3}$$

In this case, the "k" is Boltzmann's constant ($1.38 \times 10^{-23} JK^{-1}$) and T is absolute temperature. Combatting this noise is a simple matter of increasing the transconductance (g_m), which can be done by increasing the device size as well as increasing the bias current. When designing opamps, the input stage is often biased at higher

currents than subsequent stages for the purpose of minimizing thermal noise; the tradeoff is increased power consumption.

These noise sources are present in both DT and CT modulators. However, the noise generated from components are significantly different. We will discuss the DT case first. Examining the switched-capacitor integrator again (Figure 2-17), we see that the only components at the input are capacitors and switches. The switches are implemented with MOSFETs; subsequently, the modulator input is mainly affected by the thermal noise generated in the channel. This noise is small compared to the capacitor noise. While capacitors do not generate noise of their own, they do accumulate noise from other sources. In this case, noise is accumulated from the switches. It has been shown that capacitor noise power is independent of the noise power from other sources and is given by Equation 3.4. Note that since this capacitor is sampled by a switch, the capacitor noise will be slightly larger than this value due to aliasing. Nevertheless, it is still a good estimate of the thermal noise.

$$V_C^2(f) = \frac{kT}{C} \tag{3.4}$$

On-chip capacitors tend to be small (on the order of picofarads), thus the noise generated tends to be high. Furthermore, since switched-capacitor integrators contain a feedback capacitor (C_f) and an input capacitor (C_i), there are two sources of this thermal noise. Thus, when choosing capacitor sizes, the designer must trade off between component area and thermal noise level.

CT modulator designs use input resistors, feedback capacitors, and no switches to implement the integrators. So while the noise generated from the single feedback capacitor still exists, the noise from the input capacitor is absent. Instead, the input resistor (R_i) generates noise with a noise power of:

$$V_R^2(f) = 4kTR \tag{3.5}$$

In comparison to the capacitor noise, this value is significantly smaller when considering that on-chip capacitors are usually tens of picofarads (10^{-11}) while on-chip

resistors are usually tens of kiloohms (10^4). It is easy to see that when dealing with components of these sizes, CT integrators generate less noise than their DT counterparts.

3.2 Drawbacks

Although CT modulators seem well suited to solve the problem of a small, low-power, high-resolution A/D converter, they are not without their difficulties. Indeed, it is these complications that prevent CT topologies from gaining more widespread use. Such problems with CT modulators include poor absolute component tolerances, data-dependency in the feedback DAC, and clock jitter. However, the previous section should convince any Σ-Δ modulator designer that if these issues could be overcome, CT architectures would be an extremely powerful solution to the increased demands for low power Σ-Δ conversion.

3.2.1 Absolute component tolerances

In section 2.2.2, we saw that one of the more useful characteristics of switched-capacitor filters on a silicon process was that the frequency response was solely based on component ratios, which are well-controlled on silicon. Moreover, the frequency response is dependent on the sampling frequency, which is beneficial in this case, because the band of interest depends on the sampling clock as well. However, the transfer function of the CT integrator of Figure 3-1 relies on the value of the the time constant RC.

Unfortunately, since the poles and zeroes of the system are now determined by RC time constants, we can no longer accurately place poles and zeroes as in a switched-capacitor circuit. The IC process used for this project specifies the absolute tolerance of high-poly resistors to be $\pm$20% and the temperature coefficient to be -1446ppm/°C. Furthermore, the absolute tolerance of capacitors is $\pm$8%, with a temperature coefficient of -25ppm/°C. As if these problems were not difficult enough to deal with, an added headache results from the carrier band having a range anywhere from 15kHz-

25kHz, depending on physical device characteristics. Considering this frequency variation and the poor absolute tolerances of these on-chip components, it becomes clear that any CT modulator design for this application must be tolerant to large errors in pole and zero placement.

3.2.2 Sample-and-hold circuit

Another small annoyance of this topology is the need for a sample-and-hold circuit. Switched-capacitor modulators inherently sample and hold, thus eliminating the need to design one. For this topology, a separate sample-and-hold must be designed and implemented in front of the quantizer. This requirement results in increased silicon area. However, we have seen that CT modulators eliminate the need for anti-aliasing filters; considering the effectiveness of using the integrator stages for a filter, the added nuisance of a sample-and-hold is trivial, as is the area requirement. We will see later on that our particular design constraints for this SHA are extremely loose, and as a result, the area requirement is nearly inconsequential.

3.2.3 Return-to-zero DAC

The 1-bit D/A converter in the feedback path of the modulator is usually used to buffer and sometimes scale the output bitstream before feeding it back to the input. In a CT modulator, the design of this DAC becomes significantly more complicated, because a data dependency occurs when the output signal levels are directly fed back to the input [5]. Each bit transition requires a small amount of time to settle; so the energy of the pulse for a long string of "1's" is not the same as the series of pulses containing the same number of "1's." Figure 3-3 illustrates this concept.

A data dependency results because for an output with a high density of "ones," more energy is fed back to the input than for an output with equal density of "ones" and "zeroes." Therefore, the system characteristics change slightly depending on the input. This energy fluctuation is not a concern in DT modulators, because the DAC is controlled by the sampling clock and is implemented in such a way that the output

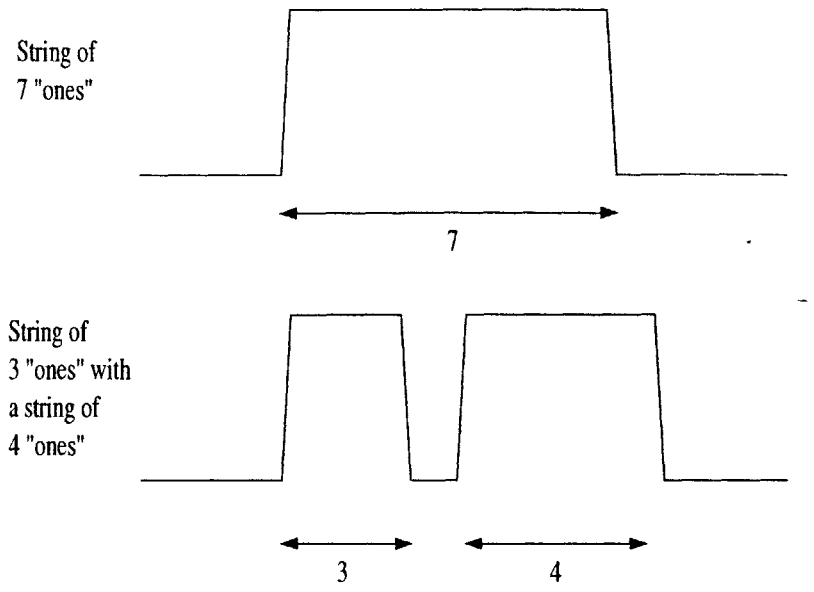

Figure 3-3: Two waveforms with 7 "ones." Due to rise and fall times, the energy (area under the curve) in the first waveform is slightly different than the energy in the second.

voltage levels are fed back to the input *after* settling.

We can rectify the situation by implementing a "return-to-zero" D/A converter in the feedback path. Such a device generates one pulse for every sampling period; thus, the energy fed back to the input is the same for every "one." Therefore, the energy of the signal in the feedback path is independent of the data. The return-to-zero waveform is illustrated in Figure 3-4.

There are other minor complications which arise from using RTZ DACs. We will discuss these further when it comes time to design a CT modulator of our own.

3.2.4 Clock Jitter

Perhaps the most problematic complication arising from the use of CT modulators is that of clock jitter. Any phase noise present in the sampling clock is very easily introduced into the modulator as a large noise source. Why are CT modulators more susceptible to this kind of error than DT modulators? Figure 3-5 helps to illustrate why [1]. The Δt term is the clock edge's time error due to jitter. The gray-shaded area represents the change in charge (Δq) fed back to the input as a result of Δt. In

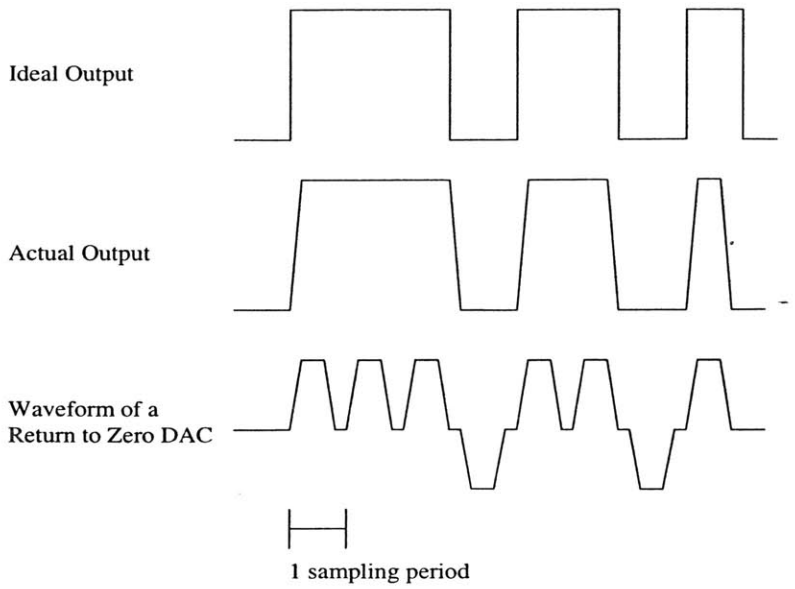

Figure 3-4: Waveform of a Return-to-Zero DAC.

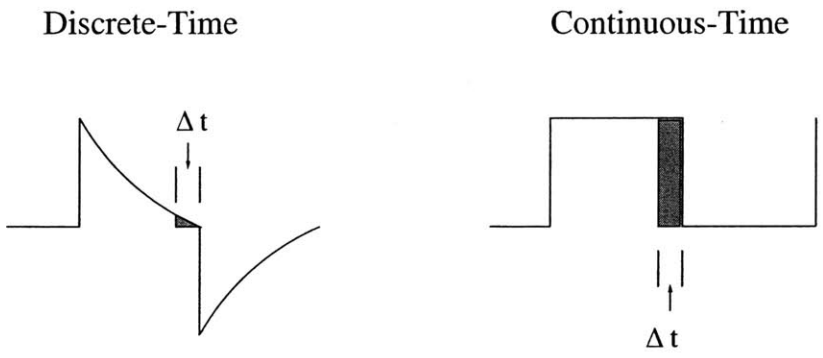

Figure 3-5: Feedback current waveforms for a single sampling period

CT modulators, this amount is very large. One can think of Δq as representing the error signal fed back to be corrected.

For DT modulators, the output voltage is sampled, then used to charge the input capacitors of the first integrator stage. Hence, most of the current usually discharges into the capacitor at the beginning of the clock period; any jitter in the sampling clock results in a negligible error. However, CT modulators supply current to the input resistors of the first integrator stage for the duration of the pulse. Clock jitter is troublesome in this case, as quite a bit more energy is fed back to the input with this kind of feedback waveform. The result is that a fair amount of random error (noise) is introduced at the modulator input. This noise can degrade SNR, especially

when trying to achieve high resolution.

Further difficulties arise from the use of a RTZ DAC, discussed in the previous section. The effects of clock jitter are exacerbated by the use of a RTZ scheme. One reason is that a RTZ pulse is inherently shorter than a non-RTZ pulse; consequently, the same Δt will result in a Δq that is a larger percentage of the total feedback charge. Another reason is that assuming that both rising and falling edges of the RTZ pulse are controlled by the clock, the jitter affects both edges, and the error becomes $2 \times \Delta q$ per clock cycle. By contrast, a non-RTZ scheme only has one edge affected (see Figure 3-6).

However, the second effect is dependent on the implementation of the RTZ pulse; if the pulse is generated using a monostable multivibrator (or "one-shot"), then the jitter effects on one of the edges can be eliminated [6]. Unfortunately, one-shots with stable pulse-widths are difficult to realize on-chip (without using offchip components).

It is difficult to accurately quantify the effects of jitter on the modulator; likewise, it is also difficult to combat its effects. In fact, the most effective way is to eliminate the jitter by using low jitter clock sources, such as crystal oscillators. Unfortunately, these are not solutions that are viable for integrated circuits; an off-chip crystal would probably be larger than the IC!

Sampling clocks are usually generated on chip by a voltage-controlled oscillator (VCO). The current third-order sigma-delta modulator ($\Sigma\Delta$M) is driven by a clock produced by a phase-locked-loop using a VCO. The most obvious solution for this design effort is to minimize the phase noise of this oscillator. There has been extensive work done on low phase-noise VCO's, and it is hoped that one can be designed that will minimize the noise injected into the $\Sigma\Delta$M . One of our objectives is to document the effects of clock jitter on our designed CT modulator and determine a reasonable phase noise specification for future VCO designs.

By this point, it should be clear that the limitations on CT Σ-Δ modulators are quite a nuisance and often prevent designers from taking advantage of their strengths. Nevertheless, it is worth the time to investigate the possibility of using a CT architecture for our specific application. The remainder of this thesis documents design

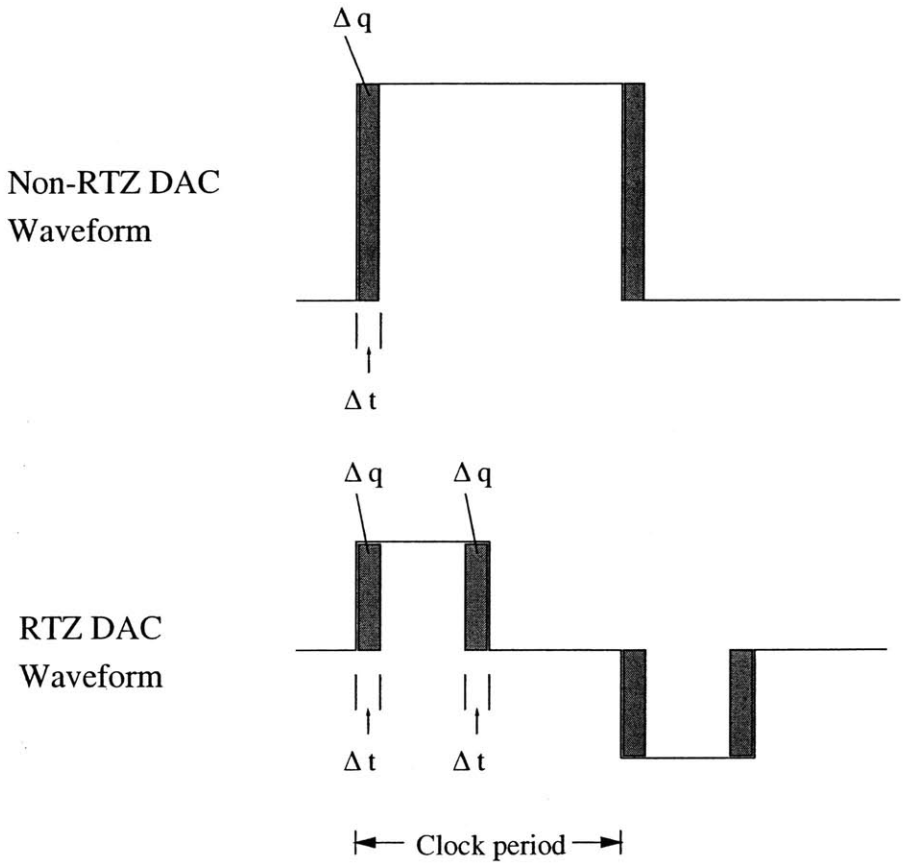

Figure 3-6: Clock jitter effects on the current waveforms of two different DAC implementations.

methodology of higher-order $\Sigma\Delta$M's and the resulting performance.

Chapter 4

Design of a third-order continuous-time Σ-Δ modulator

Our objective is to design a low-power, low-area, continuous-time (CT) Σ-Δ modulator with lower thermal noise than the discrete-time (DT) design presented in Chapter 2. The modulator is to be built on a CMOS process. Also of interest is the quantitative effects of clock jitter on the input referred noise of the modulator loop. The observed effects will allow us to obtain an upper bound for the phase noise of a voltage-controlled oscillator so that we may achieve an acceptable signal-to-noise ratio.

This chapter details the design methodology associated with designing a third-order CT modulator as well as the calculations and simulation results from this specific design.

4.1 Design possibilities

4.1.1 Direct impulse-invariance transformation

The simplest method for designing CT modulators is by first designing a DT modulator and then finding the CT equivalent. Since much software has been created to assist in DT modulator design, this method is the most straightforward. Finding

the CT equivalent can be done by using any number of transforms to go from the z-domain into the s-domain. While many such transforms exist, it is the impulse-invariant transform which cause the CT and DT loops to be the same, according to [1]. Ideally, the design of a CT modulator's noise transfer function (NTF) is no more difficult than that of a DT modulator; the process merely involves the relatively easy step of transforming $NTF(z)$ to $NTF(s)$. Since a DT modulator which meets spec already exists, does this third-order CT design simply boil down to a quick transformation?

Sadly, this method has some inherent problems for our design. While it may appear attractive in theory, it fails to work for this particular application when the time comes to build the modulator in silicon. Recall that the NTF for the DT modulator has a resonator structure that creates a notch filter at the carrier frequency (f_c). The exact frequency, as well as the bandwidth and quality factor, of this notch is determined by capacitor ratios. However, for a CT modulator, it would be determined by RC time constants. Component tolerances (and temperature coefficients) would make accurate placement of this notch filter nearly impossible. Also, the wide range of frequencies for f_c makes it difficult to know what frequency to place the notch at.

Figure 4-1 verifies this behavior. Assuming that the minor loop feedback gain ($c3$ in Figure 2-15) is altered by about $\pm 25\%$, the notch frequency varies from about 17kHz to 24kHz, which is almost the entire range of the carrier frequency. This error occurs only if one of the variables is altered; it should be obvious that this frequency can vary even more we ran this simulation varying all the resistive and capacitive components. In this case, not only would the notch frequency change, but the bandwidth and quality factor of the resonance would fluctuate as well.

The poor component tolerances make designing a reliable resonator structure difficult. This problem is unfortunate, since our information band is very small; a well designed resonator structure easily provides the noise shaping needed to filter out a significant portion of the quantization noise. This obstacle is the first of many encountered in CT modulator design.

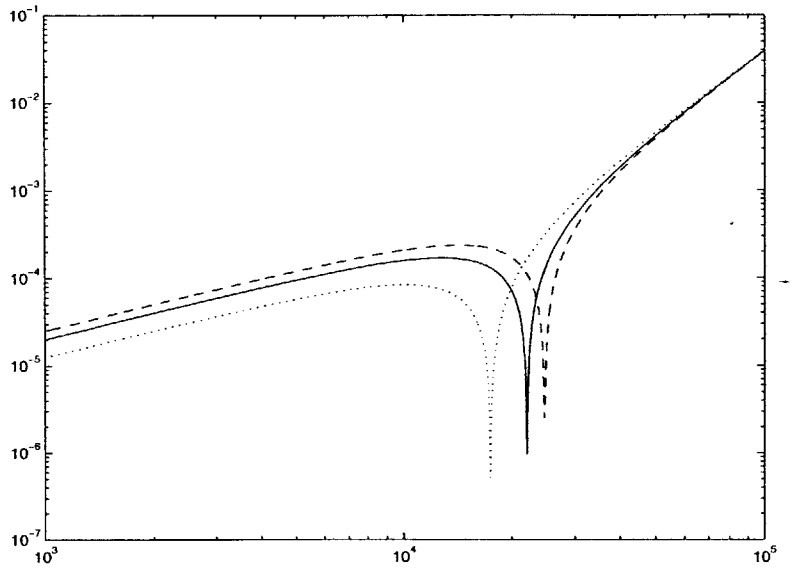

Figure 4-1: Change in notch frequency with a 25% variance in $c4$

4.1.2 Tunable on-chip components

One way to gain more accurate notch placement is by dynamically adjusting the component values so that the notch is placed at the frequency indicated by the sampling clock (which is exactly 256 times f_c). Such a scheme would involve adding or subtracting capacitance (or resistance) for several components.

This approach is daunting, to say the least. If every component were to be controlled, we would need a large amount of circuitry, especially for higher-order modulators. With all of the added circuits, the CT design could easily end up using more power and area than the DT design. Furthermore, the complexity of designing tunable components is difficult to manage. Therefore, this approach is very unattractive.

4.1.3 Continuous-time front end

Another way to overcome this problem is to use a CT filter for the first stage only, and implement DT filters for the rest of the stages as shown in Figure 4-2.

This topology allows the designer to place the poles and zeroes (and thus the notch frequency) with improved accuracy. While the first stage suffers from the large component tolerances, the successive stages are as impervious to it as the DT design.

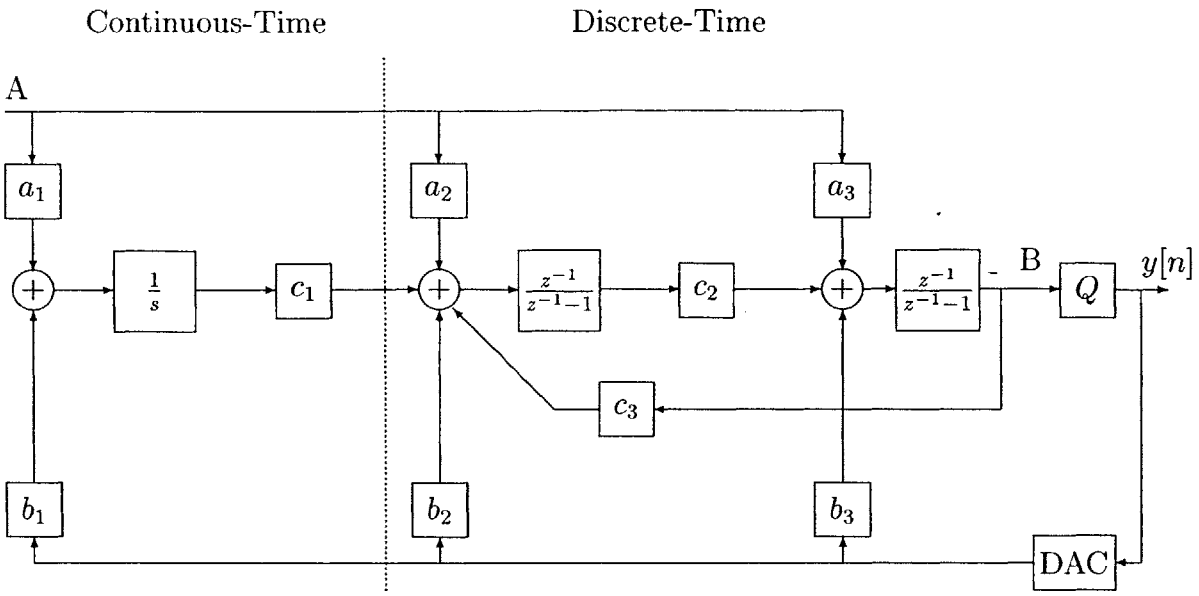

Figure 4-2: Block diagram of third order $\Sigma\Delta$ circuit.

Since the resonator is realized in the second and third stages, it should be able to be placed with reasonable accuracy, even though the pole due to the first stage will still affect the system somewhat.

This design will still eliminate excess input-referred thermal noise, because only the successive stages have capacitor-induced thermal noise, which does not significantly affect the input. Moreover, since only the current path through $b1$ is affected by clock jitter, this design may be less sensitive to jitter than a complete CT implementation. This solution has been successfully implemented in several different applications. The only drawbacks are that it does not conserve either power or area.

Only the first stage is different from the DT design, so we may be able to conserve power by biasing that opamp with lower currents. However, this opamp must have a high bias current on the input stage to limit the noise from the MOS differential pair. It is unlikely that a significant amount of power can be saved; the most effective way to do so would be to design the opamp to operate with a high input stage bias current and a low output stage bias current.

Likewise, the area will be minimized, but only for the first stage; the area due to capacitors will be the same for the second and third stage. Furthermore, one stage of

CT integration implies that the signal must be sampled and held at the input of the second stage. Hence, the anti-aliasing filter is realized by only a single-pole lowpass filter. Depending on the nature of the input, another filter may need to be added at the modulator input.

4.1.4 Lowpass modulator topology

While the application is for carrier band measurements, it is not necessary to use a notch filter in the loop to rid the output of quantization noise. It should be clear that using a resonator structure by itself is not feasible with the resolution we wish to obtain. Combining the resonator with a lowpass topology (as in Figure 2-15) is reasonable, and several solutions to solve the component tolerance problem have been suggested in this section.

Eliminating the resonator may seem like a strange proposal at first, but in fact, it is very reasonable. The SNR calculations of the DT modulator in Chapter 2 projected an extremely high SNR, much greater than the thermal noise floor would allow in practice. Hence, we can conclude that the resonator may not be necessary to achieve the needed resolution. Furthermore, the modulator loop may be easier to stabilize without the added phase contribution from the resonator.

A lowpass topology implemented completely in continuous time can take advantage of decreased thermal noise, lower power, and lower area, assuming the SNR is high enough. Fortunately, the oversampling ratio is set to be very high for this design. Consider that the oversampling ratio is set to be 25600 ($\frac{f_s}{f_o}$). By Equation 2.8, the SNR due to oversampling is:

$$SQNR(dB) = 10log(6) + 10log(25600) = 51dB \tag{4.1}$$

Targeting an SNR of 111dB (the SNR of the DT modulator), the noise shaping would be required to attenuate the quantization noise by approximately 70dB at f_c (nominally 20kHz). Taking into account the variance in the component values and the error introduced by making the white-noise quantizer assumption, it is prudent to

aim for a greater value (about 80dB). This design parameter appears to be achievable.

Choosing to implement a CT front end is probably more reliable in terms of achieving sufficient quantization noise shaping and decreasing sensitivity to component tolerances. However, a lowpass modulator design sacrifices insensitivity to component tolerances in favor of power and area savings. Since reducing these two specifications is of primary importance when building microelectronics, this solution is the most attractive.

4.2 Design Methodology

While design of first and second order modulator loops is fairly straightforward, higher order design is significantly more involved, because the modulator loop becomes *conditionally stable.* Conditionally stable loops are stable for inputs with small signal swings, but they destabilize for large swings. This problem occurs because of saturation at the integrator outputs; as the opamps saturate and the outputs become cut off at the rails, this state can introduce nonlinearities which can push the loop into instability. Modeling nonlinearities is no small feat, and developing good intuition on stabilizing nonlinear elements is even more difficult!

For this reason, it is nearly impossible to hand design high order Σ-Δ modulator loops; thus, computer simulation of Σ-Δ modulators is often necessary. This section will outline the design methodology used to design the CT third-order modulator presented in this thesis. We first use hand design tools such as root-locus analysis and Bode plots to design a modulator loop. We then use this design as a starting point to use software simulation. Using results from simulation and adjusting component values accordingly, we hope to arrive at a modulator design that meets specifications.

4.2.1 Initial hand design

Before starting the design of a third-order modulator loop, we must first examining the dynamics of a typical lowpass topology shown in Figure 4-3. Recall that the quantizer is modeled as an additive white noise source: assuming that $e[n]$ has no

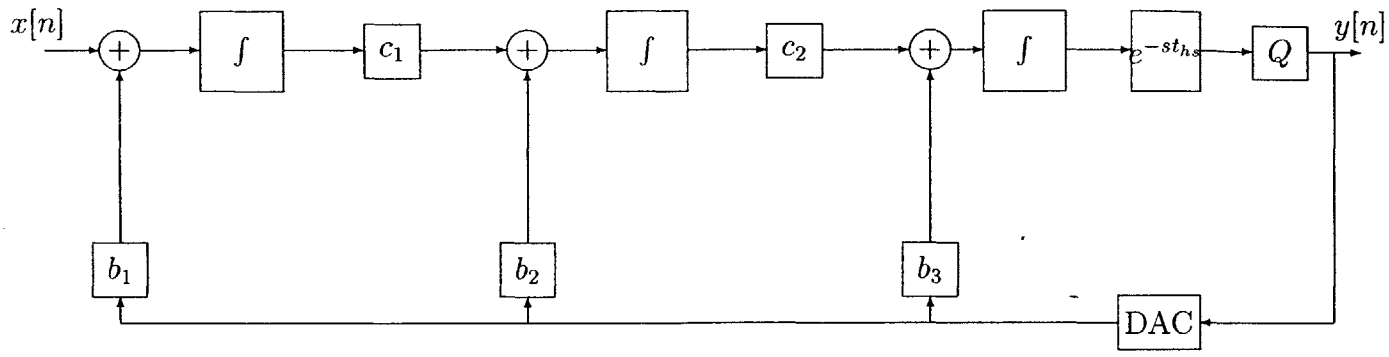

Figure 4-3: Block diagram of a third order lowpass $\Sigma\Delta$ modulator.

effect on the stability of the loop, it can be ignored. One block that has not been included before is the sample-and-hold block. It is important to model in this case, because sampling adds a delay to the system. This delay is modeled in the Laplace domain as $e^{-s\times t_{hs}}$, where t_{hs} is half the sampling period [9]. It is easy to see that while this sample-and-hold has a gain of 1, it does contribute phase as linear function of frequency.

First, we examine the NTF in detail:

$$NTF(s) = \frac{s^3}{s^3 - b_3 s^2 - b_2 c_2 s - b_1 c_1 c_2} \tag{4.2}$$

The transfer function tells us that there are three zeroes at the origin and three poles somewhere in the left half plane (for negative feedback, all b_n are negative). At low frequencies, the noise is greatly attenuated, and the amount of attenuation gradually drops off as the frequency increases.

To apply Bode plot analysis to the stability of the loop, we must determine the loop gain. Recalling Equation 2.10, the loop gain can be found in terms of the NTF.

$$L(s) = \frac{s^3 - b_3 s^2 - b_2 c_2 s - b_1 c_1 c_2}{s^3} - 1 = \frac{-b_3 s^2 - b_2 c_2 s - b_1 c_1 c_2}{s^3} \tag{4.3}$$

Equation 4.3 indicates that the loop gain has three poles at the origin and two zeroes in the left half plane, whose locations are determined by the integrator gains. Now that the loop gain has been determined, conventional bode plot analysis can be applied in placing the zeroes to keep the loop stable and maximize noise attenuation

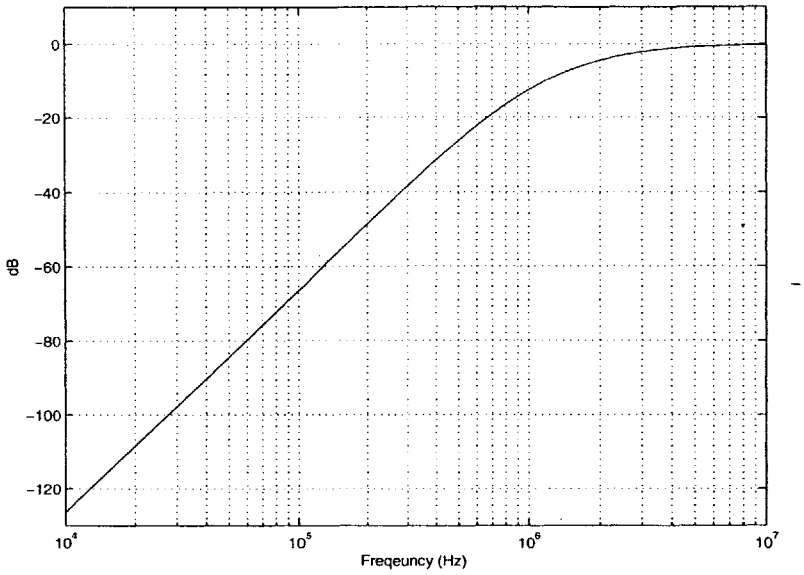

Figure 4-4: Frequency Response of the NTF with all poles located at unity-gain frequency 1.28MHz

at 20kHz.

First, consider the negative phase contributed by the half sample delay; the phase drop is given by $2\pi f \times t_{hs}$. For f_s=5.12MHz, this delay will contribute -45 degrees of phase at a frequency of 1.28MHz. Since phase drops linearly (i.e. very quickly) with frequency, it can be concluded that the loop transmission's unity-gain frequency should be no more than 1.28MHz in order to safely maintain phase margin. Also, the system zeroes should be placed lower than this frequency in order to add some phase recovery at crossover.

Since the magnitude of the NTF can be approximated as the inverse of the loop gain, it is reasonable to assume that whatever the loop gain's unity-gain frequency is, the NTF magnitude will be close to unity at that frequency as well. Thus, a rough picture of the NTF magnitude can be painted. For the NTF, the three integrators yield a 60dB/decade slope at low frequencies. It should be apparent that in order to maximize noise attenuation at 20kHz, the NTF poles should be placed as high as possible. Supposing that all three poles are located at 1.28MHz and that the gain is unity at that frequency, the noise attenuation at 20kHz would be nearly -110dB! Figure 4-4 illustrates this transfer function.

We have determined that the NTF poles should be placed as high as possible and that the loop transmission zeroes should be placed at around the crossover frequency of about 1.28MHz. How do these two conclusions relate to each other? If we assume that the loop gain is approximately equal to the inverse of the NTF magnitude, then we can surmise that the real locations of the NTF poles are roughly equivalent to the loop transmission zeroes.

It is important to note that the previous statement is certainly *not* true, as this approximation breaks down as $L(s)$ approaches unity; the fact that there are three NTF poles and only two loop gain zeroes is enough to disprove this conclusion. However, it does give us a *very* rough idea of how the locations of the NTF poles and loop transmission zeroes are related; more specifically, we can reasonably state that if the loop transmission zeroes were both located at the unity gain frequency, then the NTF poles would probably be located within a decade of that frequency. Hence, when designing the third-order modulator, the loop gain zeroes should be placed as high as possible (while still keeping the loop stable) so that the NTF poles will be high enough to maximize the noise attenuation at 20kHz. It should be evident that the upper bound on noise attenuation is limited by the possibility of the loop becoming unstable.

Now that we have an understanding of how stability and noise attenuation are related, we can easily go about hand designing the loop. We will do so by picking the zeroes for the loop gain. Since we would like the crossover frequency to be as high as possible, we will initially choose 1.28MHz to be the unity-gain frequency. Recall that the half-sample delay of the SHA will contribute -45 degrees of phase shift at this frequency. So in order to achieve 45 degrees of phase margin, we must design the rest of the loop to achieve 90 degrees of phase margin. Since the integrator poles contribute -270 degrees of phase at low frequencies, the two zeroes must contribute at least 180 degrees of phase at 1.28MHz. Therfore, the highest frequency the zeroes can be placed at is roughly 128kHz, one decade below crossover. The bode plot of this system is shown in Figure 4-5. The corresponding NTF magnitude response is shown in Figure 4-6. It is important to note that placing the zeroes *exactly* one

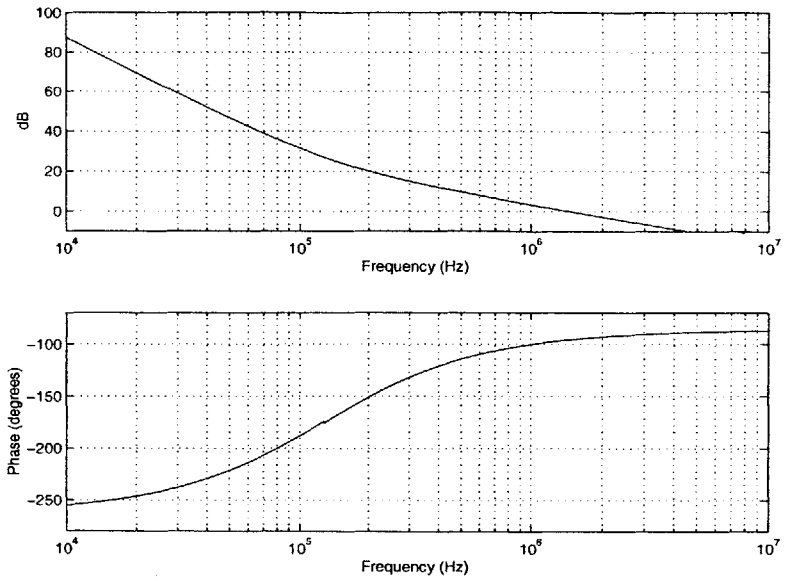

Figure 4-5: Bode plot of hand designed loop gain

decade below crossover in a third-order system actually results in 34 degrees of phase margin; however, this value is still acceptable.

While phase margin analysis is useful, it does not guarantee that the loop is stable, because the transfer function is not monotonically decreasing. Therefore, further confirmation is prudent. The root-locus plot of Figure 4-7 shows that the system can indeed be stable. Although the integrator poles initially move into the right-half plane, they gravitate towards the zeroes for increasing K. Therefore, we can be well assured that this system is stable.

This initial loop design yields a Σ-Δ modulator with 34 degrees of phase margin and nominal noise attenuation of about 70dB, which yields an SNR of about 120dB. Note that this initial design dealt only with real poles and zeroes in order to avoid the hassle of dealing with complex ones. While the SNR can most likely be optimized by designing with complex pole locations, this initial hand design is meant only to be a starting point for design with software simulation; therefore, it is not necessary to optimize SNR at this point. Furthermore, adjusting the coefficients of the loop in software will most likely result in optimal complex poles/zeroes.

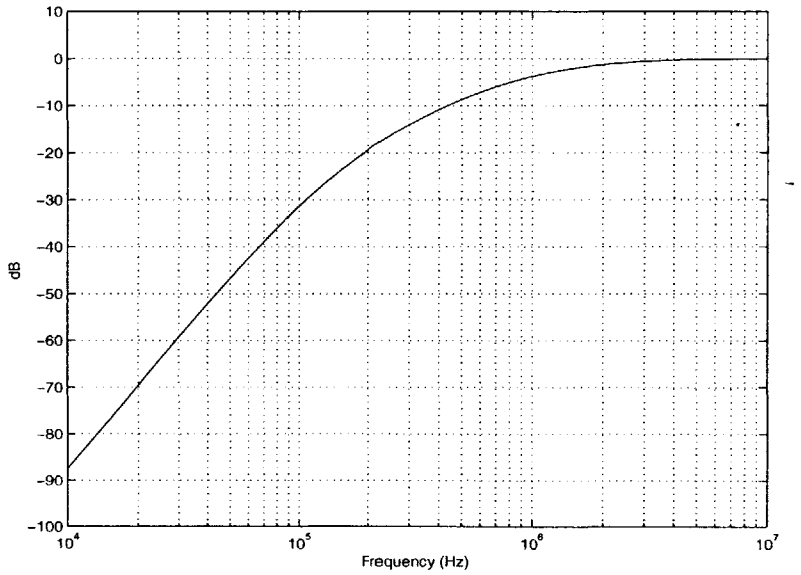

Figure 4-6: Magnitude response of the NTF corresponding to Figure 4-5

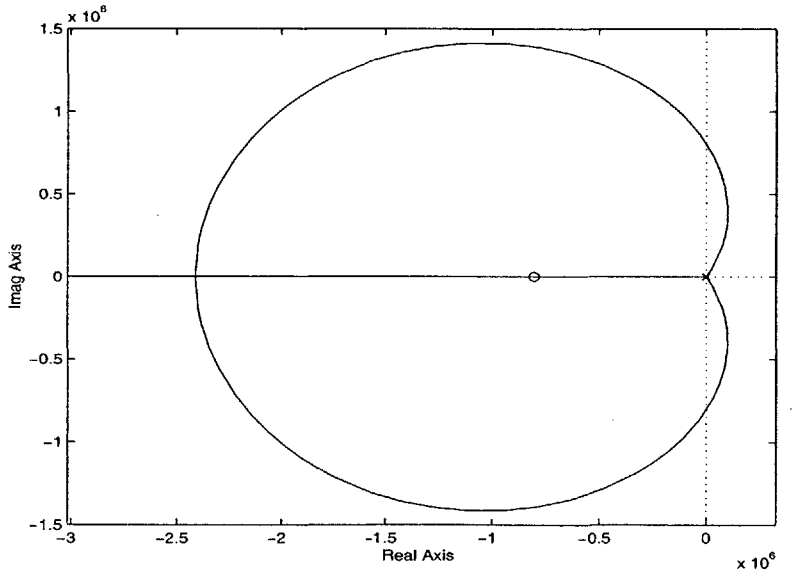

Figure 4-7: Root locus plot of the loop transmission

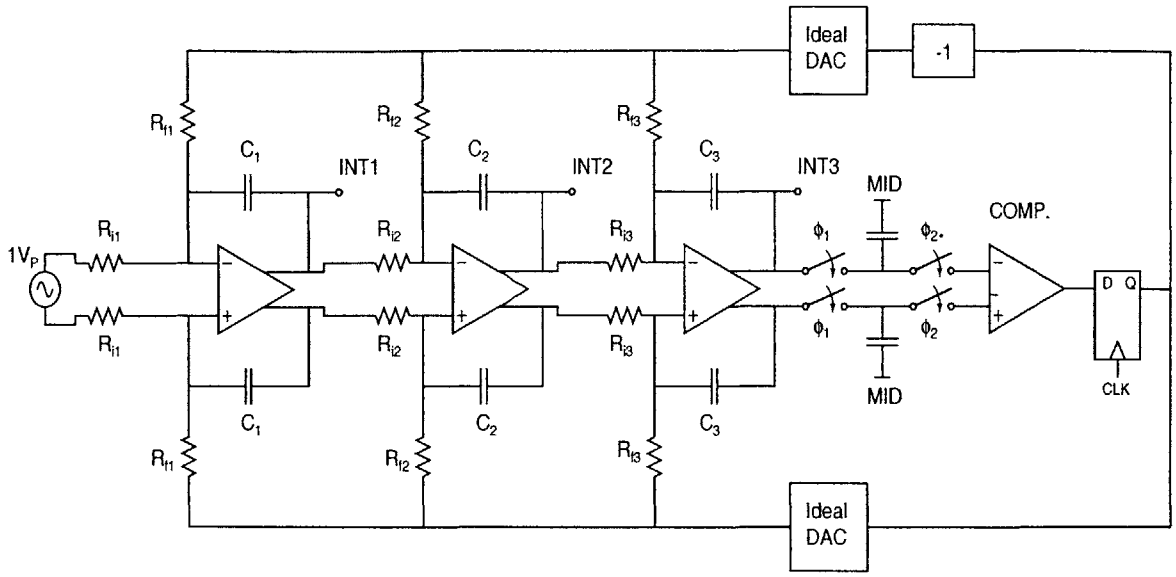

Figure 4-8: Circuit simulation of lowpass third order modulator

4.2.2 Software simulation design

Once a stable loop has been designed, the next step is to test for conditional stability. That is to say, the loop will only be stable if the input amplitude is not too large; thus, the design must be simulated with the maximum allowable input in order to verify correct operation. The circuit simulation is shown in Figure 4-8.

Since this design requires an input range of $\pm$1V full scale, the design must be simulated with at least this much voltage at the input (a $1V_p$ sinusoid is used for all transient simulations). It is easy to test for the loop's conditional stability; in order to avoid entering nonlinear operation, the opamps cannot be allowed to saturate. We can determine whether saturation has happened by probing the opamp outputs (labeled INT1-3). To prevent saturation, the integrator gains can be lowered and the feedback coefficients (b_n) can be decreased as well. Figure 4-9 and Figure 4-10 show the output of the second integrator with different values of R_{f2} to R_{i2} (whose ratio determines the feedback coefficient b_2). Altering the feedback coefficients will change the transfer function slightly, so the end result will not be exactly what our hand calculations predict; however, we should not be too far off.

Another technique that is used to keep the integrators in the linear region is to add

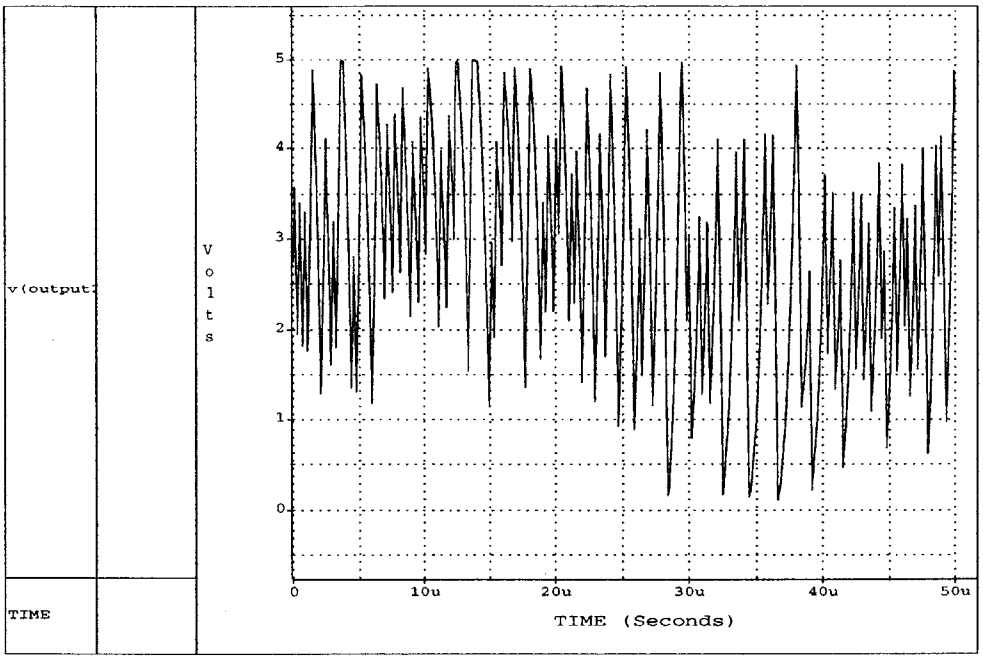

Figure 4-9: Output of integrator 2 using initial hand design

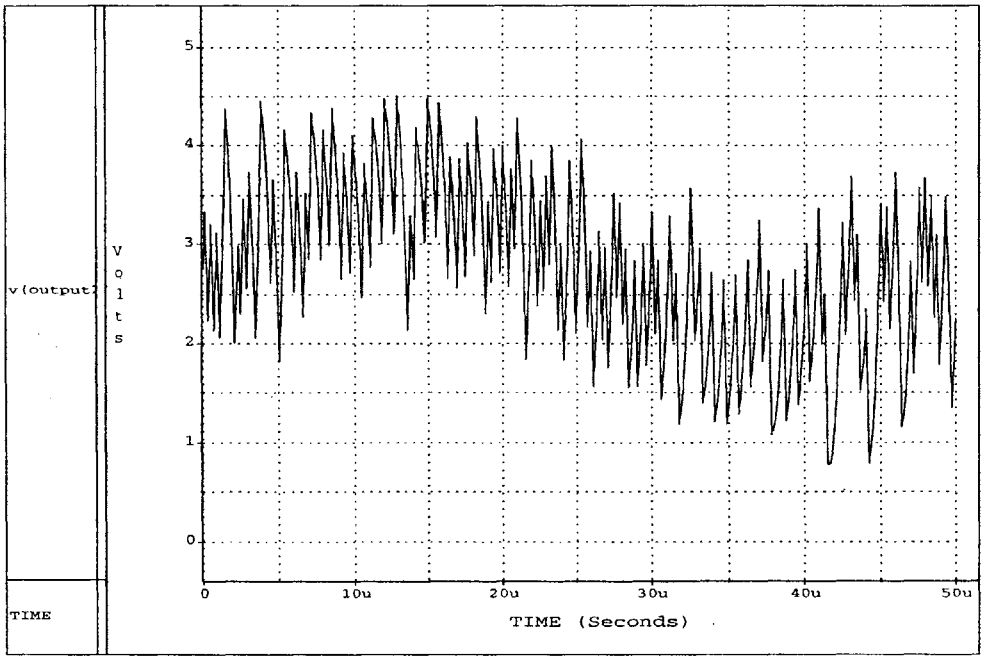

Figure 4-10: Output of integrator 2 with decreased feedback coefficient b_2

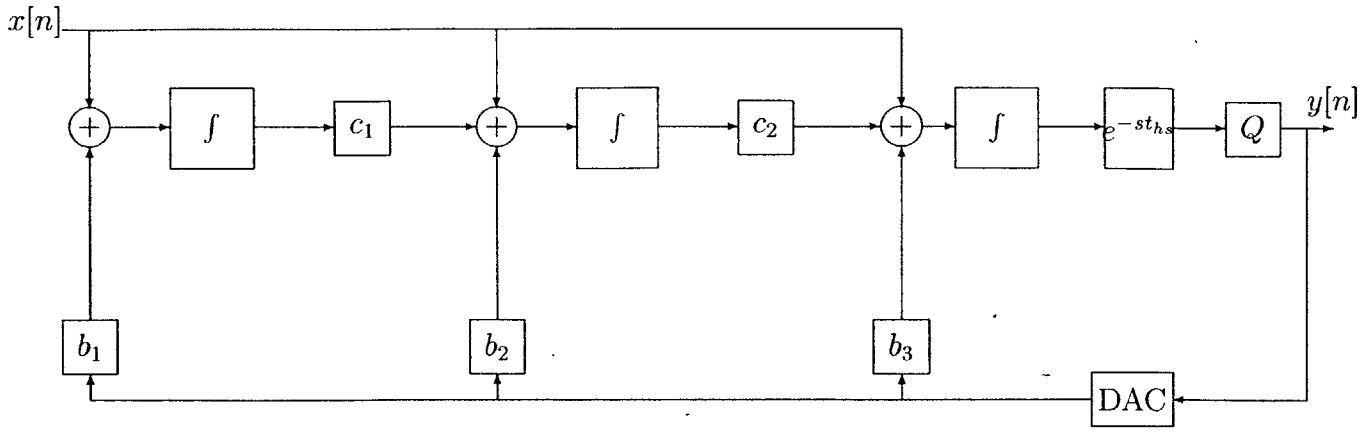

Figure 4-11: Block diagram of third order lowpass $\Sigma\Delta$ modulator with feedforward paths.

feedforward paths to the modulator topology, which has been shown to reduce internal signal swings and in some cases reduce opamp power consumption [3]. By feeding the input signal forward into the summing nodes, more signal power is injected into the loop, which has the effect of reducing the power of the feedback signals; hence, the integrator swings are suppressed [3]. The added feedforward paths are shown in Figure 4-11.

Although gain can be added in these feedforward paths to provide different swing suppressions, we will avoid those complications in this design by constraining ourselves to utilize unity-gain feedforward paths. Simulation results show that feedforward does indeed have the effect of reducing integrator swings (Figure 4-12).

Using the principles of superposition for linear systems, it should be easy to see that while the feedforward paths alter the signal transfer function somewhat (it should still approximate unity at low frequencies), they have no effect on the NTF. Hence, relieving the swings on the integrators (and thus stabilizing the loop some more) gives us more room to adjust the loop design a second time and further increase NTF attenuation. This method of simulating in software ensures that the loop is stable and the noise attenuation is near its peak.

In all these simulations, the DAC in the feedback path is realized by a unity-gain voltage-controlled voltage source. These ideal devices do not approximate the behavior of an actual D/A converter, nor do they reflect the design we will need to use for the CT modulators. The next section addresses the design of the return-to-zero

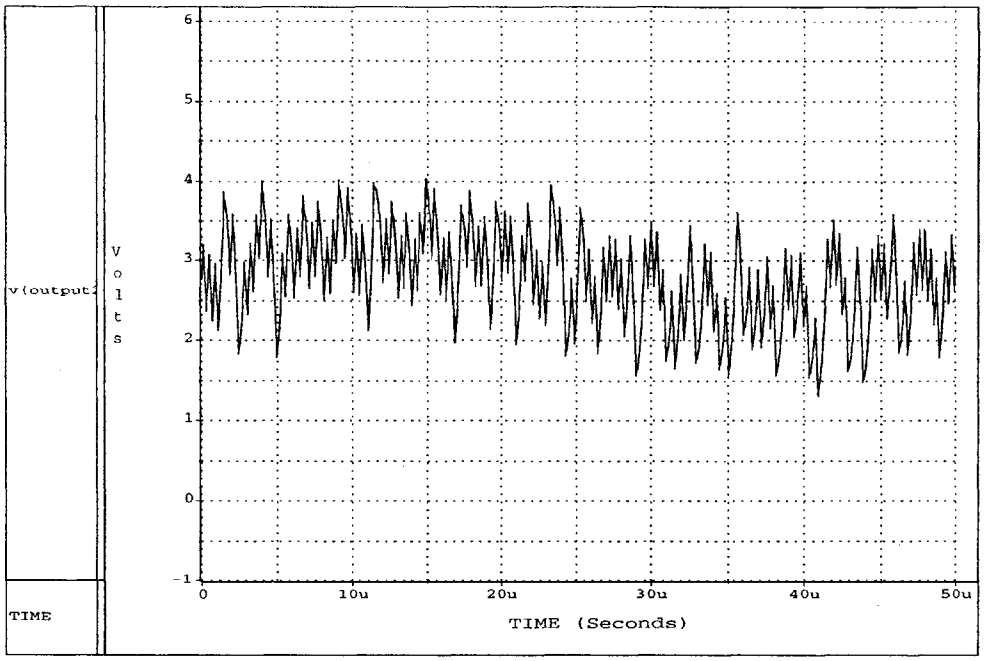

Figure 4-12: Output of integrator 2 with decreased feedback coefficient b_2 and feed-forward path from $x[n]$

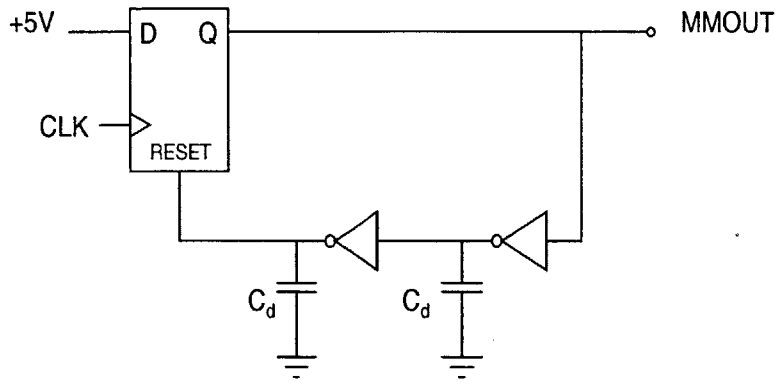

Figure 4-13: Circuit diagram of a monostable multivibrator (one-shot)

D/A converter.

4.2.3 1-bit Return-to-zero Digital-to-Analog Converter

As discussed in section 3.2.3, a RTZ pulse must return to zero before the next clock cycle. Furthermore, in this fully-differential implementation, a positive and negative pulse must be generated every clock cycle (a different polarity is needed for the two feedback paths). A simple way to implement this pulse is with a monostable multivibrator, or "one-shot." A one-shot takes a clock input and generates a short pulse when triggered by the rising clock edge. For this design, we can build a simple one-shot design that consists of a D flip-flop, a few inverters, and a few capacitors (Figure 4-13).

This digital circuit generates a short pulse, whose width is determined by the propagation delays of the inverters, which are each loaded with 30pF of capacitance; in this case, the pulsewidth is about 70ns, less than half the period of the sampling clock. This pulse, combined with some additional logic, is used to generate control signals (NSEL, PSEL, etc) for the transmission gates shown in Figure 4-14. The positive and negative pulses generated by the RTZ DAC are shown in Figure 4-15.

Since the output of the DAC returns to "zero" (2.5V is the virtual ground in this system), the feedback signals applied to the summing nodes of the modulator loop are not constant. However, in the simulations, we had assumed the DAC outputs were constant, non-RTZ waveforms. When we replace the simulated ideal DAC's with the

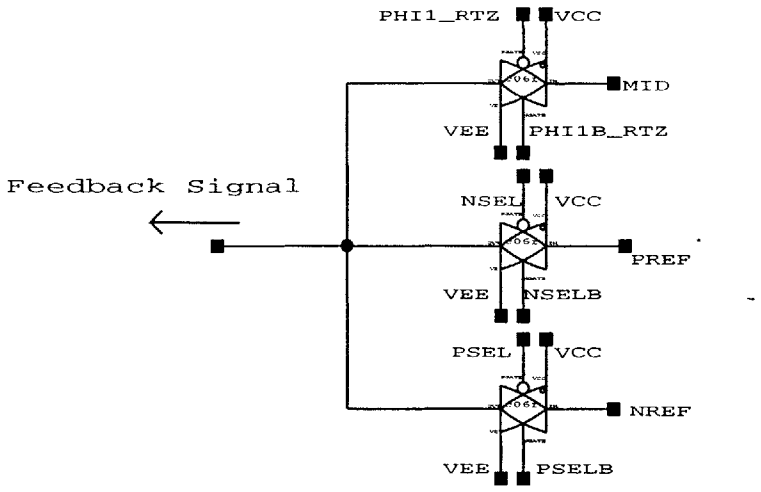

Figure 4-14: Transmission gates that switch the feedback node from NREF, MID, and PREF

RTZ DAC design, the loop characteristics will be altered and the modulator may once again be unstable. Thus, we must use the technique outlined in the previous section to readjust the loop characteristics and optimize the NTF of the modulator while remaining stable.

4.3 Design Review

4.3.1 Simulated Performance

Since the design methodology we used assures that the loop is stable for a given input range, the major things we need to simulate and measure are power consumption, silicon area, and signal-to-quantization noise ratio. While we would like to simulate the effects of thermal noise and clock jitter, the computing power necessary to model these conditions is too great; instead we must be happy with measuring these effects in the lab after the circuit is built.

With the reduced amount of digital logic and capacitor area, the size of the circuit is significantly smaller. This Σ-Δ modulator design, which we will refer to as "SD3CT," takes up 4.81mm^2 when laid out on a piece of silicon. This 50% reduction in area is a significant improvement on the DT design.

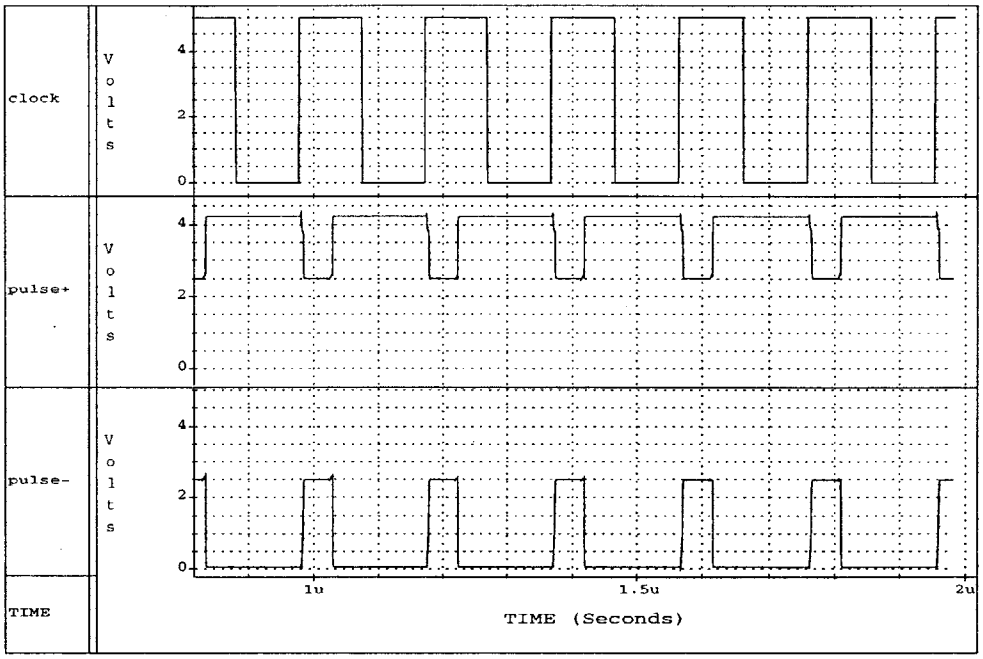

Figure 4-15: Positive and negative return-to-zero pulses generated by the DAC

However, the power consumption is somewhat larger than the DT design. Simulations show that while the power pulled from VCC is similar to the previous design, there is increased power drawn from the PREF voltage reference. The total power consumption of the SD3CT is 51.4mW. This increased power consumption can be attributed to poor design; the opamps were not optimized to draw less power, nor was any attention paid to the amount of current drawn through the PREF.

The signal-to-quantization noise ratio also falls short of the specifications. In order to roughly calculate the SNR, the NTF is simulated in SPICE, and the results are combined with Equation 2.11. The simulated NTF is plotted in Figure 4-16. At 20kHz, the noise is attenuated by 63dB; conditionally stabilizing the loop has affected the noise reduction. Therefore, the noise power can be calculated to be:

$$P_e = \int_{-f_o}^{f_o} |NTF(f)|^2 k_e^2 df = \int_{-f_o}^{f_o} |-63dB|^2 \times \frac{\Delta^2}{12} \times \frac{200Hz}{5.12MHz} \tag{4.4}$$

$$P_e = 4.07 \times 10^{-11} V^2$$

The maximum input amplitude is $1V_p$, so the SQNR is calculated to be:

$$SQNR = \frac{P_s}{P_e} = \frac{1}{\sqrt{2}} \times \frac{1}{4.07 \times 10^{-11}} \approx 102dB \tag{4.5}$$

This calculation is significantly less than our target of 111dB, especially when we consider that we have assumed the quantization noise is white. Due to this assumption, practical Σ-Δ modulators often exhibit a SNR which is lower than the theoretical value. To verify this effect, we can measure the SNR of the SD3CT by running a transient simulation in SPICE and using power spectrum estimation techniques on the output to determine the level of quantization noise. The sampling frequency is 5.12MHz, so to obtain a noise bandwidth of 80Hz, the simulation must yield 64000 points. The modulator only outputs one bit every 200ns; therefore, the simulation needs to run for 12.8ms (longer if power spectrum averaging is desired). For a circuit which uses transmission gates to sample data as well as digital logic to implement return-to-zero pulses, HSPICE will need more than two weeks to complete the simulation.

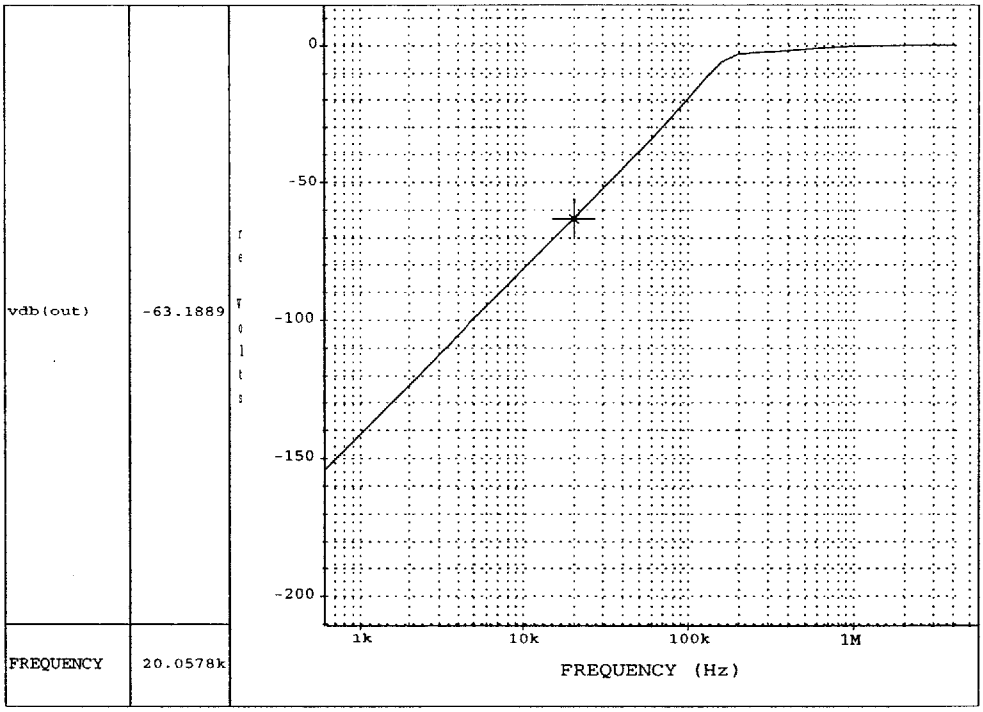

Figure 4-16: Plot of SD3CT Noise Transfer Function

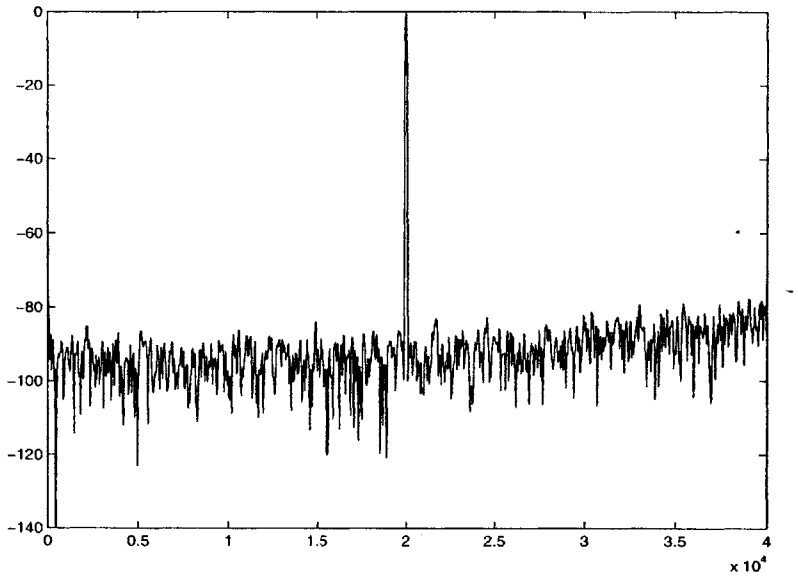

Figure 4-17: Power spectrum of SD3CT output, normalized to P_s=1V^2

This difficulty is a frequent problem in designing Σ-Δ modulators, for DT and CT systems alike. The problem is usually solved by employing the use of a high-level software intended to quickly simulate Σ-Δ modulators for long transient times. Unfortunately, such software is not available to this designer; as such, a fast circuit simulator, Avanti StarSim, is used instead. StarSim is optimized to quickly carry out long transient simulations using a different set of algorithms than HSPICE. Even with StarSim, transient simulation of the SD3CT takes about 2 days.

Once the StarSim simulation is finished, the results are read into MATLAB. The output signal is read into a vector, which takes a value every sampling period and quantizes the result into either "0" or "1." By using the "psd" command, the power spectrum can be obtained. Figure 4-17 shows the output power spectrum.

Determining the SNR from this graph is an easy task, as the plot has been normalized so that the signal power is unity. By inspection, the quantization noise floor is at 95dBc ($3.2\times10^{-10}V^2$). Dividing by the noise bandwidth (80Hz) and integrating from 19.9kHz to 20.1kHz yields the noise power in the 200Hz band. Assuming that the quantization noise is relatively flat in this 200Hz region, the normalized noise power can be calculated to be:

$$P_e = (3.2 \times 10^{-10} V^2) \times \frac{200Hz}{80Hz} = -91dBV^2 \tag{4.6}$$

Since the normalized signal power is unity, it follows that the SQNR is 91dB. This result indicates that a third-order lowpass modulator is not capable of enough noiseshaping to achieve the desired SNR.

4.3.2 Weaknesses

The previous discussion has identified two major weaknesses in the design; first, the power consumption has not been diminished and more importantly, there exists too much quantization noise at the output. A third weakness involves the RTZ DAC. While using a one-shot may cut down on the effects of clock jitter, the pulse width is not stable over temperature or process variations. Since the propagation delay of each inverter depends greatly on the load capacitor's size, temperature fluctuations can cause varying feedback pulsewidths, which destabilize the AC gain of the loop. Moreover, the signal-to-noise ratio may also have a larger temperature dependence as a result of this dependence.

Even before building the SD3CT, it is apparent that it falls far short of meeting several major specifications. However, it does demonstrate that a CT Σ-Δ modulator takes up less space than its DT counterpart, and the effects of clock jitter from a monostable multivibrator RTZ pulse can still be studied if the chip is built. The next step in designing a CT Σ-Δ modulator is to examine techniques to overcome the obstacles encountered in this design.

Chapter 5

Design of a fourth-order continuous-time Σ-Δ modulator

The major flaw of the Σ-Δ modulator presented in the previous chapter was the inability of the loop to shape the quantization noise to an acceptable level. This chapter explores a few ways to solve this problem; ultimately, we choose to design a fourth-order loop. Also, we will discuss redesign of the RTZ DAC and taking advantage of the continuous-time (CT) topology to minimize power consumption.

5.1 Techniques to decrease quantization noise

5.1.1 Increasing sampling frequency

The SQNR can be improved two ways: increasing the oversampling ratio and increasing the noise attenuation. We will first examine the effects of changing the sampling frequency. Recalling the discussion in Chapter 2, doubling the sampling frequency (and thus the oversampling ratio) will increase the SQNR by approximately 6dB. In order to add another 30dB to the SQNR of the SD3CT, the sampling frequency would have to be increased to 163.84MHz!

However, since increasing the sampling frequency changes the amount of phase that the SHA contributes to the loop, the stability of the modulator becomes better

as a higher sampling frequency is used. Doubling f_s enables the crossover frequency to be placed at 2.56MHz (roughly a third of a decade); it follows that the zeroes can be placed at 256kHz. Since we are using a third-order modulator, the initial slope of the Noise Transfer Function (NTF) is 60dB/decade. Thus, the amount of noise shaping at f_c can increase about 20dB! Thus, doubling f_s can yield a significant improvement in SQNR. To achieve 111dB of SNR with a third-order modulator, we would probably have to increase the sampling frequency by a factor of 3.

Unfortunately, changing the sampling frequency is not an easy task. Since the sampling clock is generated by a phase-locked loop, the sampling frequency is limited to the range of the VCO used in the loop. The VCO used in this system is not designed to generate a frequency significantly higher than previously required; therefore, a new VCO would need to be designed in order to make the improvements discussed here.

5.1.2 Fourth-order topology

Another way to improve SQNR is to design a fourth-order modulator, which can shape the quantization noise to a lower level. A fourth-order modulator's NTF would have an initial slope of 80dB/decade. Assuming we can place the three loop transmission zeroes at a frequency of 128kHz as in the SD3CT, the attenuation at 20kHz should be approximately 98dB. Our initial design of the SD3CT predicted nearly 70dB of noise shaping; we can therefore conclude that a fourth order design can achieve an extra 28dB of SQNR. This improvement is barely enough to achieve our goals.

There are some drawbacks to this approach. First, adding another integrator complicates the method of stabilizing the loop. With higher-order loops, guaranteeing conditional stability becomes much more difficult, and we may need to explore new techniques (such as resetting integrators) to do so. Secondly, another integrator means increased power consumption. While we can still hope to use less power than a discrete-time (DT) modulator, a fourth-order design will most likely draw more power than a third-order design. Finally, we will need to increase the area of the chip to include a fourth integrator.

The desirable approach is to design a new VCO and increase the sampling frequency, as doing so will conserve power and area. However, since we will most likely need to design a VCO to meet clock jitter requirements later, it makes sense to wait until a suitable phase noise specification is made before starting a VCO design. Since obtaining such a spec would take too long, we choose to explore the design of a fourth-order modulator at this point in time.

5.2 Design methodology

The same design methodology discussed in the previous chapter is applied to designing a fourth-order lowpass modulator. Starting with an initial hand design to ensure loop stability, HSPICE simulations are run, providing a basis to stabilize the loop by altering the integrator gains and feedback coefficients. As before, we will begin by discussing the initial design.

5.2.1 Initial hand design

Since the sampling frequency remains the same, the crossover frequency of the loop transmission must still be no greater than 1.28MHz. With four integrators, the phase at low frequencies is -360 degrees. The zeroes need to contribute 270 degrees of phase recovery at 1.28MHz in order to achieve 45 degrees of phase margin (due to the SHA). Thus, we place all three zeroes at roughly 128kHz (slightly lower if we wish to achieve full phase recovery). Figure 5-1 shows the bode plot of the loopgain, and Figure 5-2 shows that the NTF is nearly -90dB at 20kHz. The root locus plot in Figure 5-3 verifies that the system is stable, since all the origin poles move into the left half plane as DC gain increases.

Using this design, we can once again use the techniques discussed before (feedforward paths, software simulation) to ensure conditional stability for an input range of $1V_p$.

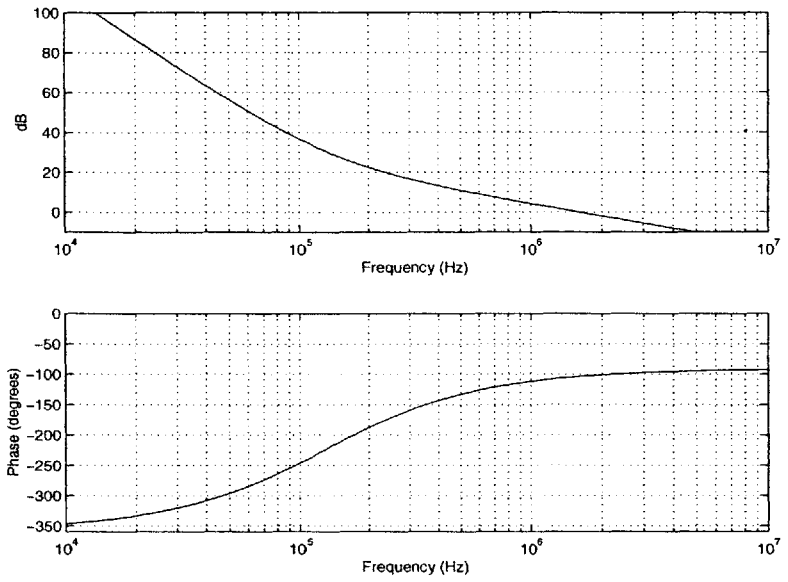

Figure 5-1: Bode plot of the hand designed fourth-order loop gain

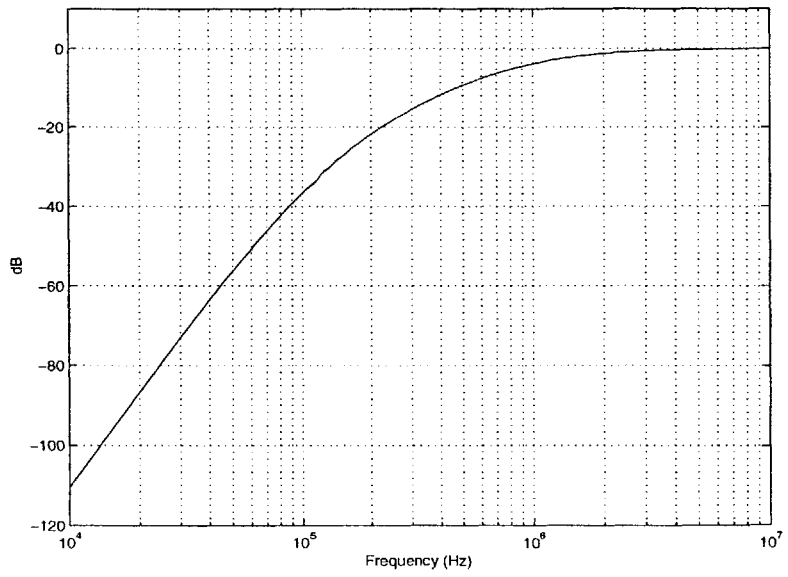

Figure 5-2: Frequency Response of the NTF corresponding to Figure 5-1

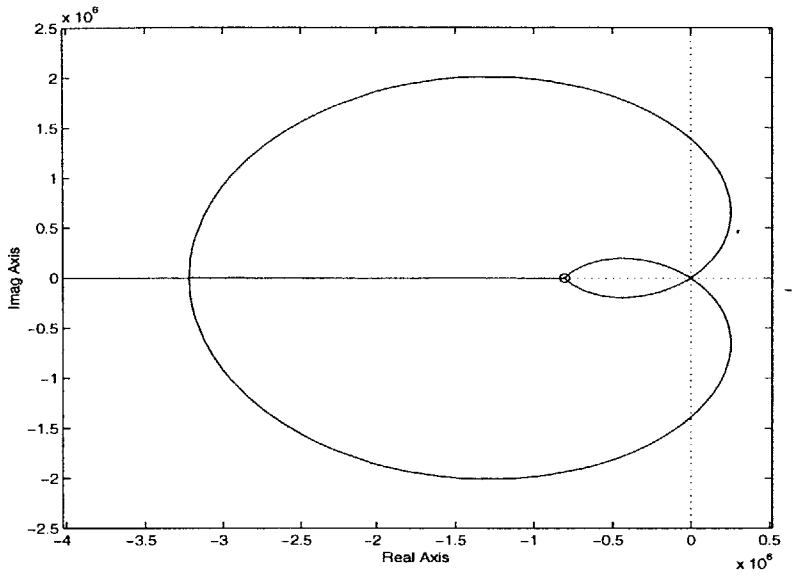

Figure 5-3: Root locus plot of the fourth order loop transmission

5.2.2 Minimizing power consumption

As discussed in Chapter 3, the main way we can conserve power with a CT design is by using opamps which do not need to supply as much output current. Although it is desirable to design opamps designed for this new set of requirements, limited time prevents us from doing so. Instead, we will set the second, third, and fourth integrators to draw a smaller bias current; we keep a large current running through the first integrator to quiet the first stage's opamp noise.

As long as the opamps maintain a relatively high gain and bandwidth, the performance of the modulator should not suffer. Also, the slew rate of the opamp must not be limited—the opamp needs to be able to drive the feedback capacitance in order to avoid entering the nonlinear region. The nominal set current of these CMOS opamps is 30μA ; simulations show that 12μA is the smallest set current that does not effect change in modulator performance. To give ourselves a safety margin, we bias the opamps at 15 μA.

Another problem identified in the SD3CT was the additional current drawn from the voltage reference (PREF). Further inspection shows that the easiest way to reduce this current is to increase the size of the feedback resistors. We must also increase

the size of the other resistors to maintain the feedback coefficients we picked in our initial hand design. Likewise, this change requires the capacitor size to be decreased, which helps to further minimize silicon area.

5.2.3 RTZ DAC redesign

Since we are very much interested in a pulsewidth that is insensitive to temperature or process variations, we take a different approach to designing the return-to-zero DAC. Instead of using a "one-shot" to generate the RTZ pulse, this design generates a pulse using the falling edge of the sampling clock. Thus, the RTZ pulse lasts for exactly half the sampling period.

Generating this pulse is easily done by using the output of the modulator and the clock signal as two inputs to a NAND logic gate. When the clock signal is high, the return-to-zero pulse is activated; the output of the NAND (and its inverse) are used as the signals to a transmission gate connecting the voltage reference to the feedback node. When the clock is low, the pulse turns off and the feedback node is switched back to virtual ground (MID); the clock signal and its inverse are used as the control signals for that transmission gate. The output of the modulator determines which transmission gates (PREF or NREF) are switched on during the first half of the clock cycle. Because there are two possible different transmission gates for each feedback path, we will actually need two NAND gates to generate all the control signals. Figure 5-4 illustrates this circuit and Figure 5-5 shows the waveforms.

We must still worry about possible shorts from PREF to NREF; therefore, we must use a "break-before-make" generator as before to generate control signals that avoid switching two transmission gates on at the same time. In this case, we use the sampling clock as the input to the BBM generator (see Appendix A).

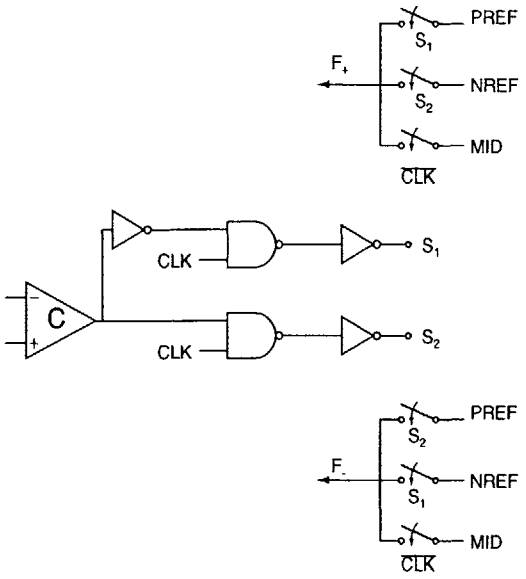

Figure 5-4: Circuit implementation of the redesigned DAC

5.3 Design Review

5.3.1 Simulated Performance

Even though an extra integrator was added, redesigning the resistor and capacitor sizes resulted in more area conservation. The layout for this modulator, which we will refer to as "SD4CT," takes up about $5.15mm^2$ of silicon. This area is a little more than the SD3CT but is still a significant improvement from the DT modulator.

HSPICE simulations show that the power usage of this circuit also exhibits a marked improvement. By decreasing the set currents on the integrators and increasing feedback resistors to limit the current drawn from PREF, the power is decreased considerably. HSPICE simulations show that 5.0mA is drawn from VCC (5V) and 195μA is drawn from PREF. The total power consumption is estimated to be 26mW, compared to the 40mW used by the DT modulator. This 35% reduction is an improvement from that of the SD3CT as well.

We calculate SNR using the same techniques as before. Figure 5-6 shows that the quantization noise is attenuated 83dB at 20kHz (the resonance at crossover indicates a low phase margin, though we can probably tolerate it). Hence, the noise power is given by:

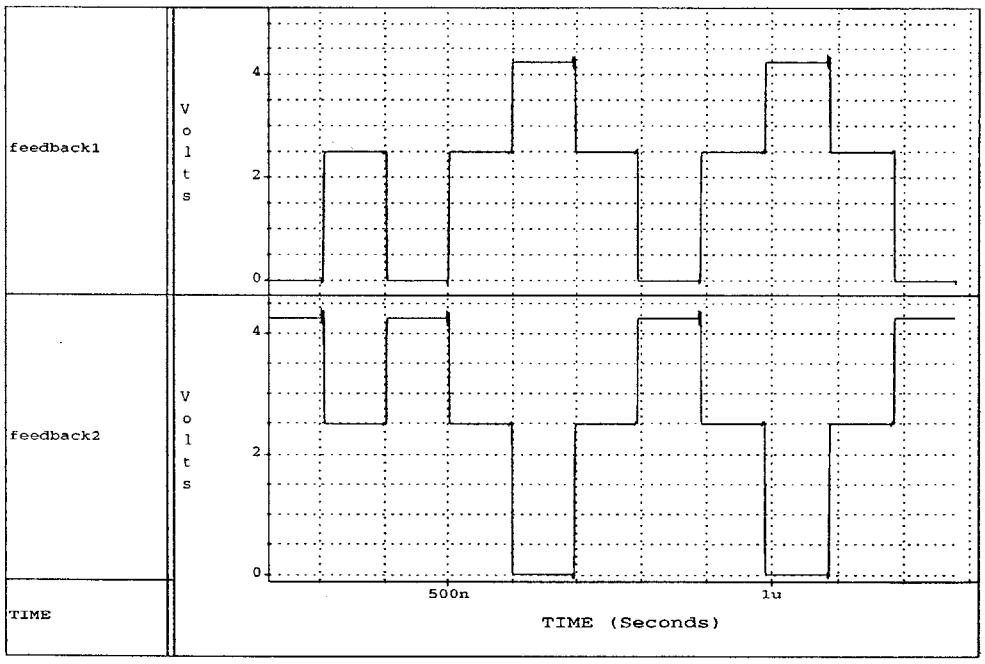

Figure 5-5: Return-to-zero waveforms for the redesigned DAC

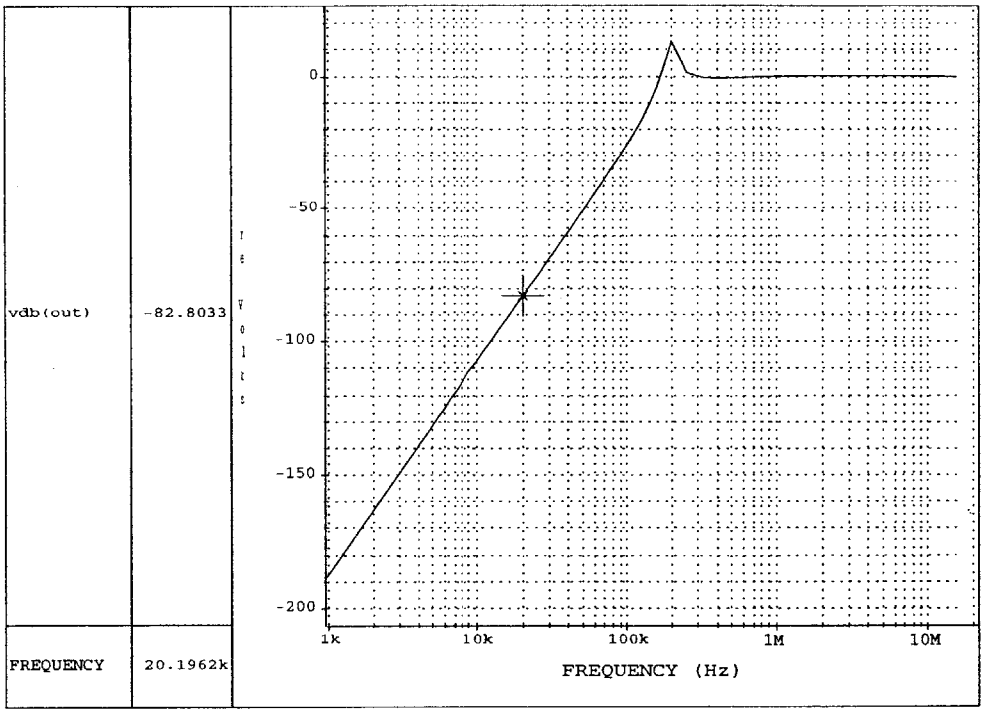

Figure 5-6: Magnitude response of the SD4CT noise transfer function

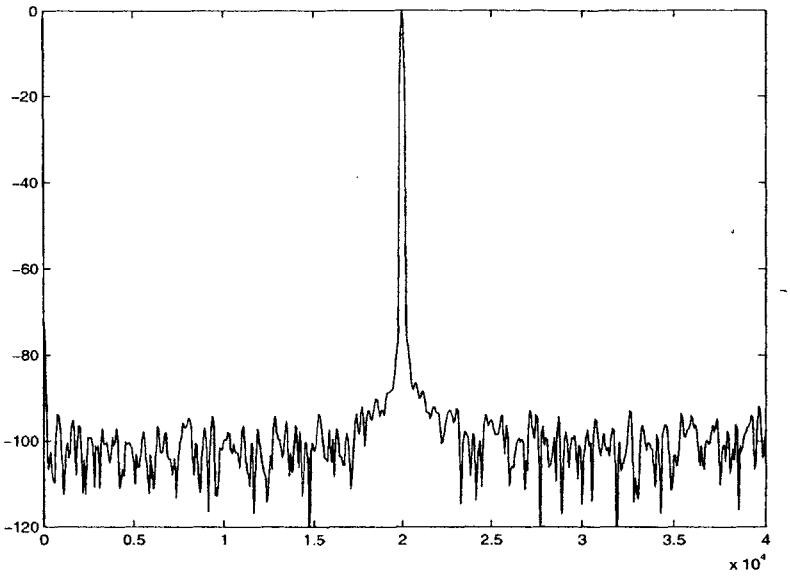

Figure 5-7: Power spectrum of the SD4CT output, normalized to P_s=1V^2

$$P_e = \int_{-f_o}^{f_o} |-83dB|^2 \times \frac{\Delta^2}{12} \times \frac{200Hz}{5.12MHz} = 4.08 \times 10^{-13}V^2 \tag{5.1}$$

With a full scale input of $1V_p$, the SNR is calculated to be 122.3dB. As before, we can expect the actual modulator to exhibit a SNR which is somewhat lower than this value. To verify this expectation, we again use StarSim to generate transient information, and use MATLAB to determine the power spectrum of the output, shown in Figure 5-7.

Surprisingly, the noise floor does not appear to be significantly lower than that of the SD3CT. By inspection, the noise power appears to be -105dBc($3.2\times10^{-11}V^2$), which corresponds to a SQNR value of:

$$SQNR = \frac{P_s}{P_e} = \frac{1}{(3.2 \times 10^{-13}V^2) \times \frac{200Hz}{80Hz}} = 101dB \tag{5.2}$$

This SQNR is not the improved value we had projected for a fourth-order modulator, nor does it coincide with the value calculated from the simulated NTF. There are few possible explanations for this result. First, the actual quantization noise may become more colored as the order of the modulator increases. With a fourth order modulator, the white-noise approximation used to estimate SNR may break down

further than in the third-order case; thus, the quantization noise power around the 20kHz band may be much larger than anticipated.

Secondly, the simulator may be giving erroneous results. One reason for this occurence may be that this simulator does not model circuit noise. Without circuit noise at the input of the comparator, the quantization noise may be more correlated to the input frequency (or other frequencies) than it would be for a real circuit. Such a correlation would increase the noise power in the band of interest as well.

5.3.2 Analysis

With the exception of the questionable power spectrum results, the SD4CT performs better than the SD3CT, and it also demonstrates some improvements over the DT design as well. However, even if the simulation results are flawed, the Σ-Δ The SQNR can be imp modulator will almost surely fall short of the 122.3dB value, though it is extremely difficult to predict exactly how much. As such, it is possible that when built, this modulator may not outperform the DT modulator's SNR of 111dB.

Even so, this design demonstrates that comparable performance can be achieved using significantly less power and silicon area, provided that a low jitter clock is used to drive the SD4CT. The viability of this design hinges on our ability to design a low phase-noise VCO capable of achieving the sampling frequency of 5.12MHz. Lab testing of the manufactured IC's should provide us with all we need to determine whether or not a continuous time modulator can be used for this application.

Chapter 6

Measurement Results of On-chip Continuous Time Σ-Δ Modulators

6.1 Measured performance

6.1.1 SD3CT measurements

Measuring the performance of the continuous-time (CT) Σ-Δ modulator chips is similar to measuring performance in simulation. In order to determine the signal-to-noise ratio of the modulator, the power spectrum of the output must be measured. The rise and fall times of the output transitions makes using a spectrum analyzer undesirable; therefore, we simply acquire each individual sample of the output into software by using a data acquisition card interfaced with a PC. Once several thousands of samples of the output are contained in software, simple power spectrum estimation algorithms (in software such as MATLAB) can be used to obtain an accurate power spectrum.

Using this measurement technique, the SD3CT is measured to have a full scale SNR equal to 95dB (this value is obtained when using a low-jitter clock reference to a crystal). This measurement is slightly better than what the software simulation predicted; however, because part of these measurements is obtained by visual inspection of the power spectrums, some slight discrepancies should not be surprising. The

Characteristic	Specification	Simulated	Measured
Clock Frequency	$256^* f_o$	—	—
Input Range	$> \pm 1V_p$ Full scale	$\pm 1V_p$	$\pm .8V_p$
Full Scale SNR	110dB	91dB	95dB
Second Harmonic Distortion (from 10kHz input)	< -80dBcFS	—	-91.8dBcFS
AC Gain Stability at 20kHz	1% over temperature	—	15% over temperature
Spurious Tones	< -80dBcFS at 20kHz$\pm$4kHz	No tones in output spectrum	No tones in output spectrum
Power Consumption	—	51.4mW	52.2mW
Chip area	—	—	4.81mm^2

Table 6.1: Measured performance for A/D conversion using SD3CT

remaining test measurements are detailed in Table 6.1.

Another discrepancy we see with the SD3CT is that the input range is slightly less than the specification. We define the full scale input range to be the maximum amount of voltage swing on the input that maintains the same amount of noise shaping (i.e. prevents the integrators from saturating). Since we carried out all of our simulations using a $1V_p$ input sine wave, it is likely that process variations or StarSim errors resulted in larger internal voltage swings when the design was built in silicon. In the future, it would be wiser to use a slightly larger voltage (like $1.2V_p$) in design simulations to provide a margin for such error.

Another concern of the SD3CT is its sensitivity to component fluctuations. Though it is difficult to measure the effects of component tolerances, it is a simple matter to make measurements over a wide temperature range. Temperature change will alter the absolute component values; hence, we will be able to gain insight into how such component changes affect the output power spectrum. At the temperature extremes of 85 and -40 degrees Celsius, the quantization noise power density around the carrier frequency changes by about 5dB (the power spectrums at different frequencies can be found in Appendix B).

This behavior seems to indicate that component tolerances do indeed affect the

Characteristic	Specification	Simulated	Measured
Clock Frequency	$256*f_o$	—	—
Input Range	$> \pm 1V_p$ Full scale	$\pm 1V_p$	$\pm .9V_p$
Full Scale SNR	110dB	101dB	99dB
Second Harmonic Distortion (from 10kHz input)	< -80dBcFS	—	-92.8dBcFS
AC Gain Stability at 20kHz	1% over temperature	—	2.2% over temperature
Spurious Tones	< -80dBcFS at 20kHz$\pm$4kHz	No tones in output spectrum	No tones in output spectrum
Power Consumption	—	26.0mW	28.4mW
Chip area	—	—	$5.2mm^2$

Table 6.2: Measured performance for A/D conversion using SD4CT

noise transfer function of the modulator; however, considering the effects of components on a bandpass design (as shown in Figure 4-1), the CT lowpass topology proves to be less susceptible to fluctuations in component values.

The last major problem we observe is the that AC gain stability is awful. As discussed in Section 4.3.2, the varying pulsewidths of the RTZ DAC have a large effect on AC gain. Over temperature, the SD3CT displays a 15% change in scale factor! This result emphasizes the need for stable pulsewidths in the feedback DAC and confirms that a one-shot is a poor way to implement the DAC.

6.1.2 SD4CT measurements

The SD4CT performance data is obtained in a similar manner as before. As with the third-order modulator, the SD4CT exhibits the same limited dynamic range as before. One marked improvement is the 44% decrease in power from the SD3CT, which we had predicted. The performance data is shown in Table 6.2.

Another thing we notice right away is that the AC gain stability is much improved, as the SD4CT does not use a one-shot to implement the DAC. Unfortunately, it is still not as good as we would like; the gain instability can most likely be attributed to the effects of temperature on component values, which alter the dynamics of the

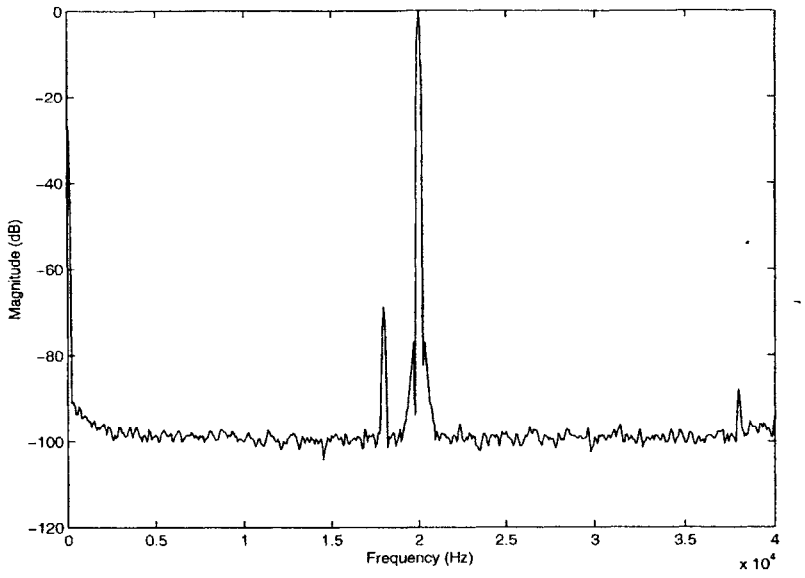

Figure 6-1: Nominal output spectrum of SD4CT

noise shaper, and thus the scale factor.

The transient simulations from the previous chapter indicated that quantization noise floor of the SD4CT modulator is higher than expected. Since this result conflicts with the AC simulations done on the same circuit, we put forth the possibility that the noise was due to some sort of simulation error. Unfortunately, the simulated behavior is verified by our lab tests. The SD4CT output spectrum is shown in Figure 6-1. It is consistent with the simulated spectrum shown in Figure 5-7.

Since the spectrum confirms that the quantization noise is indeed shaped at higher frequencies (Figure 6-2), it appears that there is some noise floor at -100dBc. There are several possibilities; first, the flat response at low frequencies could be caused by thermal noise on the input or by noise generated from the lab equipment. Second, the integrators could be saturating and introducing a nonlinearity in the modulator loop, thus decreasing the effectiveness of the noise shaper at low frequencies. As the SNR of this modulator falls fall short of expectations, it is worthwhile to investigate these possibilities.

Circuit noise in the modulator

In the DT modulator design, the noise floor is presumed to have originated from

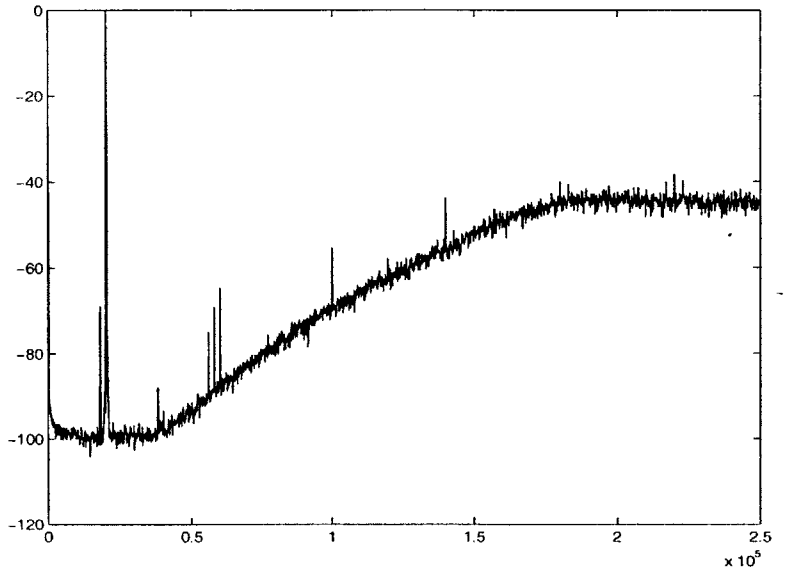

Figure 6-2: Nominal output spectrum of SD4CT (zoomed out)

the thermal noise of the input capacitors. As previously discussed, replacing the capacitance with a resistance should significantly cut down on the thermal noise. With an input amplitude of .9V_p, the noise power density that corresponds to -100dBc is calculated to be 775 $nV/\sqrt{Hz}$. The first stage input resistors are 20K; the thermal noise density due to them is calculated to be 18 $nV/\sqrt{Hz}$ from Equation 3.5. Likewise, the worst cased referred-to-input noise of the integrator opamp (biased at 30μA) has been verified to be 92 $nV/\sqrt{Hz}$. These estimates lead us to believe that this noise floor is not due to thermal noise.

Other possible noise sources include the power supplies and voltage references used to operate the chip. While it is virtually impossible for noise from VCC (5V) to show up on the Σ-Δ modulator output (due to power supply rejection in both the opamps and the modulator loop), noise from the voltage reference will appear directly on the input, as would noise from the virtual ground of MID (2.5V). While passive RC filters serve to attenuate the noise around 20kHz, noise may still be able to feed through to the output. The supply noise proves not to be the problem however, as the noise density for both supplies are measured to be about $25nV/\sqrt{Hz}$ at 20kHz.

Circuit noise is most likely not the cause of the unexplained noise floor in the output spectrum. Further evidence to support this assertion is the fact that the

-100dBc noise floor shows up in simulation, where circuit noise is not modeled.

Integrator saturation

While transient simulations show that the integrators do not saturate, it is worthwhile to verify this behavior in the real world. Fortunately, test pads were placed on all of the opamp outputs during layout of the SD4CT chip. By probing these pads, we can check to see if saturated integrators are introducing nonlinearities into the noise shaper.

The outputs of the integrators are exactly what we expect them to be when a 20kHz sinusoid is fed through the modulator. This result eliminates the possibility of nonlinearities injecting low frequency noise into the output and further reinforces our faith in the simulations. Oscilloscope pictures of each integrator output can be found in Appendix C.

While the cause of the noise floor has not been discovered at the time of this writing, we have narrowed the possibilities down somewhat. The AC simulations show that the noise shaper does attenuate noise at lower frequencies; transient simulations of the noise shaper without the nonlinear quantizer (comparator) or the feedback DAC also confirm this effect. Because this noise appears in simulation, we can be reasonably certain that the noise is not due to thermal noise or other problems with the laboratory setup. Hence, the noise must be caused by either the RTZ DAC or the comparator. Future work is necessary to determine the actual cause.

6.2 Effects of Clock Jitter

One of the major factors in determining the viability of a CT Σ-Δ modulator for this particular application is the amount of noise that clock jitter will introduce into the system. Since it is difficult to derive a quantitative relationship between the amount of phase noise on the sampling clock and the amount of noise that appears on the modulator output, perhaps the most effective method of predicting the effects of clock jitter is to empirically determine the relationship. This section concerns itself

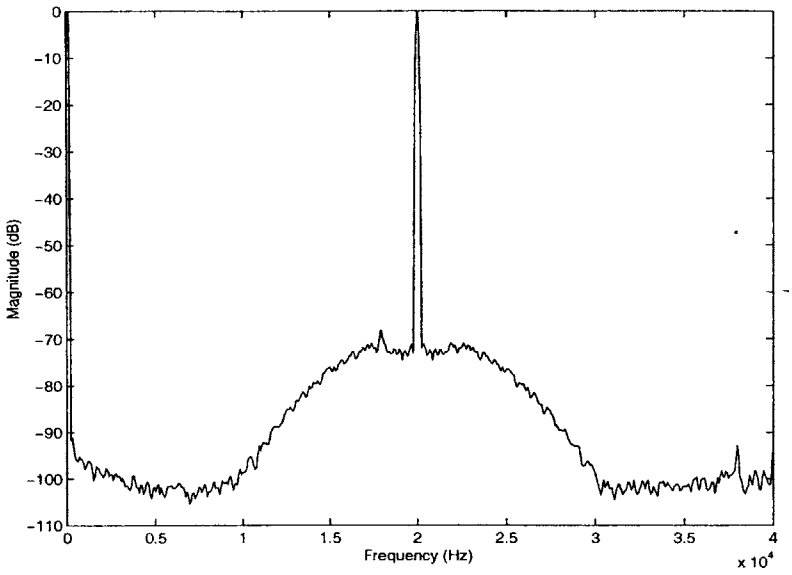

Figure 6-3: Output spectrum of SD4CT using a clock with white jitter (bandlimited to 3kHz)

with examining how phase noise appears on the modulator output and arriving at a maximum allowable phase noise requirement to achieve the SNR we would like with a CT modulator.

6.2.1 Phase noise modulation

The first thing we do in examining the effects of jitter is to get a good qualitative feel of how phase noise shows up at the output; thus, the output of the Σ-Δ modulator is observed when clocked with a 5.12MHz square wave phase modulated with white noise. If the noise is large enough, the effects of phase noise are visible on the Σ-Δ modulator output. In fact, it appears that the phase noise of the sampling clock is *modulated* onto the carrier frequency of the input! Figure 6-3 shows the output spectrum of the modulator when the clock is phase modulated with white noise bandlimited to 3kHz. The noise modulated on the carrier frequency is also white noise bandlimited to 3kHz.

To verify this phenomenon, we phase modulate sinusoids of different frequency onto the sampling clock and observe the output of the modulator with these clocks. The resulting spectrums can be found in Appendix D. They indicate that the phase

noise is indeed directly modulated onto the carrier frequency. For example, phase modulating 1kHz onto the clock results in noise concentrated at 1kHz away from the carrier frequency of 20kHz. The same effect takes place with sinusoids of different frequencies. Moreover, the phase noise does not appear to be shaped; that is, white phase noise appears as white noise on the output.

Hence, the only phase noise of interest is that contained within $\pm$100Hz from the sampling frequency. This noise should be as small as possible if the maximum SNR is to be obtained. Therefore, our goal is to determine the maximum amount of phase noise *power* that can be tolerated in that frequency band.

This maximum phase noise requirement is determined in the following manner. By introducing white phase noise into the sampling clock, we can observe the amount of noise that appears on the output. By varying the power density of this white phase noise, we can ascertain the amount of phase noise that results in negligible noise at the output. One problem with this test is that the quantization noise floor of both the SD3CT and the SD4CT is well above the specified limit, and it is impossible to "see" the noise from clock jitter below that floor. Therefore, we must extrapolate the available data to obtain a rough estimate of the amount of phase noise we can tolerate to obtain 111dB SNR.

While phase noise is usually expressed as the noise power contained at some offset from the carrier (e.g. 70dBc at 10kHz), this representation does not accurately describe the phase noise requirement. We are not concerned with the noise at a single discrete frequency; instead, we are interested in the noise within a specified band. Hence, the phase noise specification is expressed in phase noise power within $\pm 100Hz$ from the sampling frequency (excluding the power at 0Hz).

6.2.2 Clock jitter tolerances for SD3CT and SD4CT

The phase noise density is altered at regular intervals while the output spectrum of the modulator is monitored. The phase noise is white in order to obtain a flat frequency response, making it easier to estimate the noise power in a specified frequency band. The data obtained from the SD3CT is shown in Table 6.3.

Phase noise density ($mV/\sqrt{100Hz}$)	Noise floor about f_c ($dBV^2/80Hz$)
100	-72 dBc
50.0	-79 dBc
41.1	-80 dBc
32.3	-83 dBc
24.0	-85 dBc
12.2	-90 dBc
7.2	-93 dBc
3.9	-95 dBc

Table 6.3: Relationship between clock jitter and output noise floor for SD3CT

Phase noise density ($mV/\sqrt{100Hz}$)	Noise floor about f_c ($dBV^2/80Hz$)
100	-72 dBc
52.0	-77 dBc
39.0	-80 dBc
29.2	-83 dBc
23.9	-85 dBc
17.0	-88 dBc
7.0	-94 dBc
3.0	-98 dBc

Table 6.4: Relationship between clock jitter and output noise floor for SD4CT

By extrapolating the data out to a noise floor of -116dBc (which corresponds to 111dB SNR), the phase noise density can be no greater than $31.6\mu V/\sqrt{Hz}$. Thus, the maximum phase noise power the modulator can tolerate in a 200Hz band is -67.0dBV^2.

Since the SD4CT has a different RTZ DAC implementation, it is reasonable to guess that its sensitivity to clock jitter may be different than the SD3CT. Also, its modulator loop has different characteristics, which could also slightly affect sensitivity to jitter. The data obtained from the SD4CT is detailed in Table 6.4. It shows that the SD4CT sensitivity to jitter is similar to the SD3CT.

Extrapolating the data out to -116dB again, the phase noise density must be

$40\mu V/\sqrt{Hz}$ at most. This value corresponds to 65.0dBV^2 in a 200Hz band around the carrier (again, with an amplitude of $1V_p$). This result is essentially the same as that obtained for the SD3CT.

The phase noise of a VCO is determined mainly by two characteristics: 1) circuit noise in the oscillator, and 2) the shape of its resonator structure. The shape of phase noise in a VCO usually resembles that of a resonator (having skirts on the sides of the carrier); it is almost never white (as in our tests). Most of the noise power is concentrated at frequencies close to the carrier. So in order to minimize the phase noise power within a small 100Hz band, a high-Q resonator structure would probably be needed. Such designs are difficult to realize on chip; indeed, the difficulty of designing reliable resonator structures led us to abandon them when designing the modulator loop!

So while achieving such low phase noise may be possible, it does appear to be extremely difficult. It is reasonable to conclude that such a design effort would take up a large amount of time and manpower, perhaps more than is worthwhile.

6.3 Future Work

Aside from possibly investigating designs for low phase noise oscillators, there are other aspects of this design that could benefit from more detailed research. First, redesigning the integrator opamps for this specific purpose would be beneficial in reducing power consumption (and even cutting down on silicon area). Since these particular opamps were not optimized for driving resistive loads, using them in a CT modulator is somewhat clumsy.

Furthermore, the possibility of using Gm-C filters (transconductance amplifiers with a capacitive load) to realize the integrators is an attractive one. Considering that operational transconductance amplifiers (OTA's) use significantly less power than traditional opamps, using them could result in even more power savings.

Another design topic to be explored is the continuous-time front-end topology suggested in Chapter 4. While this design will probably not reduce power consump-

tion, it could definitely reduce area and improve SNR. Also, since only one feedback path would be in continuous-time, the effects of jitter may be significantly diminished. Since we earlier considered this option to be attractive for a modulator design, it would be beneficial to research this design further.

Chapter 7

Conclusion

In this document, we have researched the pros and cons of implementing a continuous-time bandpass Σ-Δ modulator for the specific purpose of digitizing carrier band measurements. It has been found that continuous-time modulators usually take up less silicon area and result in lower power consumption than their discrete-time counterparts. Furthermore, the circuit noise of continuous-time integrators is found to be significantly less than those implemented with switched-capacitors; thus, the SNR (and the resolution of the data conversion) can be larger.

On the other hand, designing on-chip continuous-time modulators to be robust against the effects of temperature and process variations proves to be more difficult. We showed that the component's absolute tolerances made bandpass loop topologies unreliable and that lowpass topologies made it more difficult to effectively shape quantization noise. However, we also demonstrated that theoretically, a CT modulator's quantization noise could be adequately shaped to achieve comparable SNR to a third-order DT modulator. We presented a continuous-time, fourth-order, single-loop design in which the theoretical SNR is similar to the DT modulator designed for the same purpose. Moreover, this design achieves this performance with roughly 40% reduction in both power consumption and silicon area.

While the fourth order design appears to be an improvement, we found that when fabricated on silicon, the quantization noise floor was significantly higher than expected for reasons unknown at this time. More importantly, we found that clock

jitter degrades the modulator's performance considerably. In studying the effects of jitter on the sampling clock, we discovered that VCO phase noise feeds through to the output and is modulated on the carrier frequency of the input. Our tests show that phase noise cannot be more than -65dBV^2 within $\pm$100Hz of the oscillator frequency if the modulator is to achieve 111dB SNR. Such low phase noise can be reached with either an off-chip crystal oscillator or an on-chip VCO with extremely low phase noise; neither of these options appears to be viable at the time of this writing.

It appears that using a continuous-time modulator for this specific application is not viable, solely because of the effects of clock jitter. However, it is not clear how clock jitter would affect a *single* continuous-time feedback path; therefore, the possibility of using a hybrid topology remains attractive. Implementing the first stage in continuous time and the subsequent stages in discrete time has benefits which may still prove useful for this application.

Appendix A

Circuit Schematics of SD3CT and SD4CT

- Schematic of SD3CT noise shaper
- Schematic of SD3CT RTZ DAC (one-shot)
- Schematic of Break-before-Make Generator
- Schematic of SD4CT noise shaper

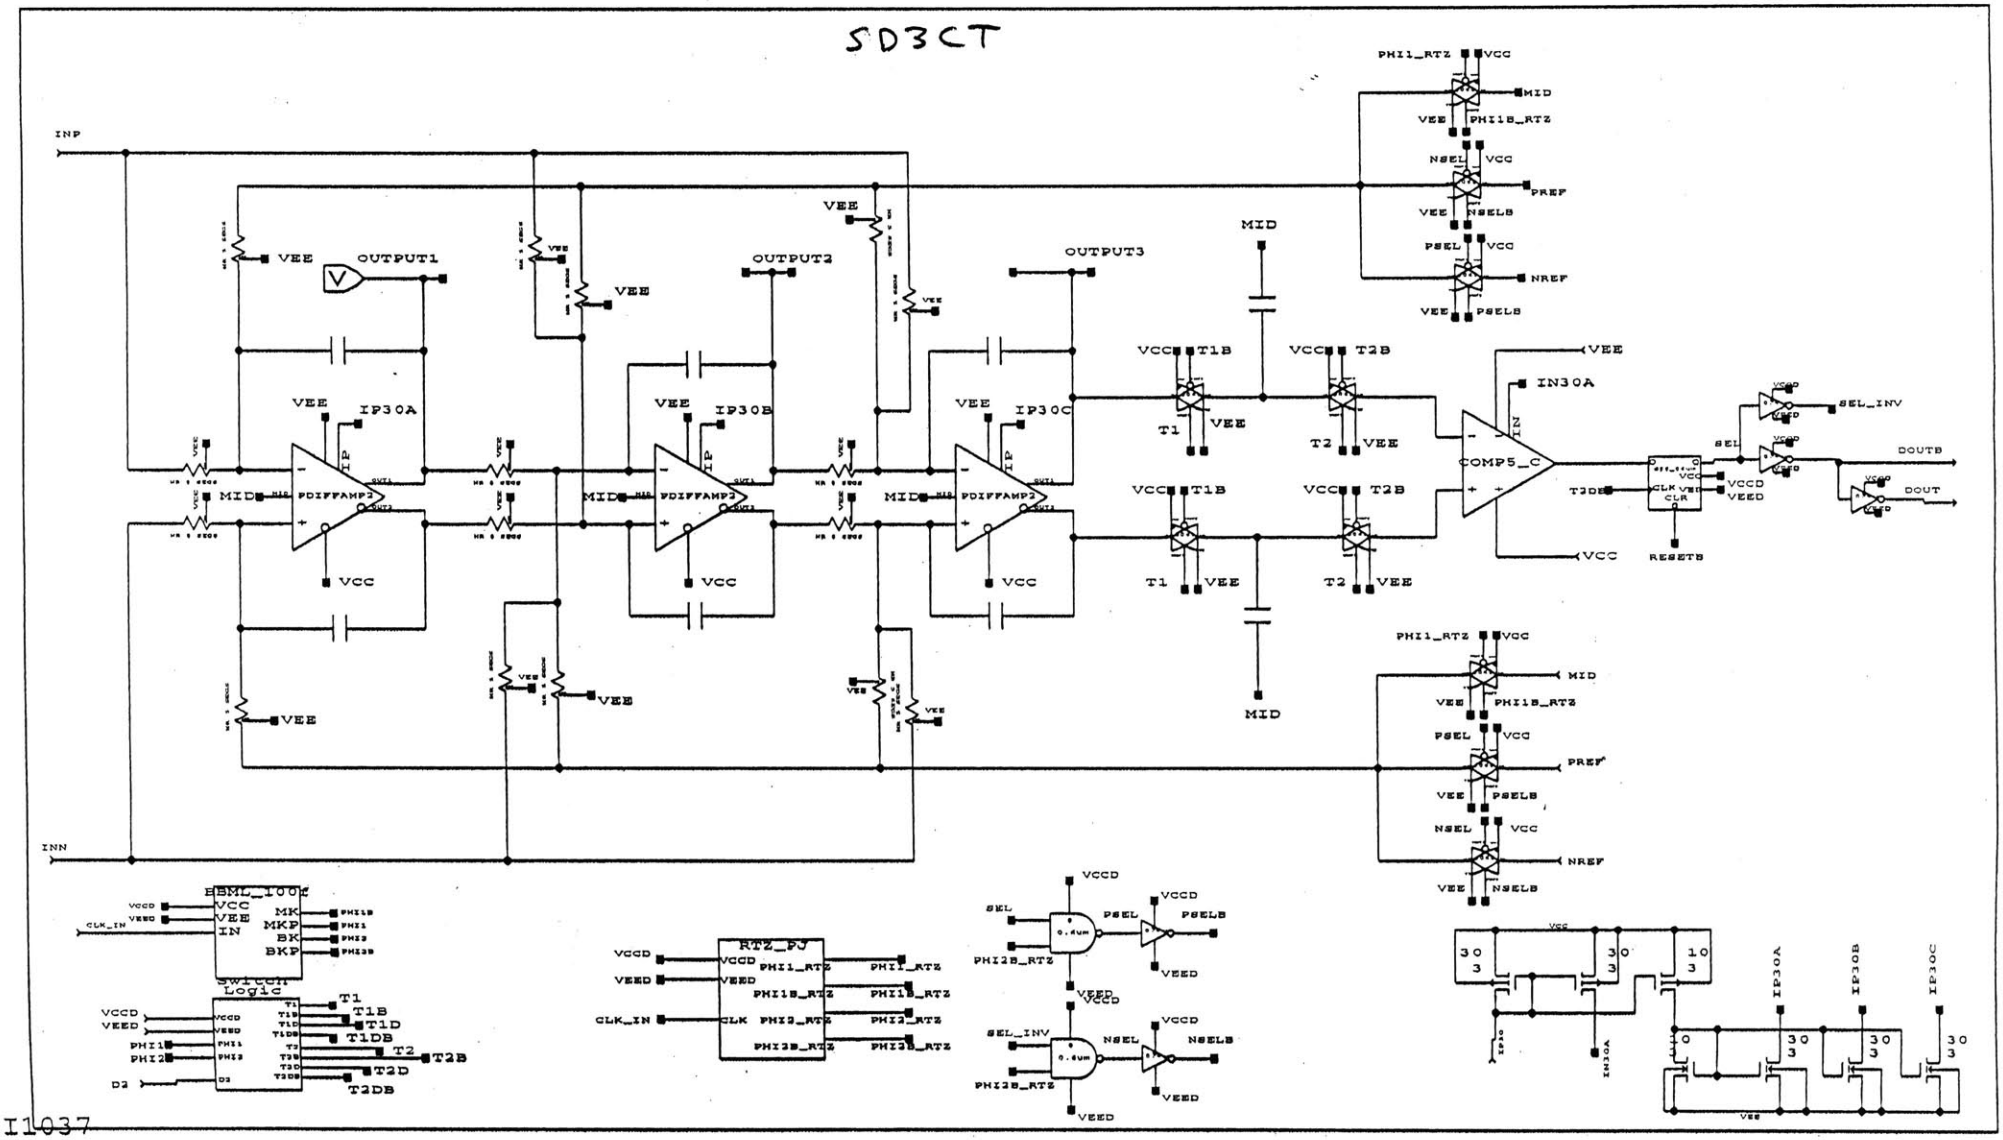
5D3CT
INP
INN
VEE
VCC
OUTPUT1
OUTPUT2
OUTPUT3
IP30A
IP30B
IP30C
MID
PDIFFAMP3
COMP5_C
IN30A
T1
T1B
T2
T2B
PHI1_RTZ
PHI1B_RTZ
PHI2_RTZ
PHI2B_RTZ
PSEL
NSEL
PSELB
NSELB
PREF
NREF
SEL
SEL_INV
DOUT
DOUTB
RESETB
T3DB
CLK
CLR
VCCD
VEED
CLK_IN
RTZ_PJ
BBML_100E
Switch Logic
T1D
T1DB
T2D
T2DB
PHI1
PHI2
D2

$1I1037

RTZ_PJ

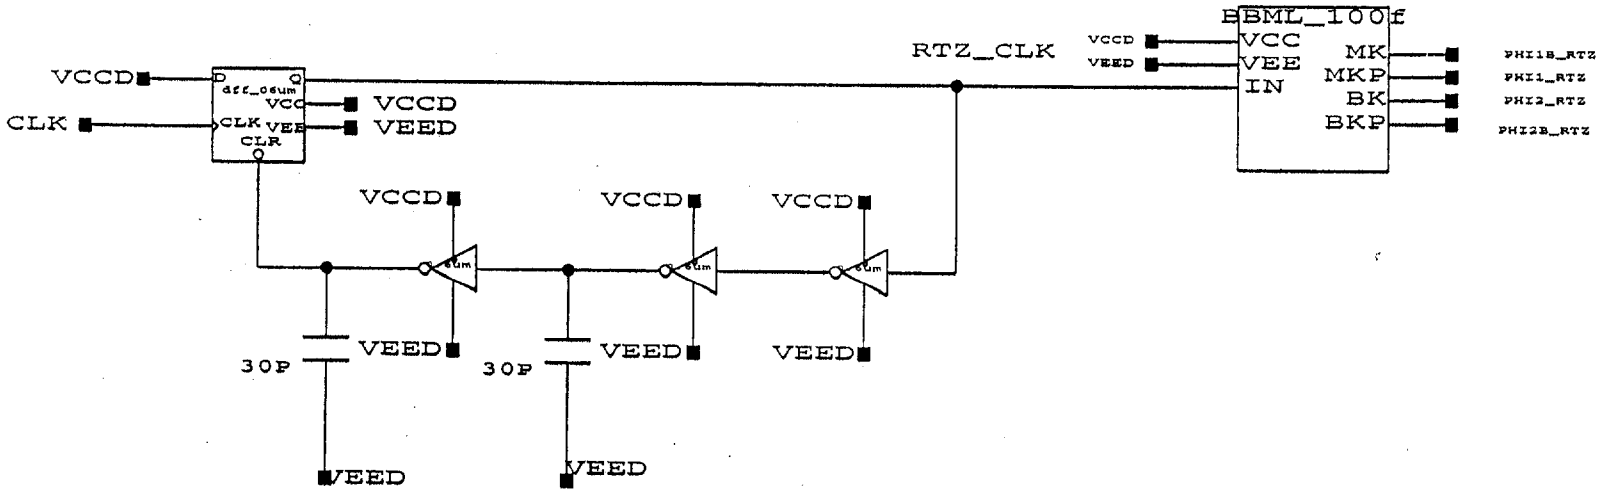
VCCD
CLK
dff_06um
VCC
CLK
VEE
CLR
VCCD
VEED
RTZ_CLK
VCCD
VCCD
VCCD
06um
06um
06um
30P
30P
VEED
VEED
VEED
VEED
VEED
BBML_100f
VCCD
VEED
VCC
VEE
IN
MK
MKP
BK
BKP
PHI1B_RTZ
PHI1_RTZ
PHI2_RTZ
PHI2B_RTZ

BREAK-BEFORE MAKE GENERATOR

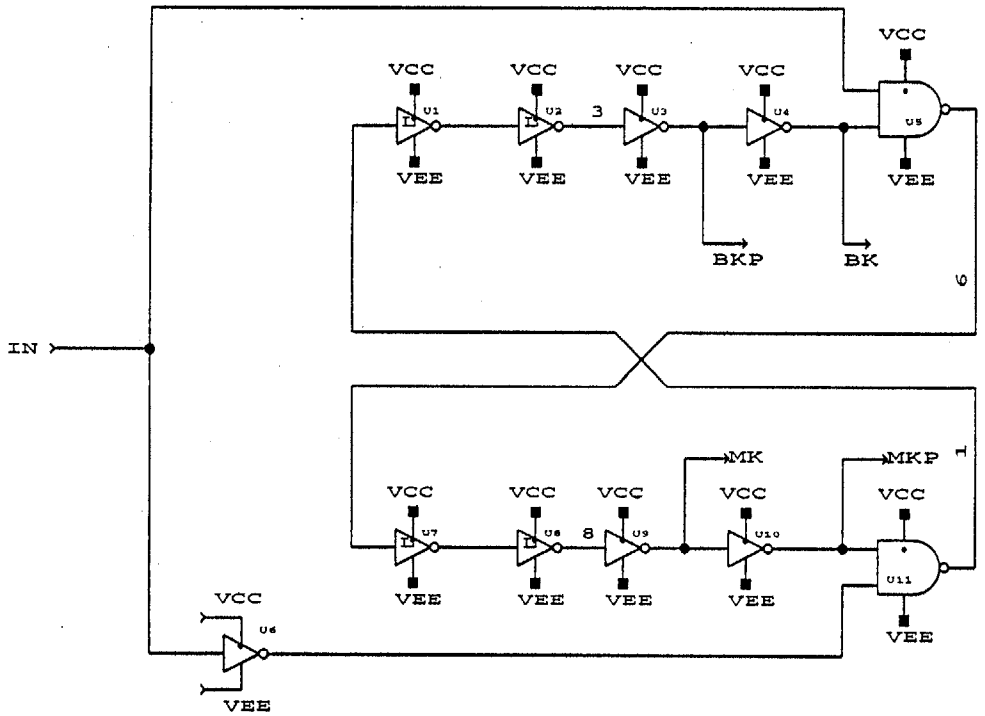

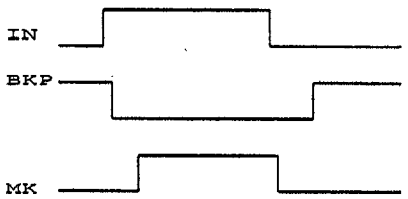

BBML_100f

Draper Lab	Dwg. Number: 383760
Ed Balboni	REV:
RSS 0.6um translated to	Date: 2-22-2000_10:58
	Name: c5_bbml_1 of Sheet of

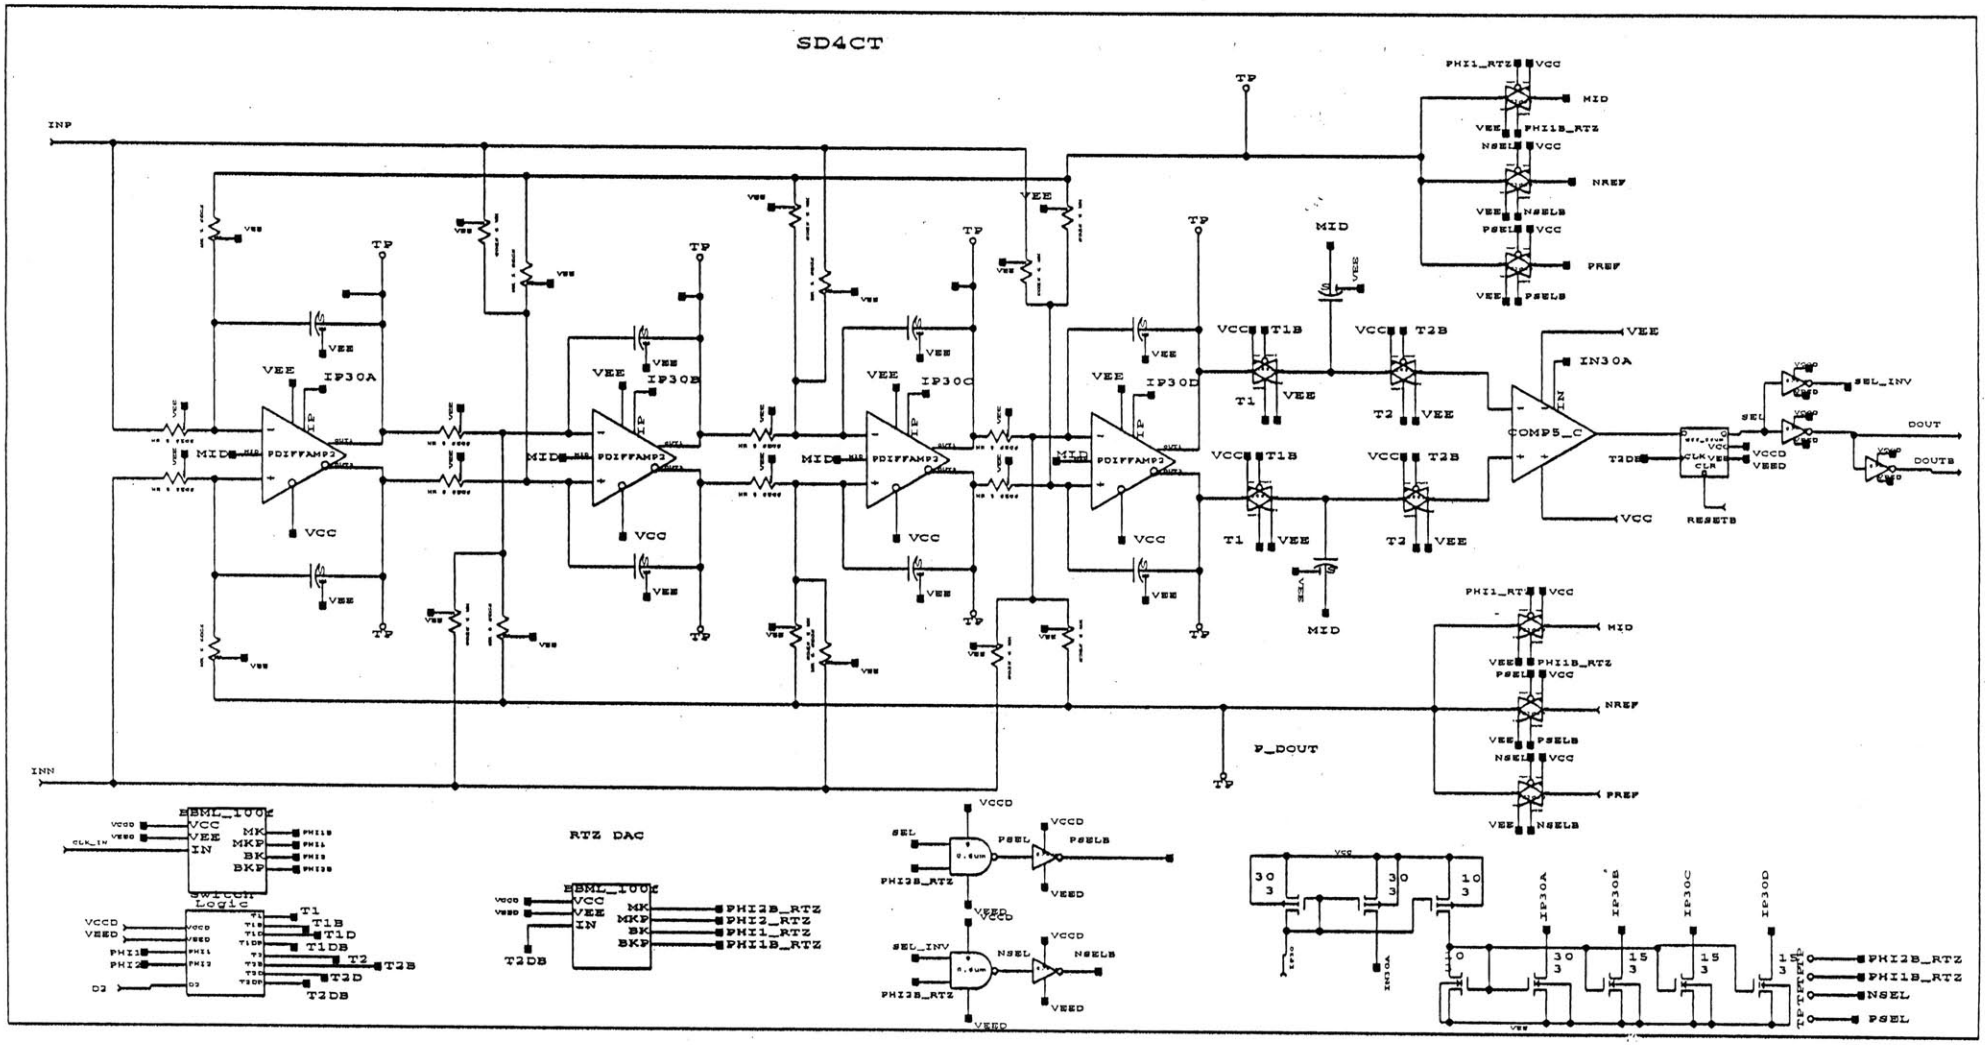
SD4CT
INP
INN
TP
VEE
VCC
MID
IP30A
IP30B
IP30C
IP30D
PDIFFAMP2
T1B
T1
T2B
T2
COMP5_C
IN30A
SEL_INV
DOUT
DOUTB
RESETB
T2DB
PHI1_RTZ
PHI1B_RTZ
NSELB
NREF
PSELB
PREF
P_DOUT
EBML_100C
Switch Logic
RTZ DAC
PHI2B_RTZ
PHI2_RTZ
PHI1_RTZ
PHI1B_RTZ
NSEL
PSEL

Appendix B

SD3CT Simulation and Lab Test Data

- AC Simulation Schematic of Noise Transfer Function (NTF)
- Result of AC Simulation of NTF
- Simulated spectrum of SD3CT at 27 degrees C (Using StarSim)
- Simulated spectrum of SD3CT at -40 degrees C
- Simulated spectrum of SD3CT at 85 degrees C
- Measured spectrum of SD3CT at 27 degrees C
- Measured spectrum of SD3CT at -40 degrees C
- Measured spectrum of SD3CT at 85 degrees C

BODE PLOT | BIAS POINT | AMI C5 NOMINAL | TEMPERATURE | OPTIONS

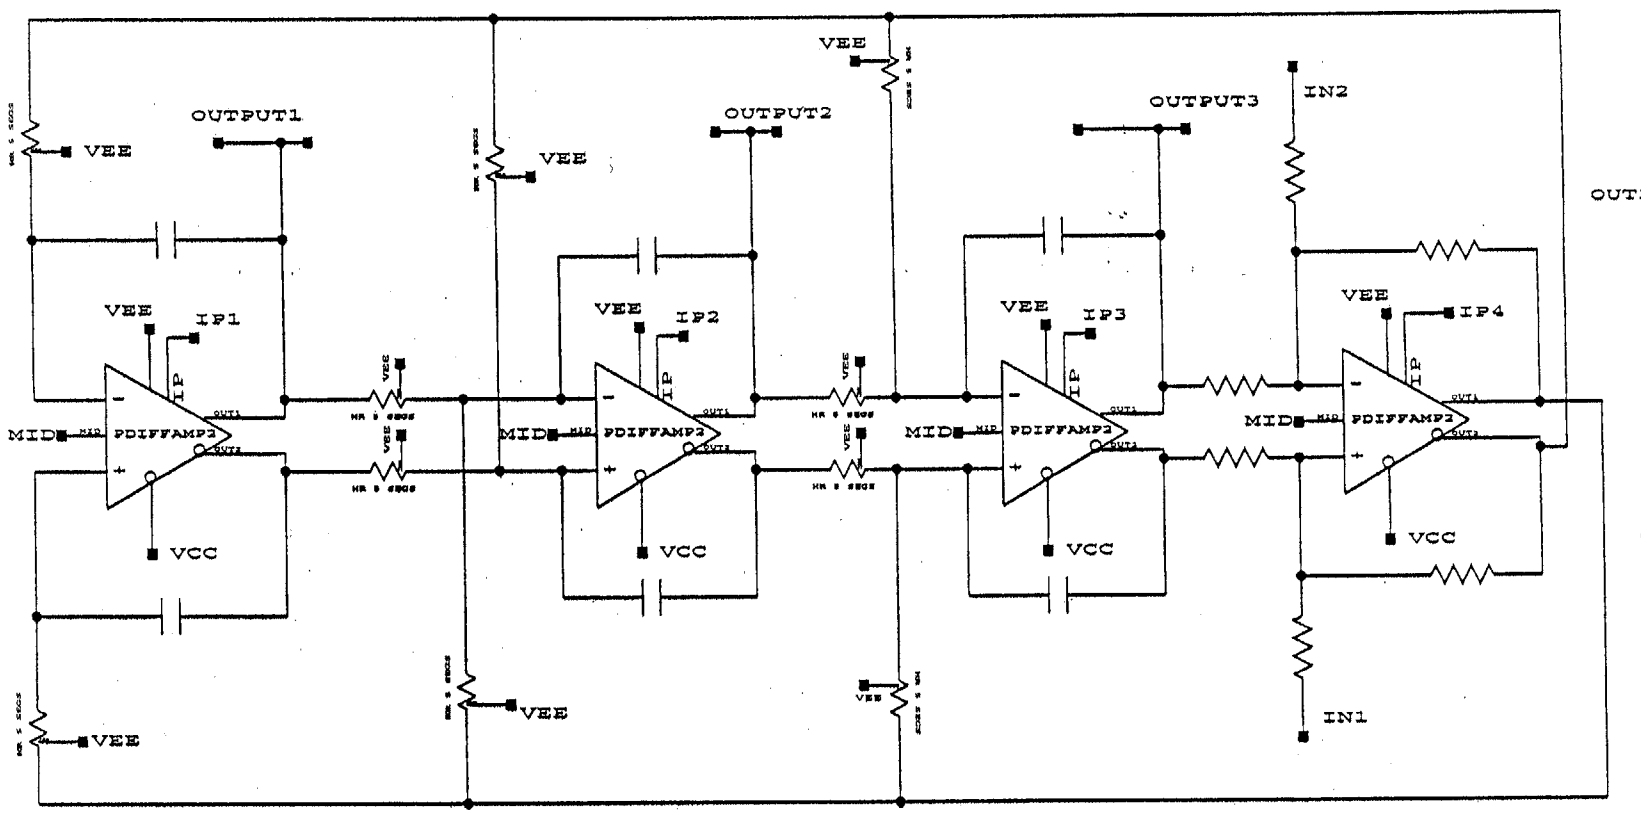

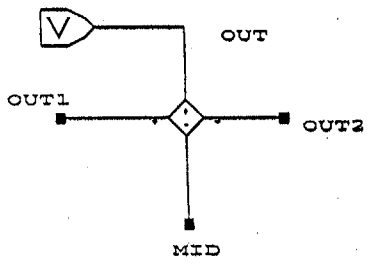

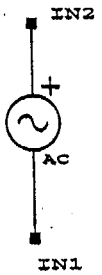

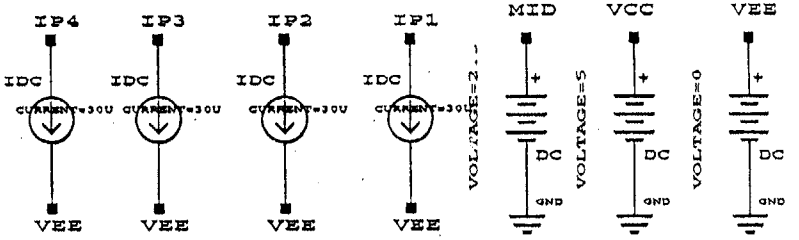

Simulation schematic of NTF

vdb(out)	-64.4728
FREQUENCY	19.9526k

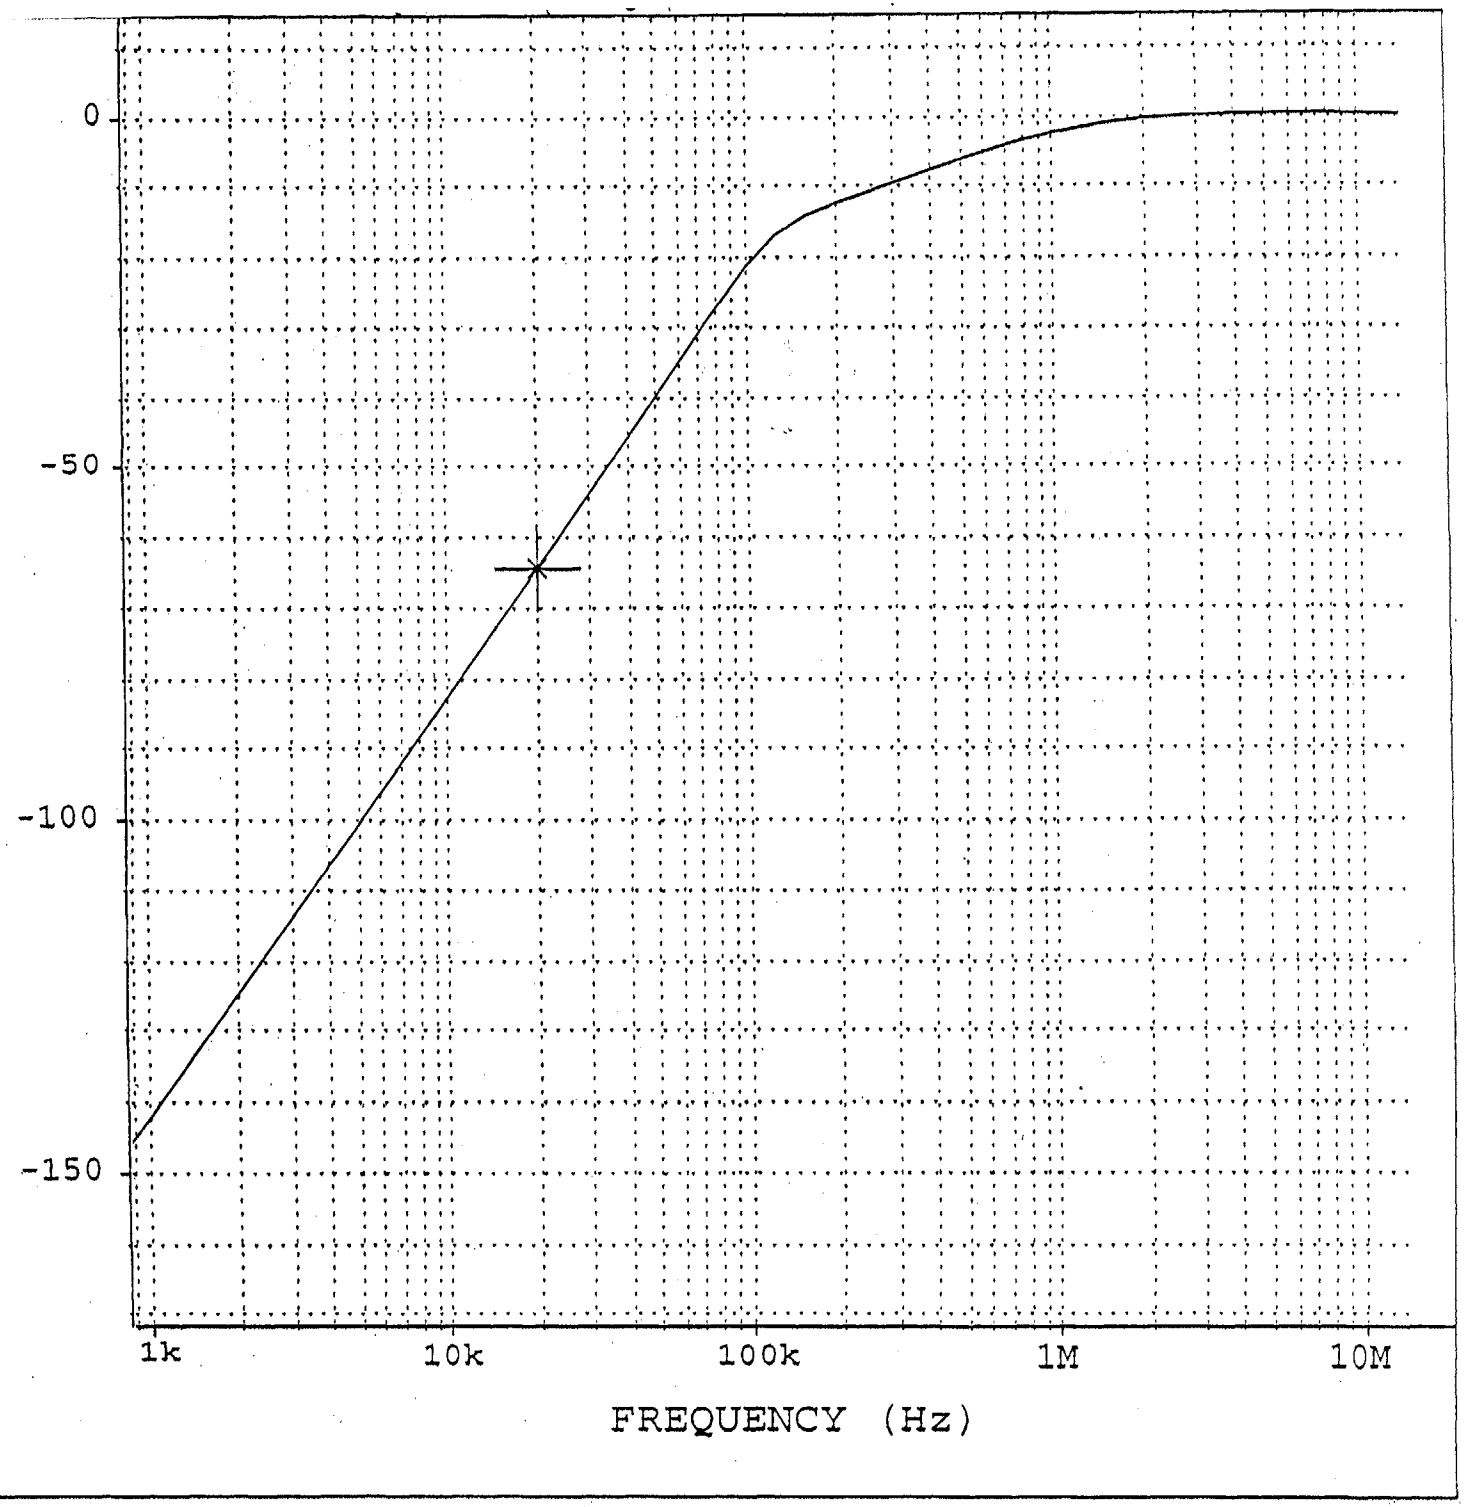

Simulated Output

27°C

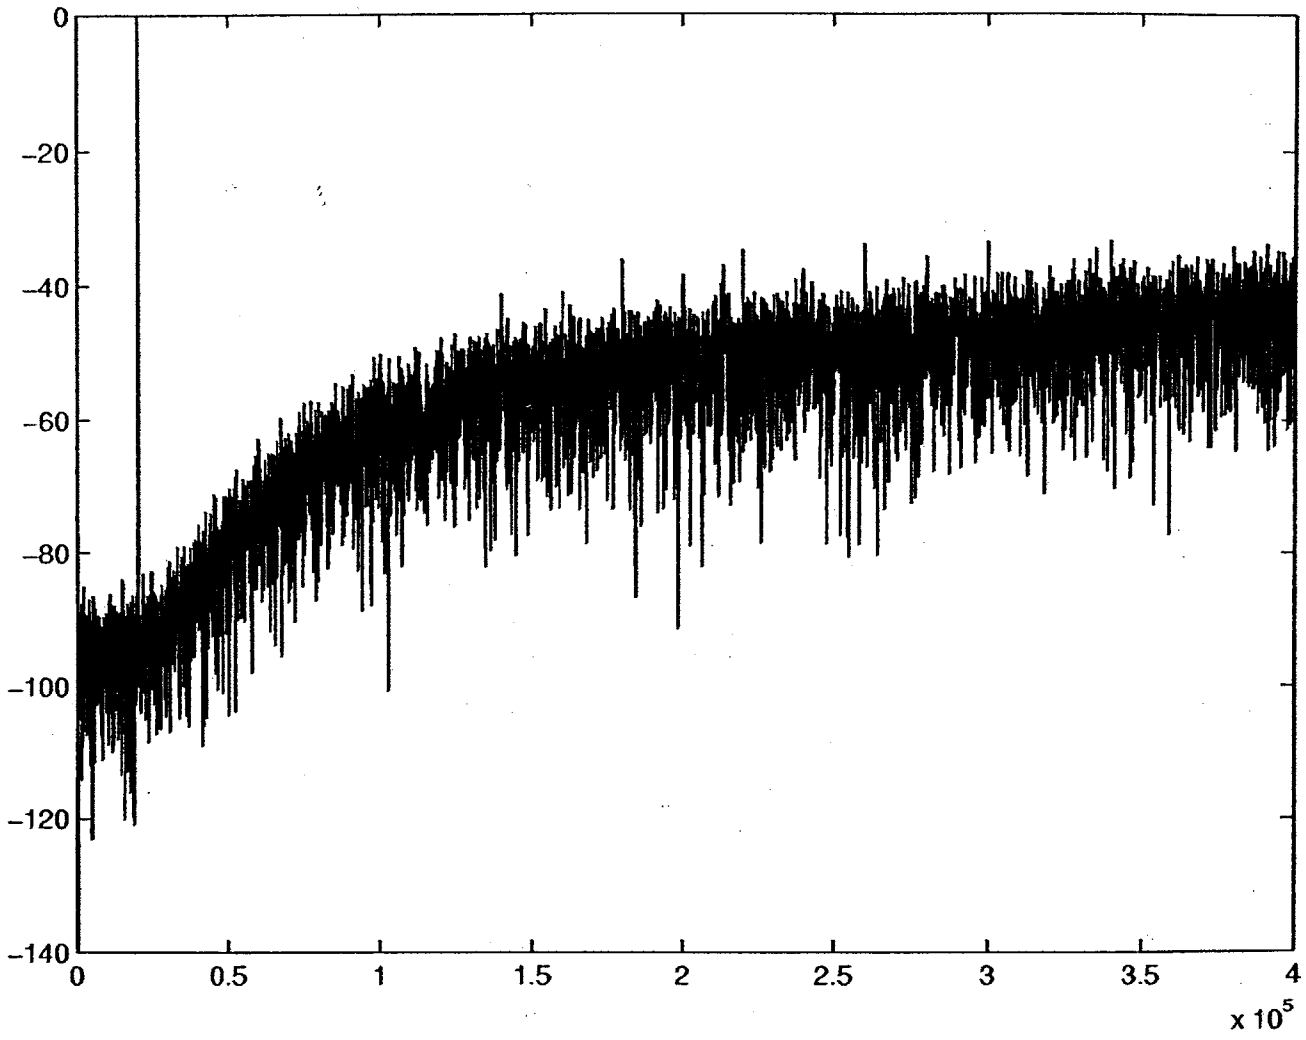

Simulated Output

-40° C

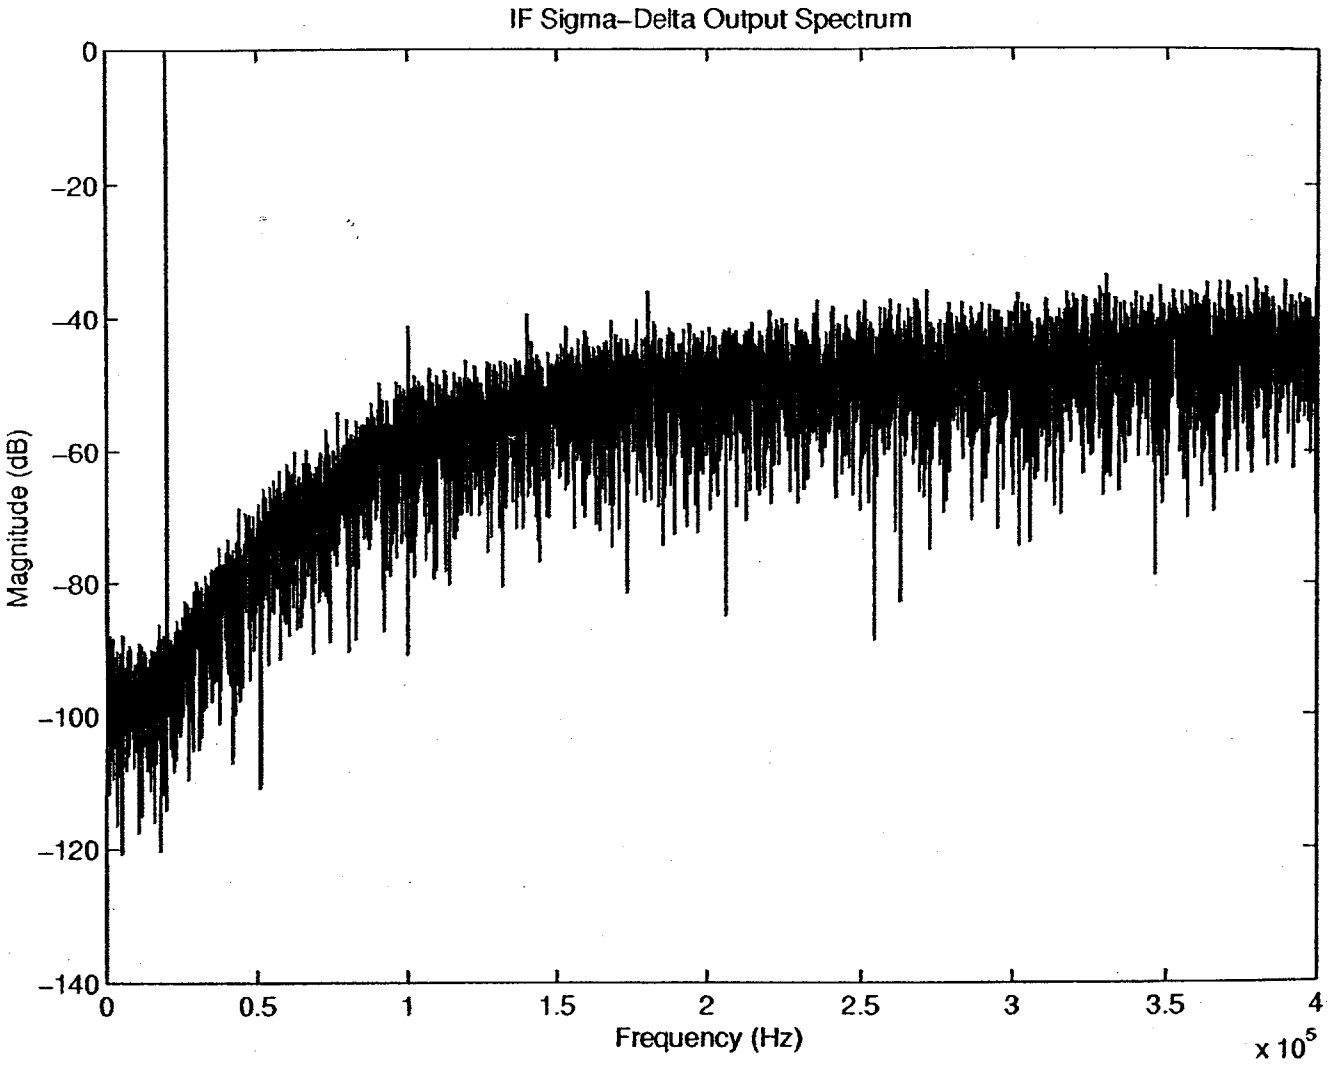

Simulated Output

85°C

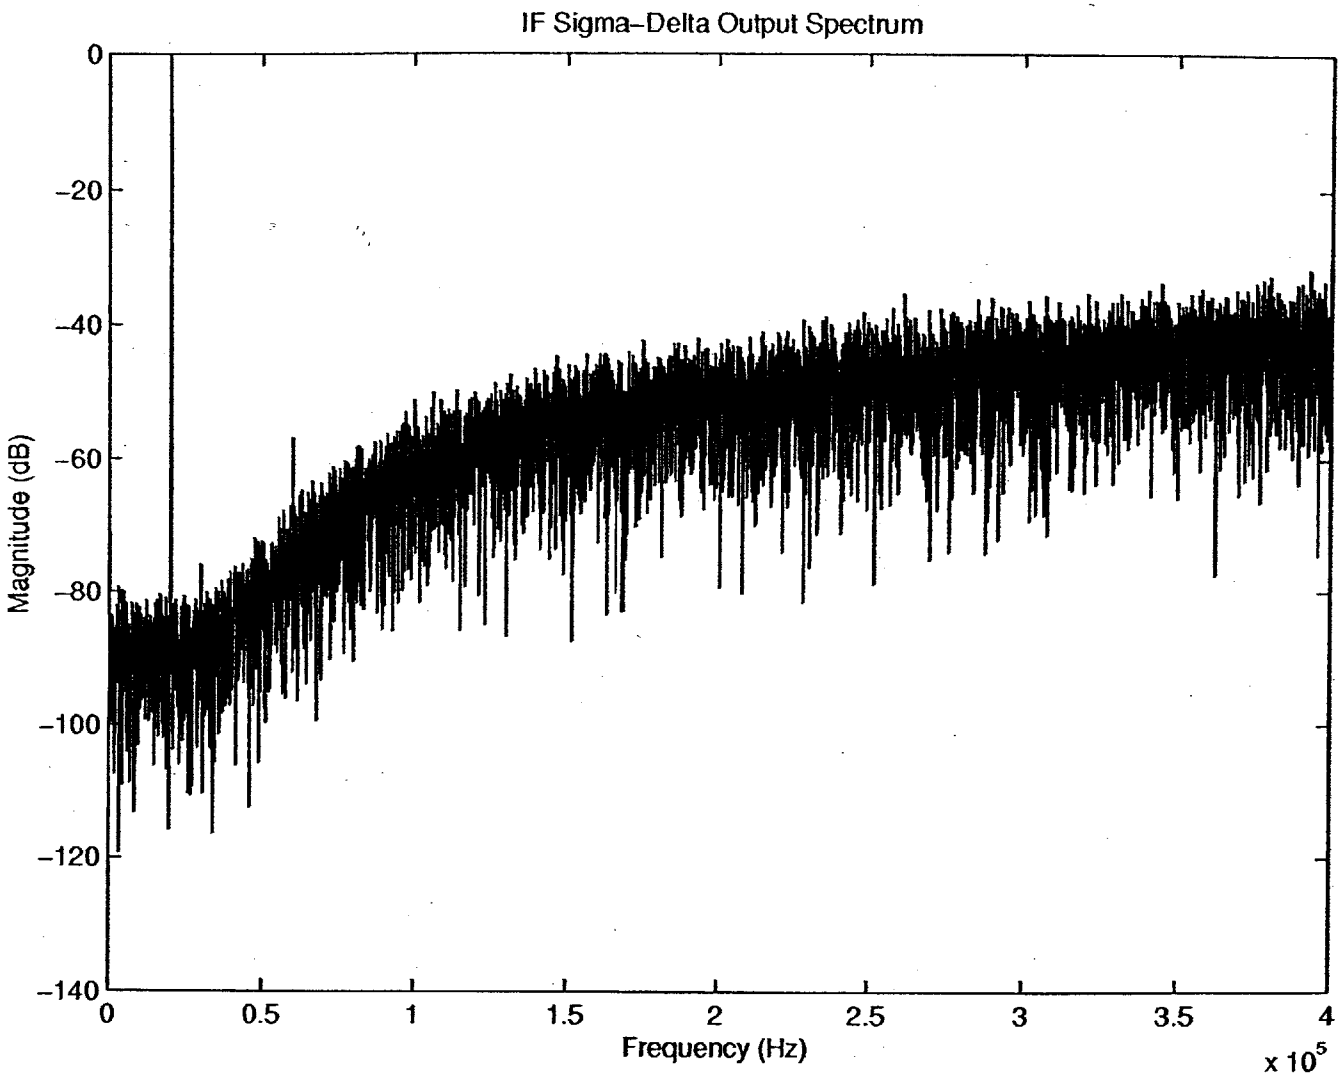

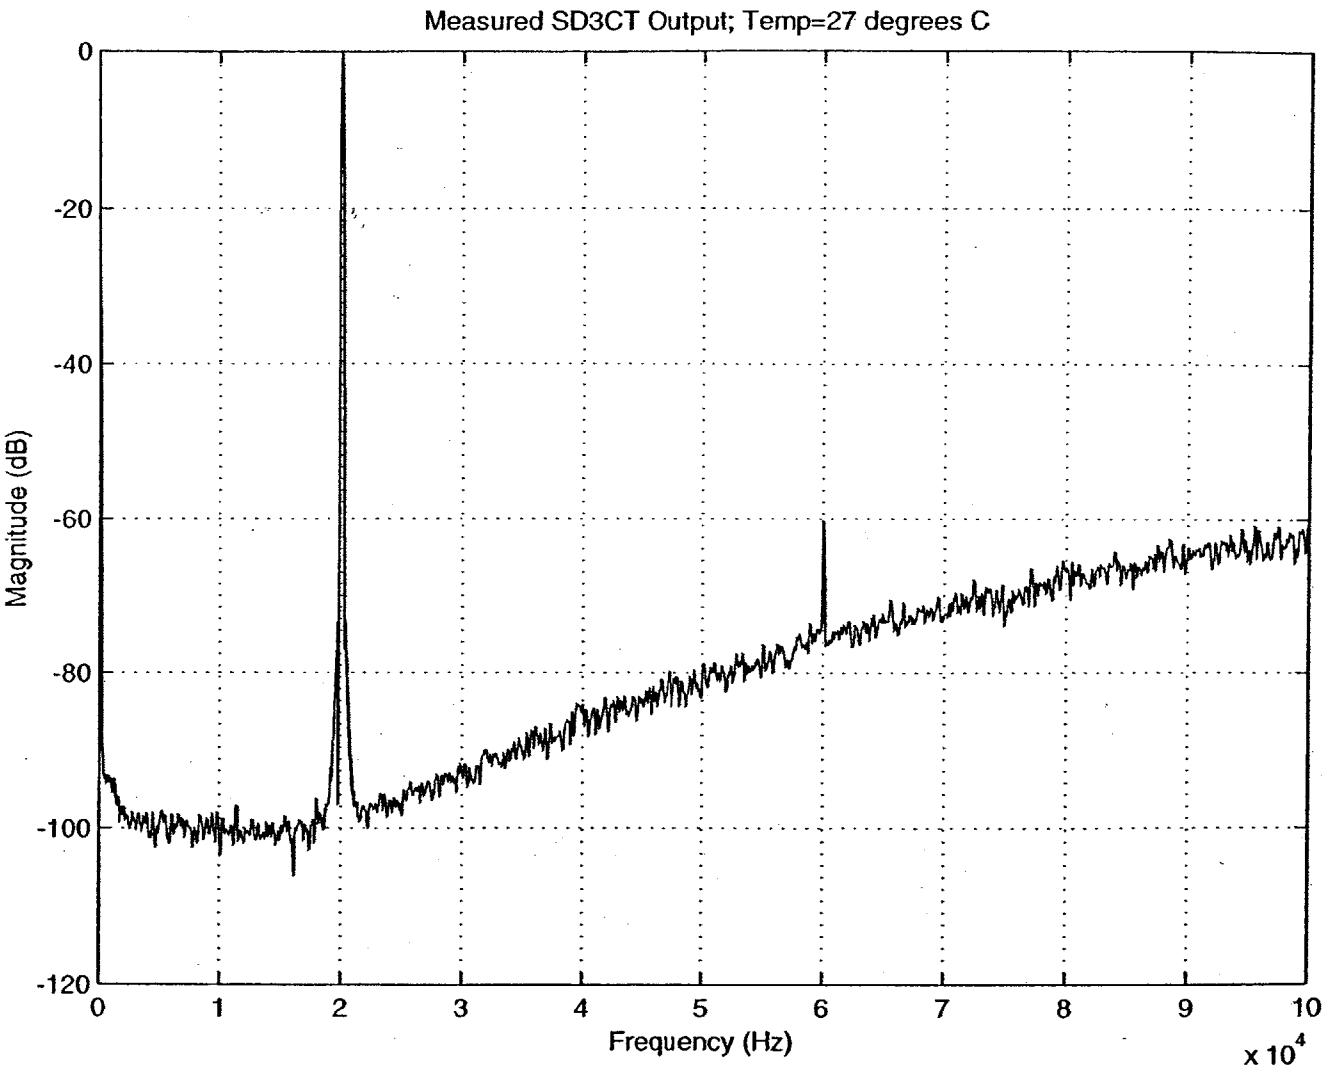
Measured SD3CT Output; Temp=27 degrees C
Magnitude (dB)
0
-20
-40
-60
-80
-100
-120
0
1
2
3
4
5
6
7
8
9
10
Frequency (Hz)
x 10^4

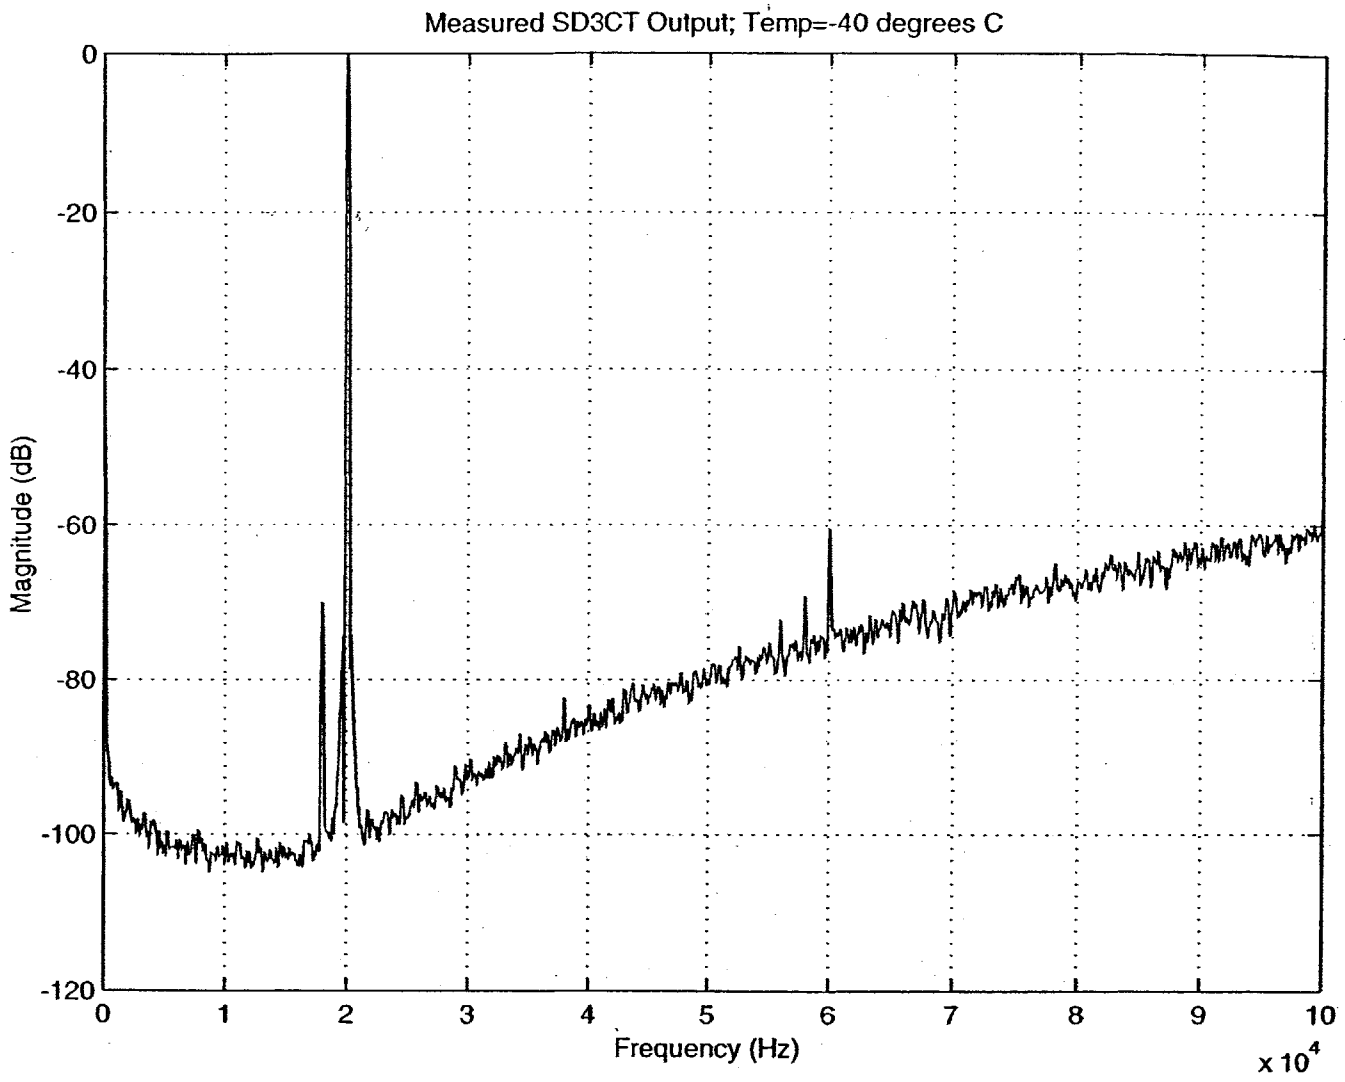
Measured SD3CT Output; Temp=-40 degrees C
0
-20
-40
-60
-80
-100
-120
Magnitude (dB)
0
1
2
3
4
5
6
7
8
9
10
Frequency (Hz)
x 10^4

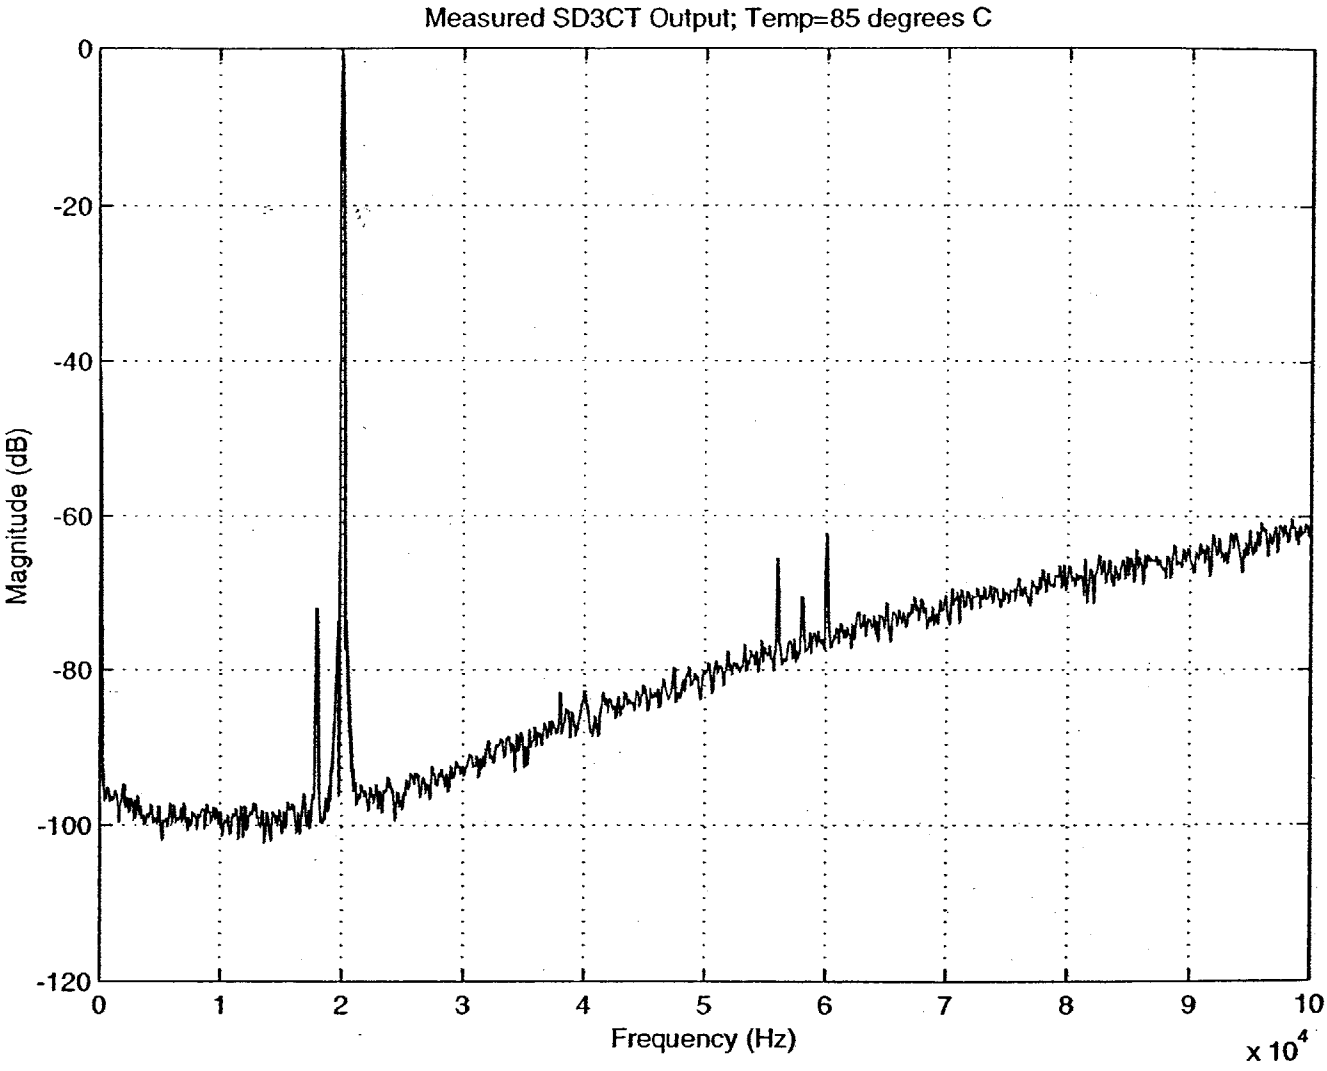
Measured SD3CT Output; Temp=85 degrees C
Magnitude (dB)
0
-20
-40
-60
-80
-100
-120
0
1
2
3
4
5
6
7
8
9
10
Frequency (Hz)
x 10^4

Appendix C

SD4CT Simulation and Lab Test Data

- AC Simulation Schematic of Noise Transfer Function (NTF)
- Result of AC Simulation of NTF
- Simulated spectrum of SD4CT at 27 degrees C (Using StarSim)
- Measured spectrum of SD4CT at 27 degrees C
- Measured spectrum of SD4CT at -40 degrees C
- Measured spectrum of SD4CT at 85 degrees C
- Measured voltage output of first stage integrator
- Measured voltage output of second stage integrator
- Measured voltage output of third stage integrator
- Measured voltage output of fourth stage integrator

OSCILLOSCOPE

BIAS POINT

AMI C5
NOMINAL

TEMPERATURE

OPTIONS

INCLUDE

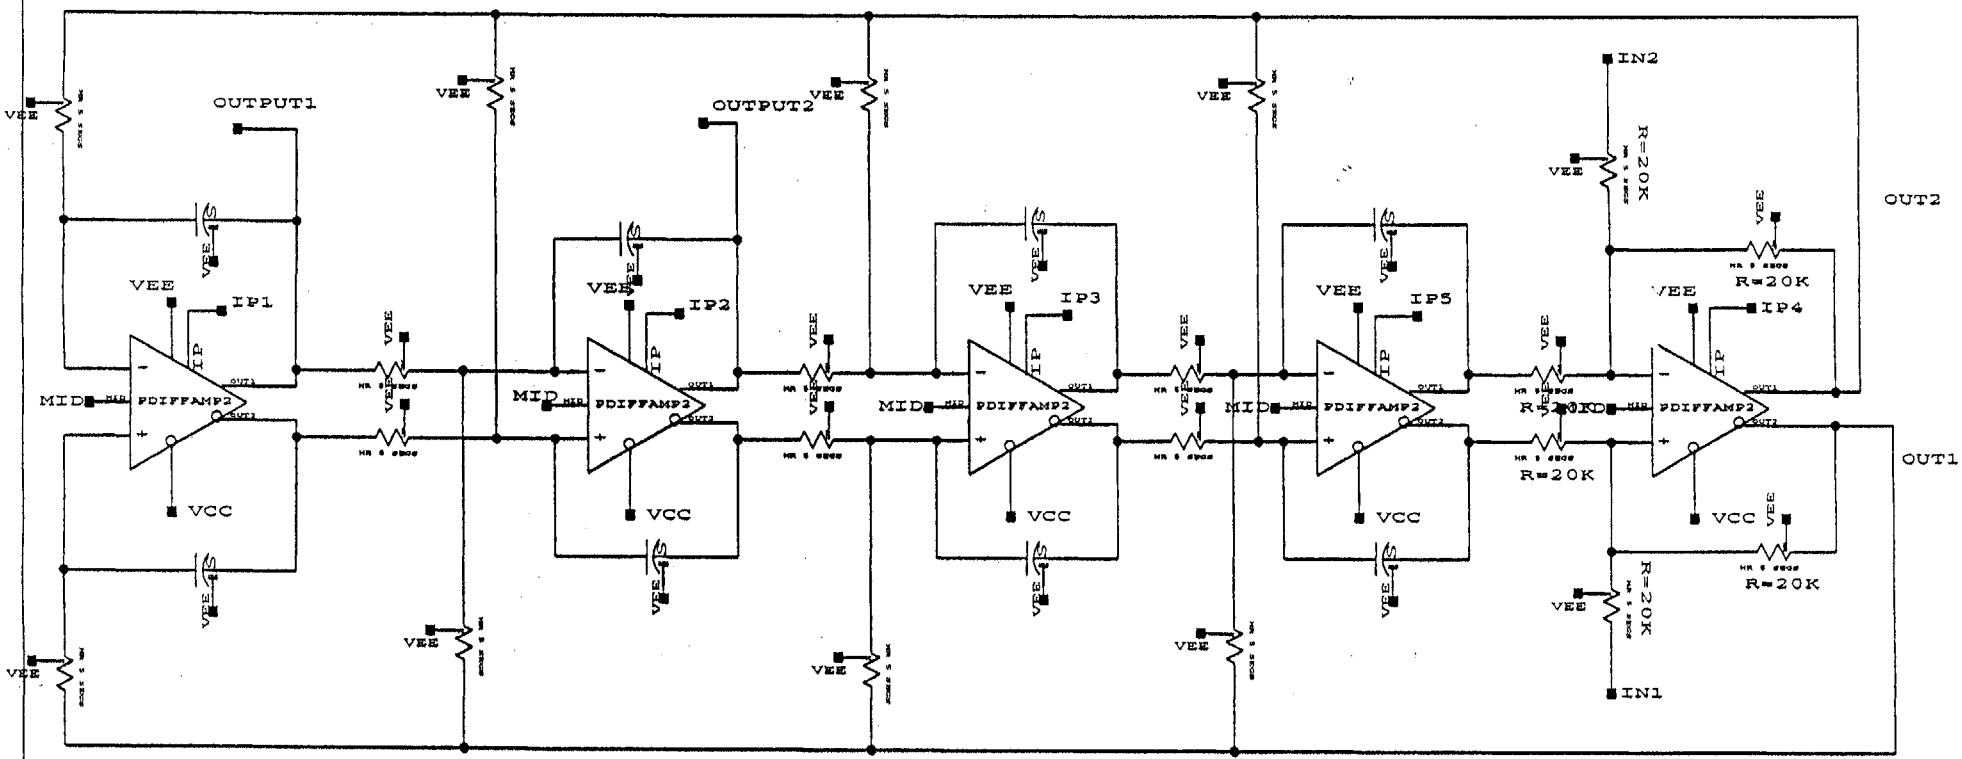

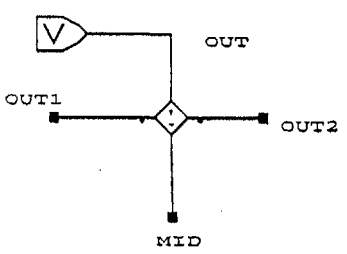

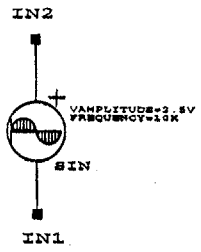

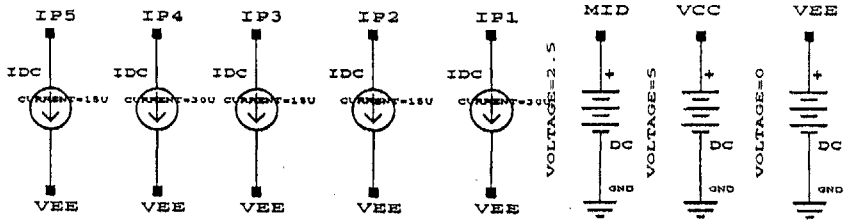

Simulation Schematic of NIA

vdb(out)	-82.8396
FREQUENCY	20.1751k

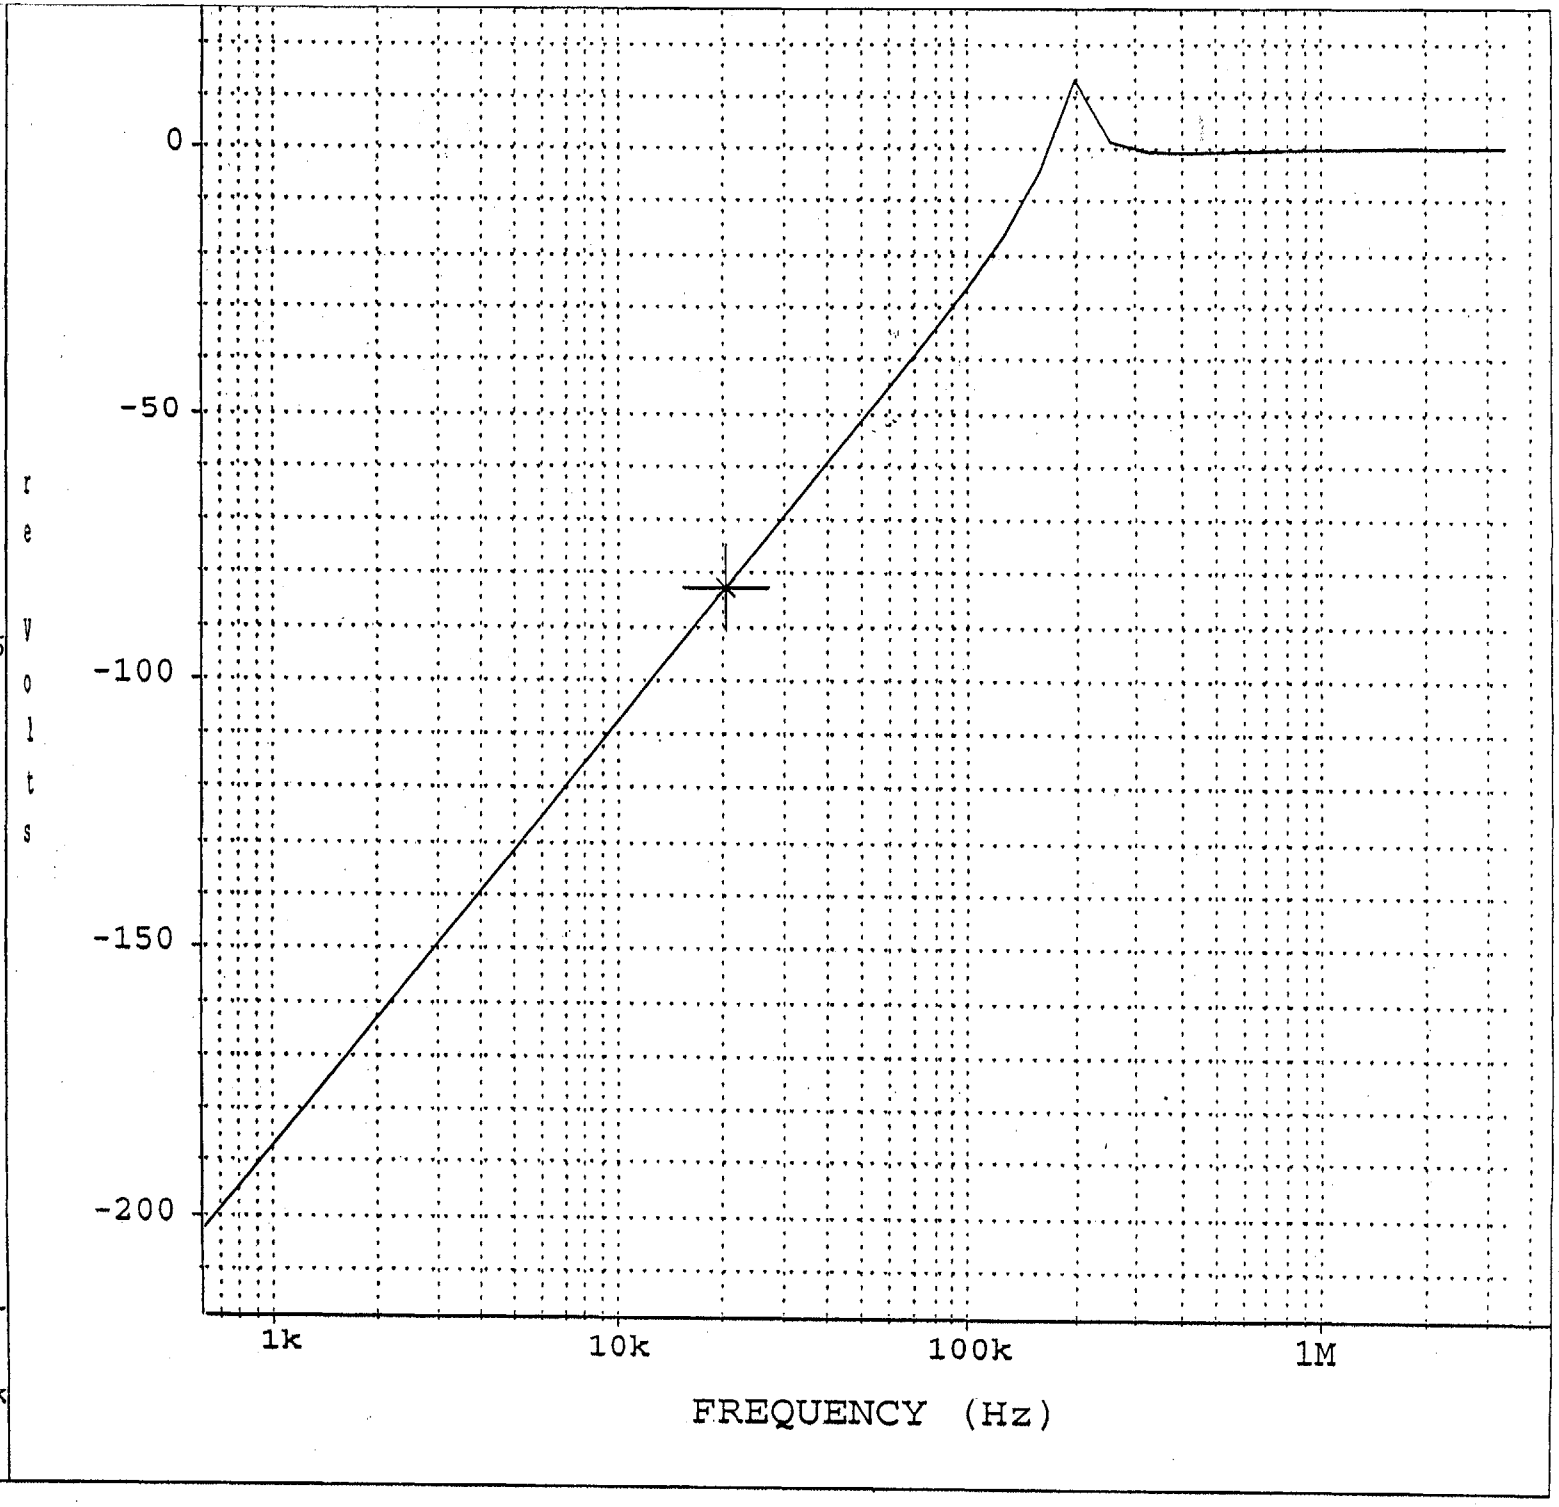

Simulated Output

27°C

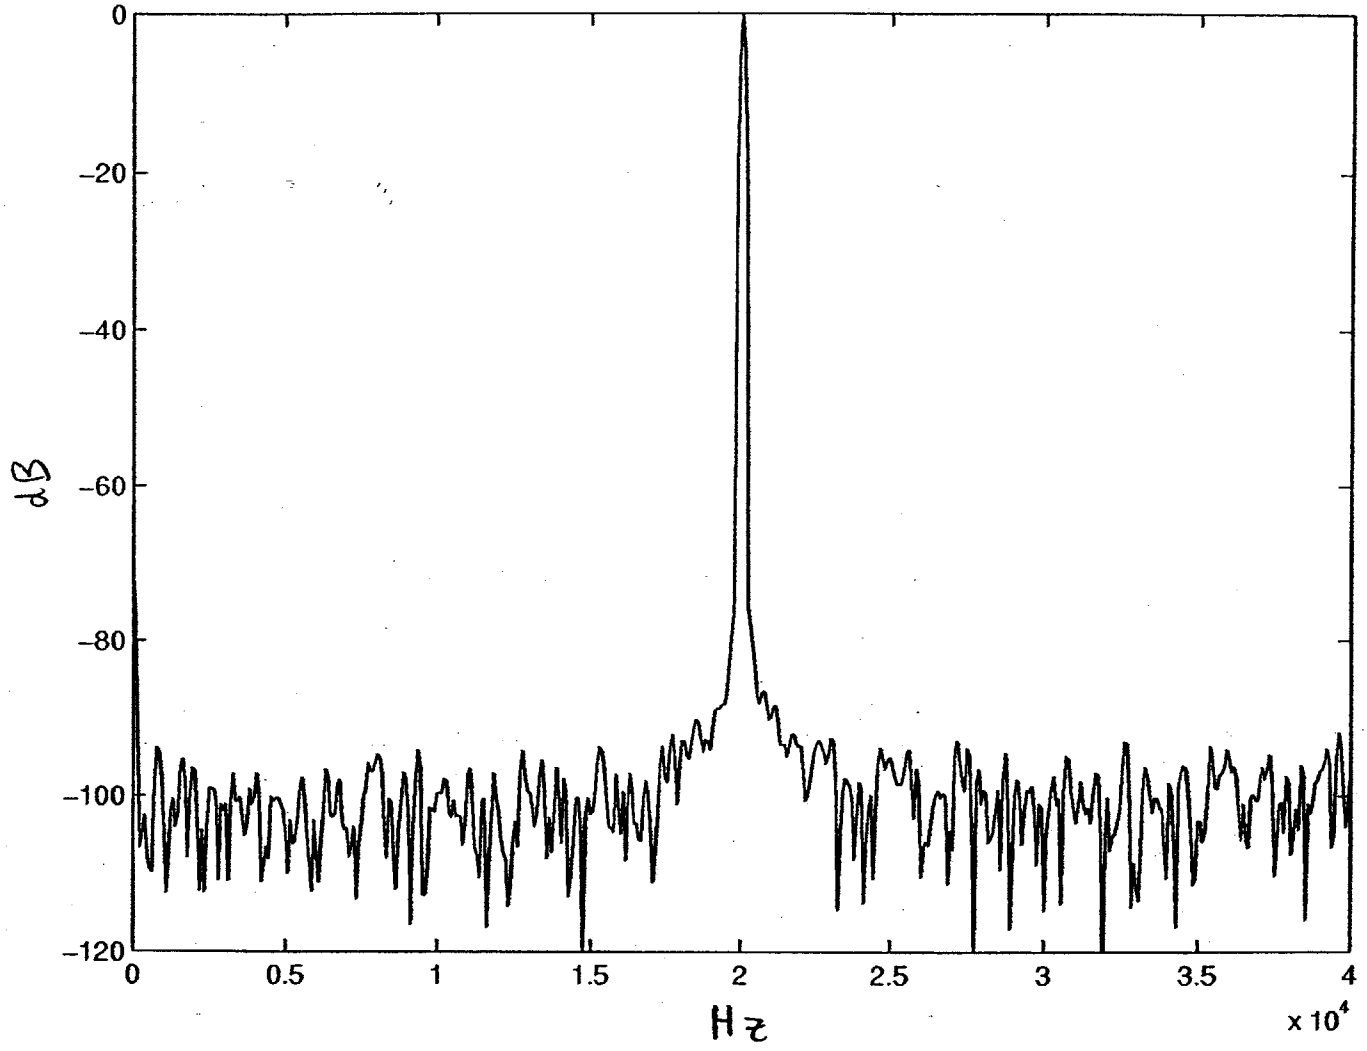

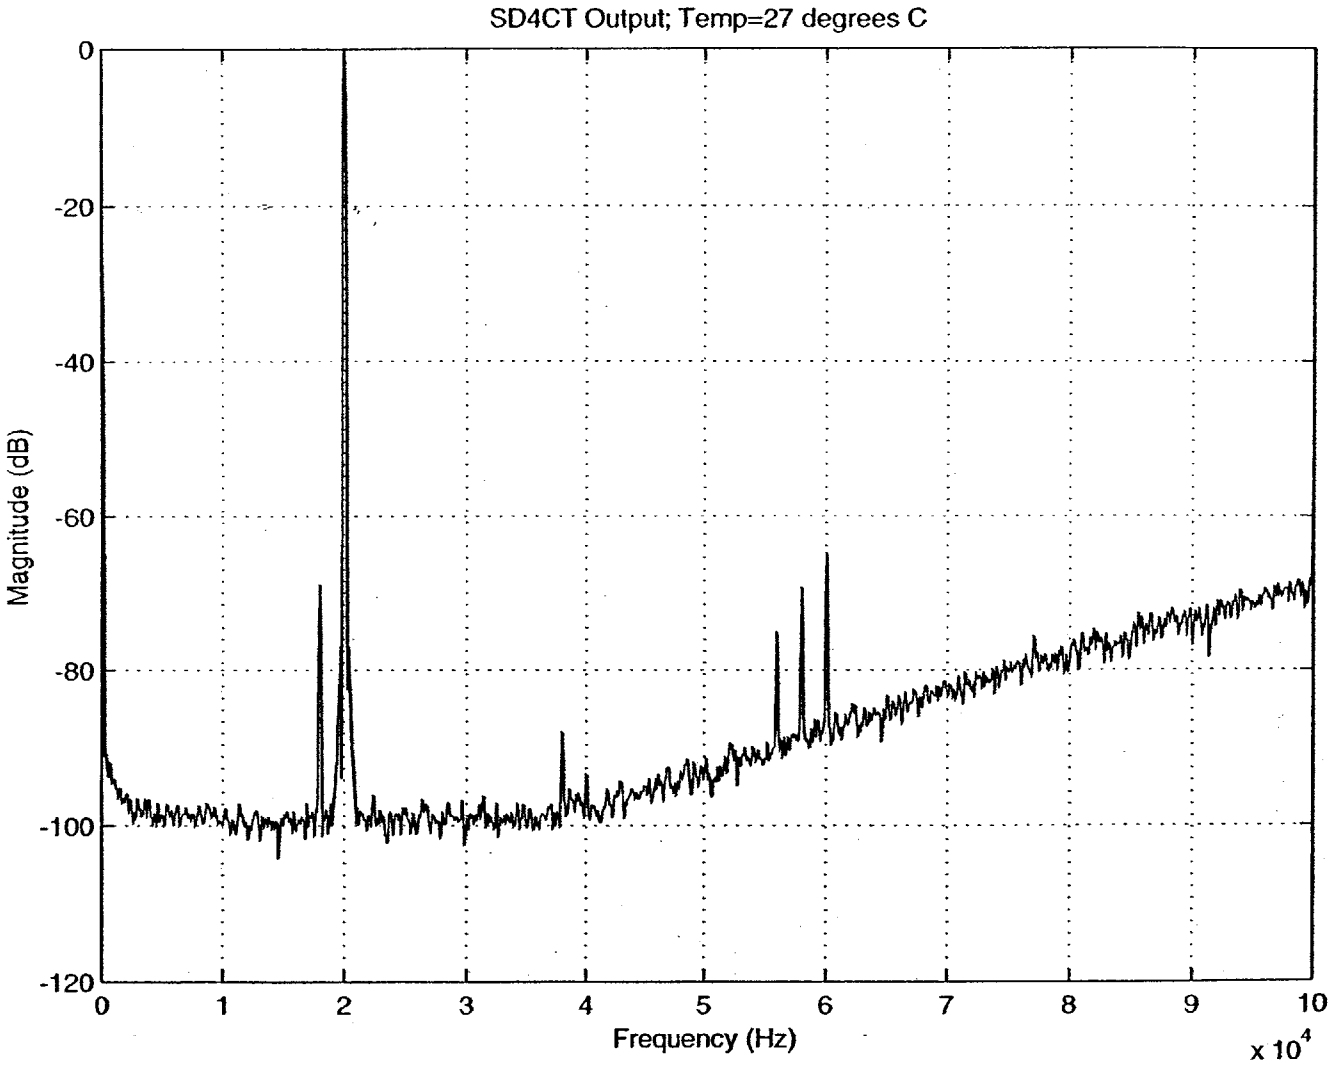
SD4CT Output; Temp=27 degrees C
0
-20
-40
-60
-80
-100
-120
Magnitude (dB)
0
1
2
3
4
5
6
7
8
9
10
Frequency (Hz)
x 10^4

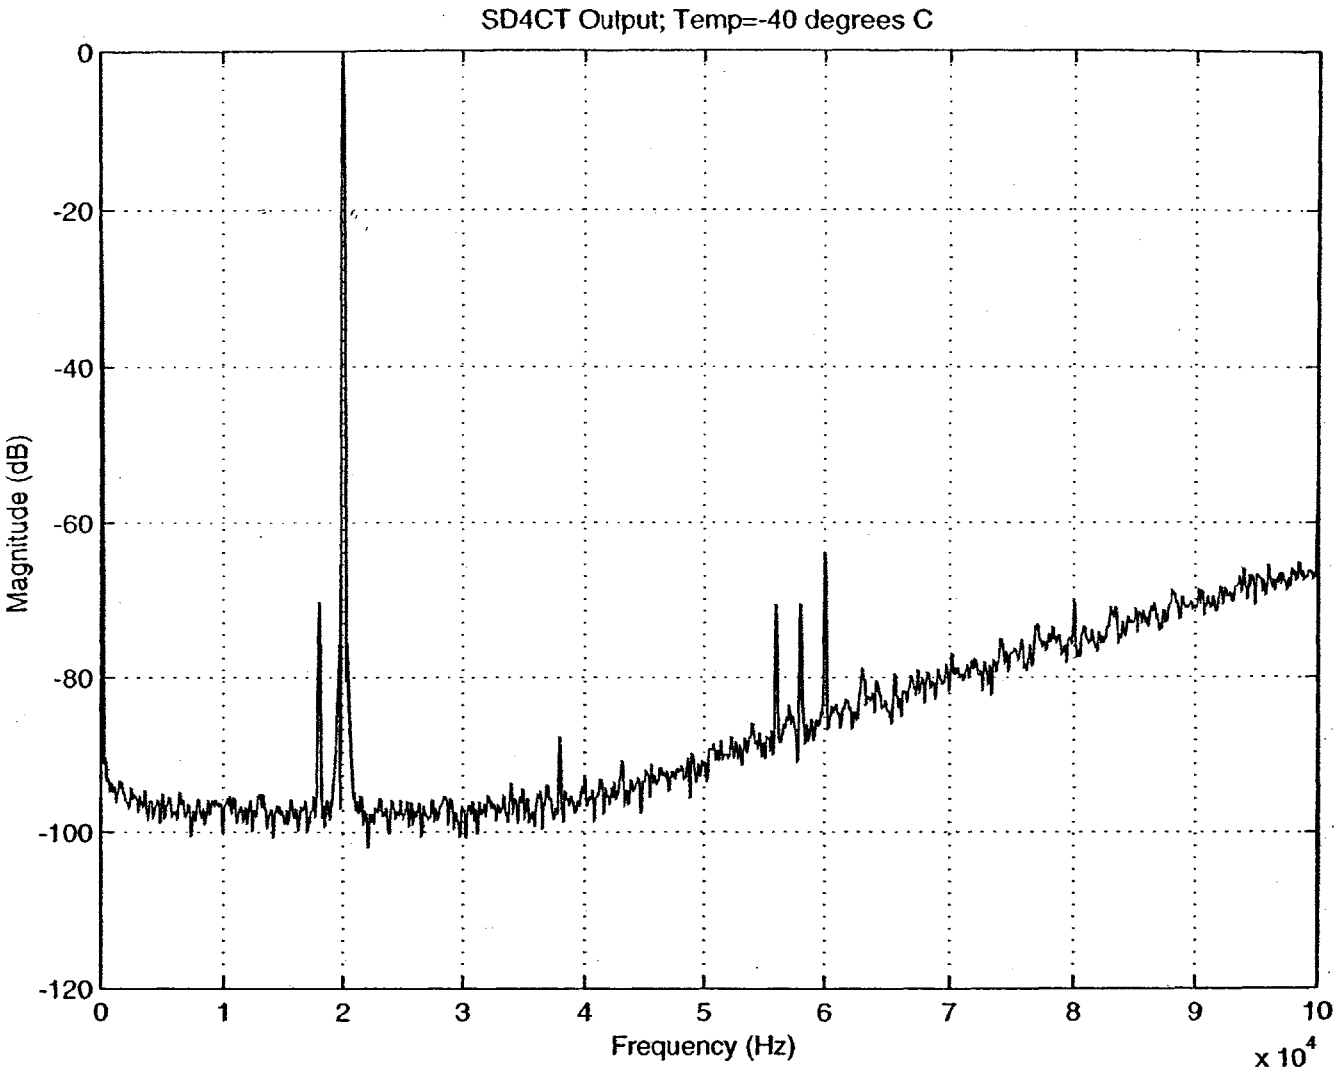
SD4CT Output; Temp=-40 degrees C
Magnitude (dB)
0
-20
-40
-60
-80
-100
-120
0
1
2
3
4
5
6
7
8
9
10
Frequency (Hz)
x 10^4

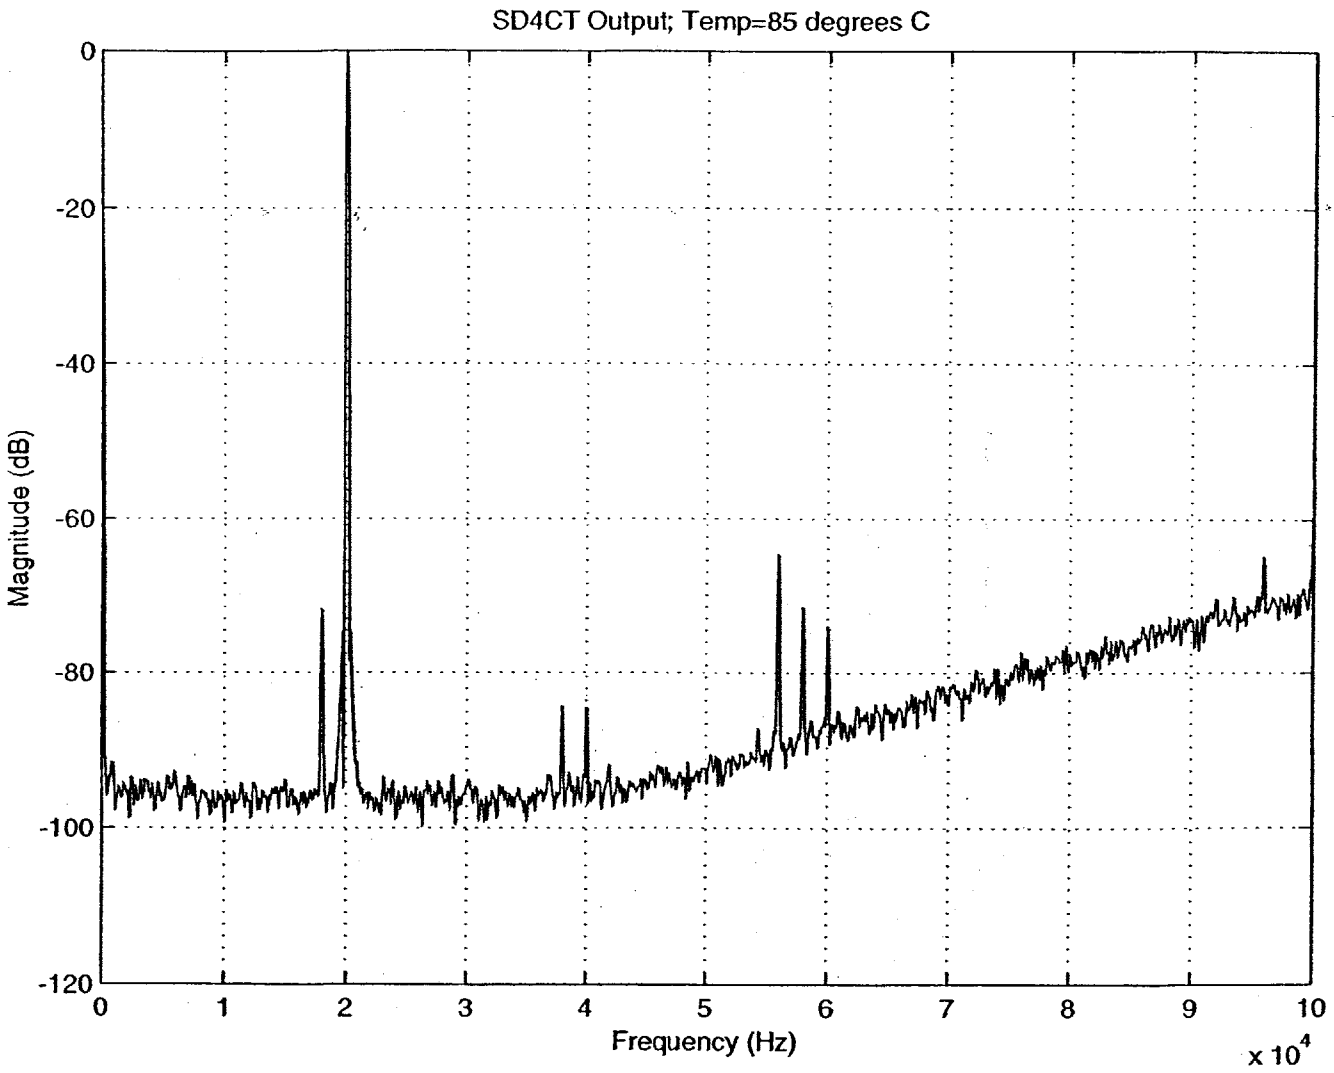
SD4CT Output; Temp=85 degrees C
0
-20
-40
-60
-80
-100
-120
Magnitude (dB)
0
1
2
3
4
5
6
7
8
9
10
Frequency (Hz)
x 10^4

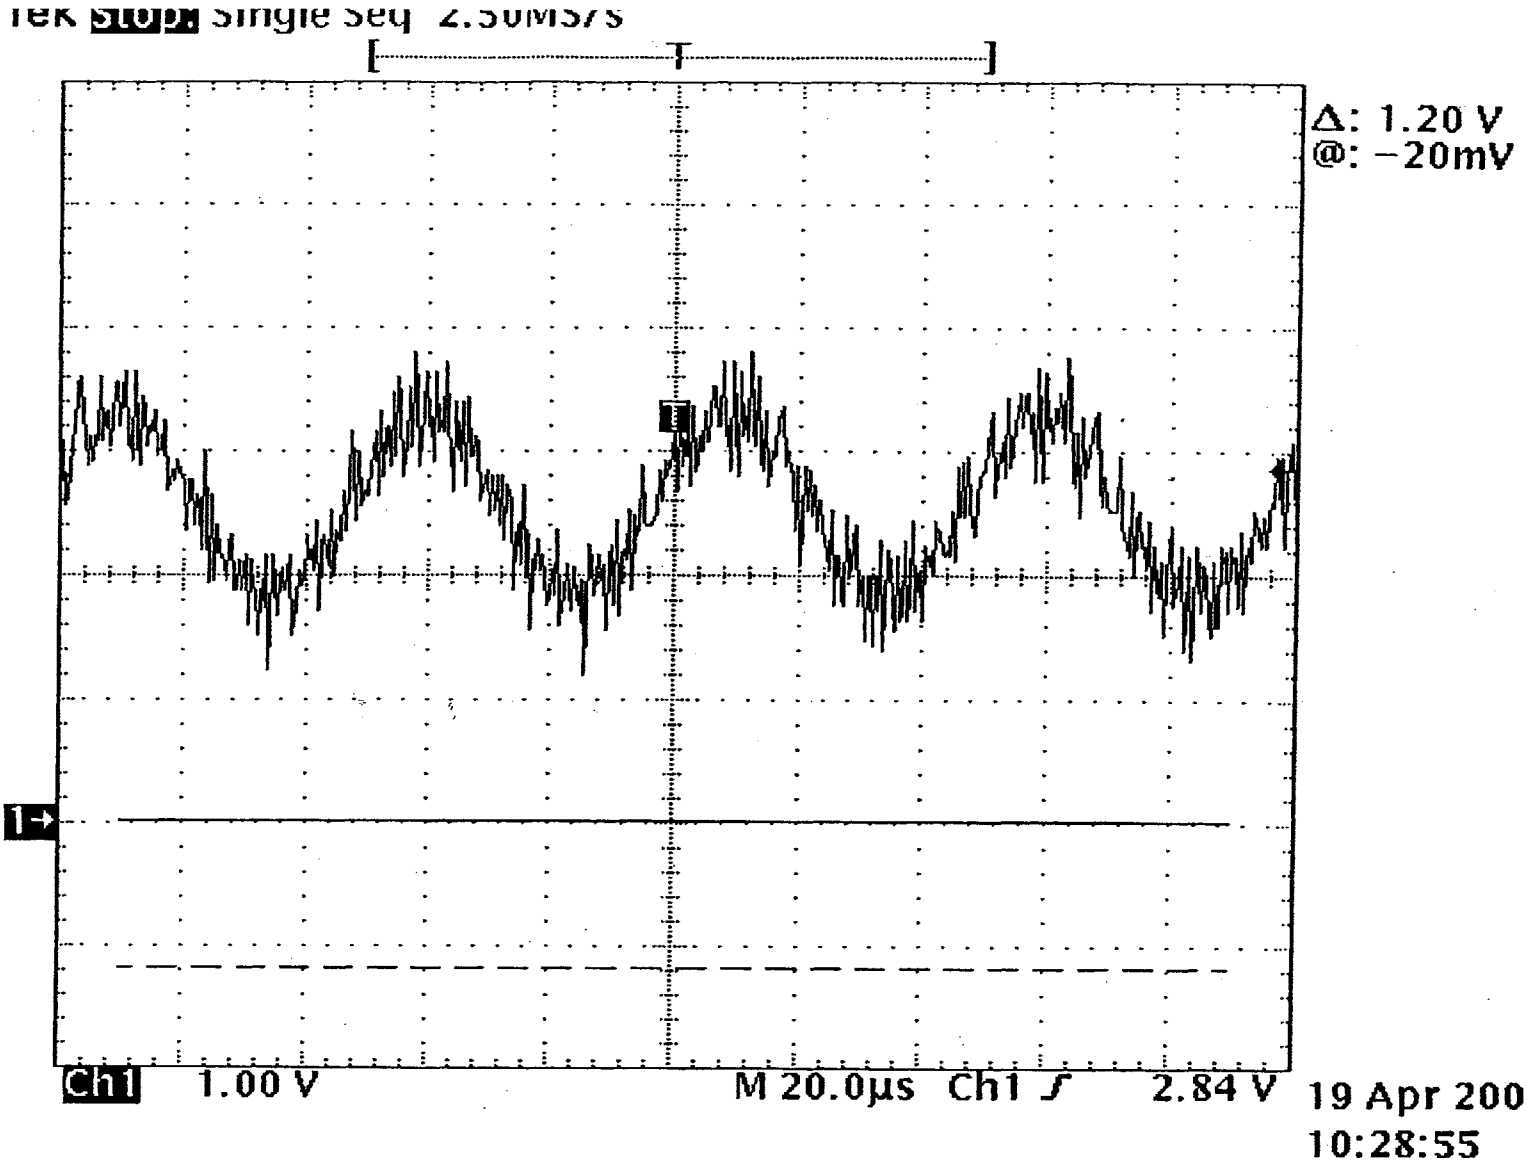

Measured Output of 1st Integrator

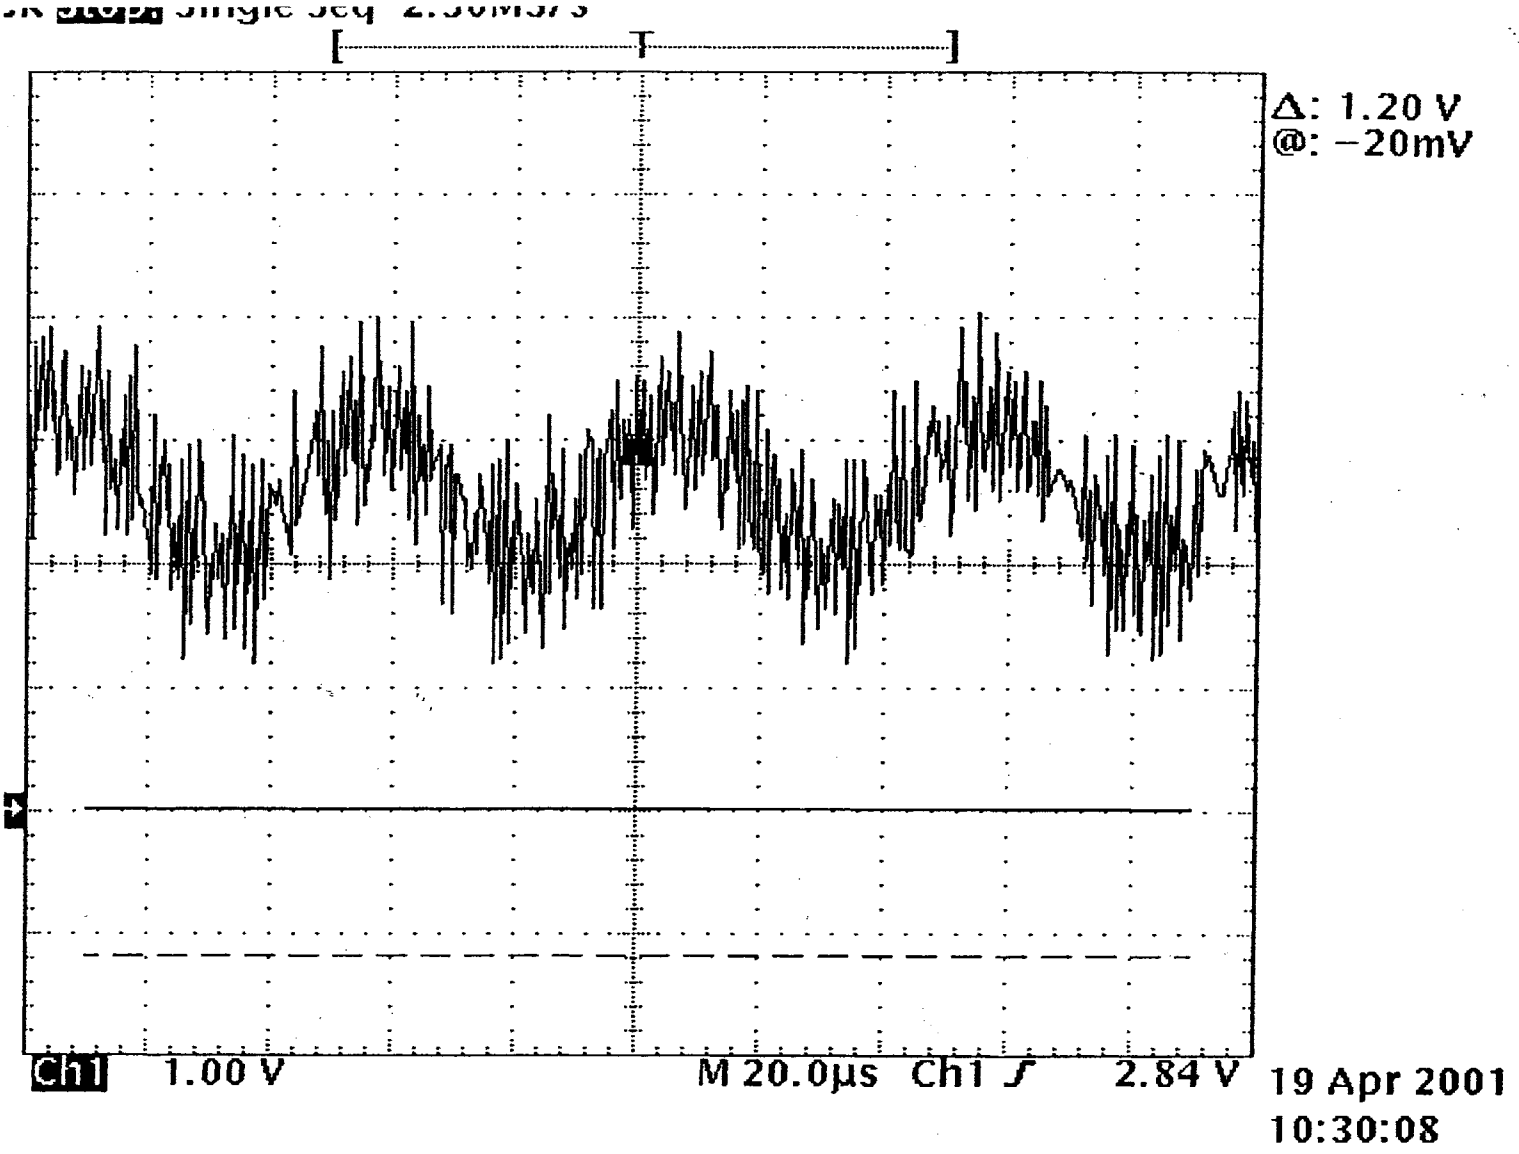

Measured Output of 2nd Integrator

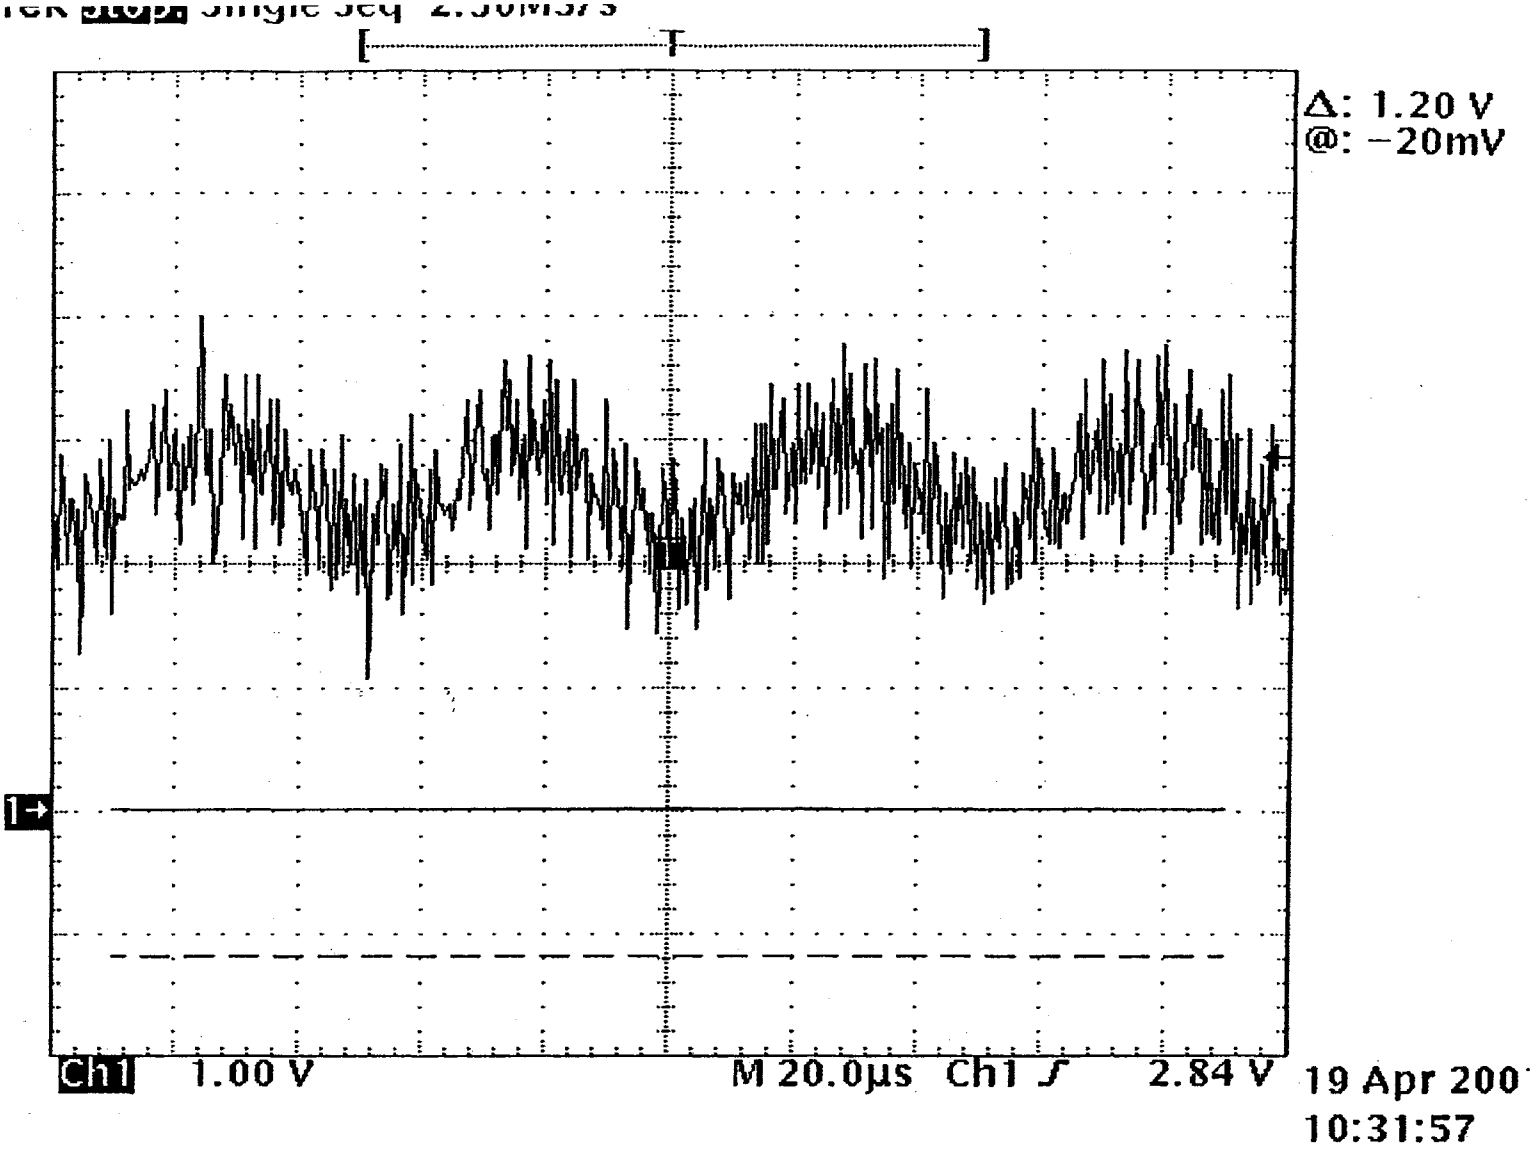

Measured Output of 3rd Integrator

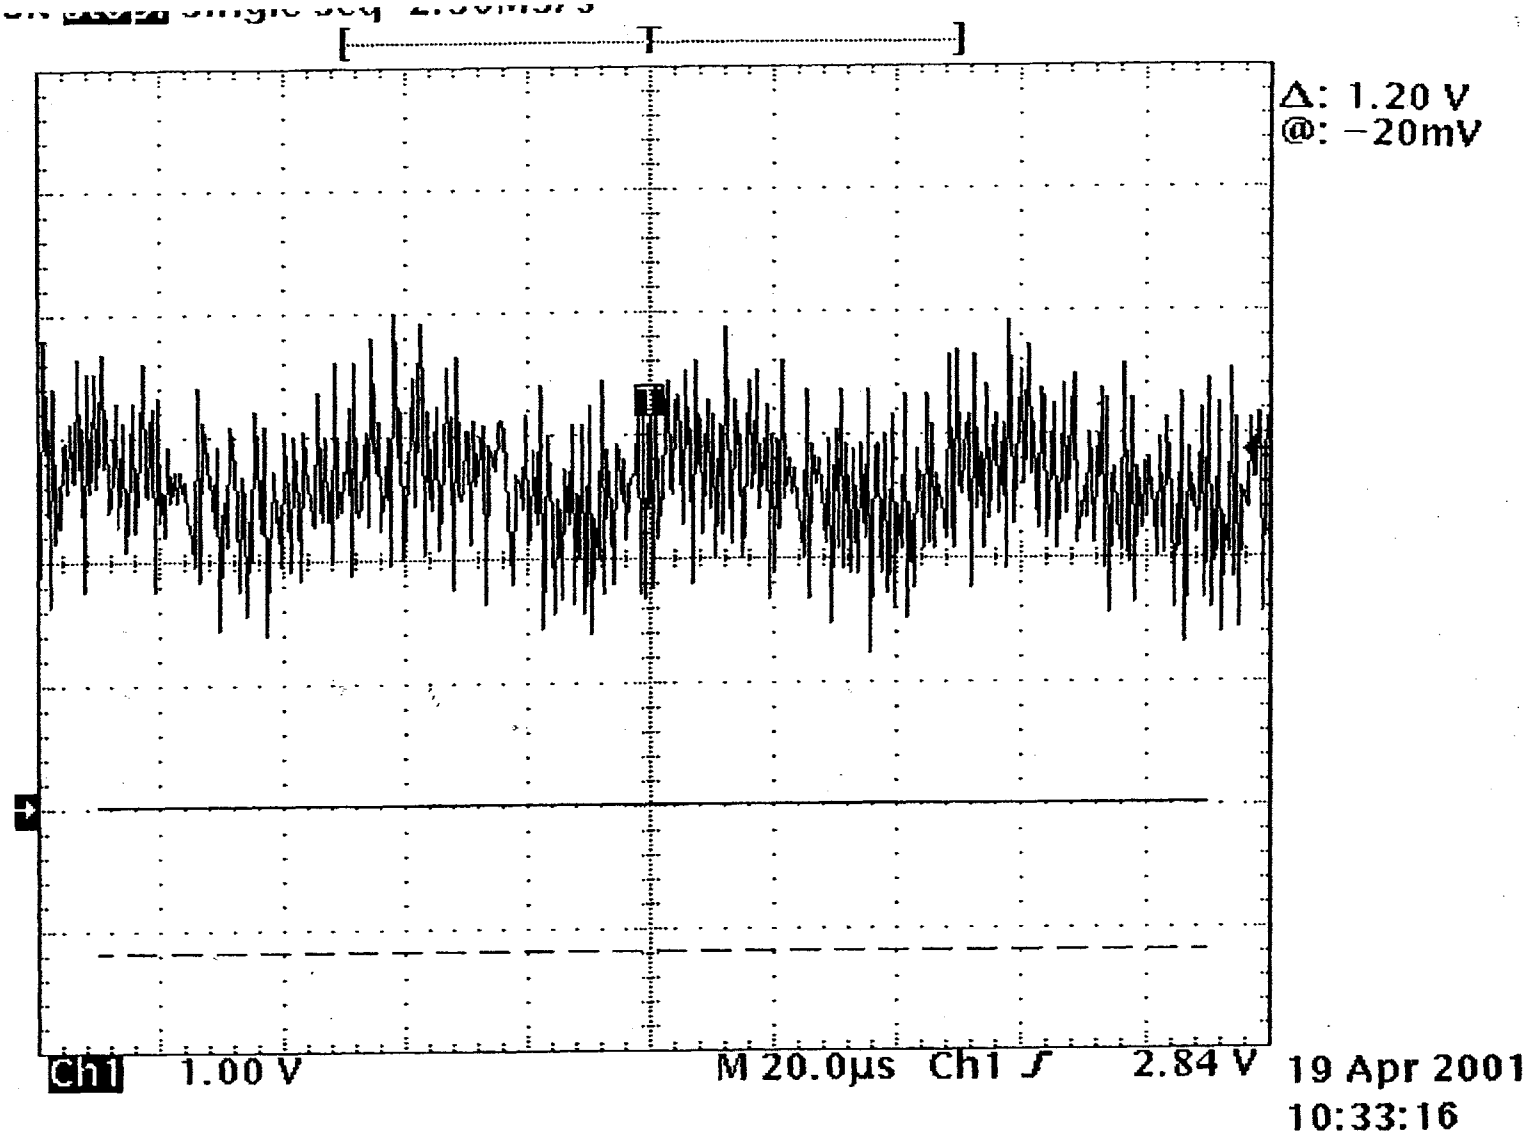

Measured Output of 4th Integrator

Appendix D

Clock Jitter Test Results

- Spectrum of sampling clock phase modulated with white noise (bandlimited to 3kHz)
- Measured SD4CT spectrum using jittered clock
- Measured SD4CT spectrum using 1kHz phase modulation on clock
- Measured SD4CT spectrum using 3kHz phase modulation on clock
- Measured SD4CT spectrum using 5kHz phase modulation on clock

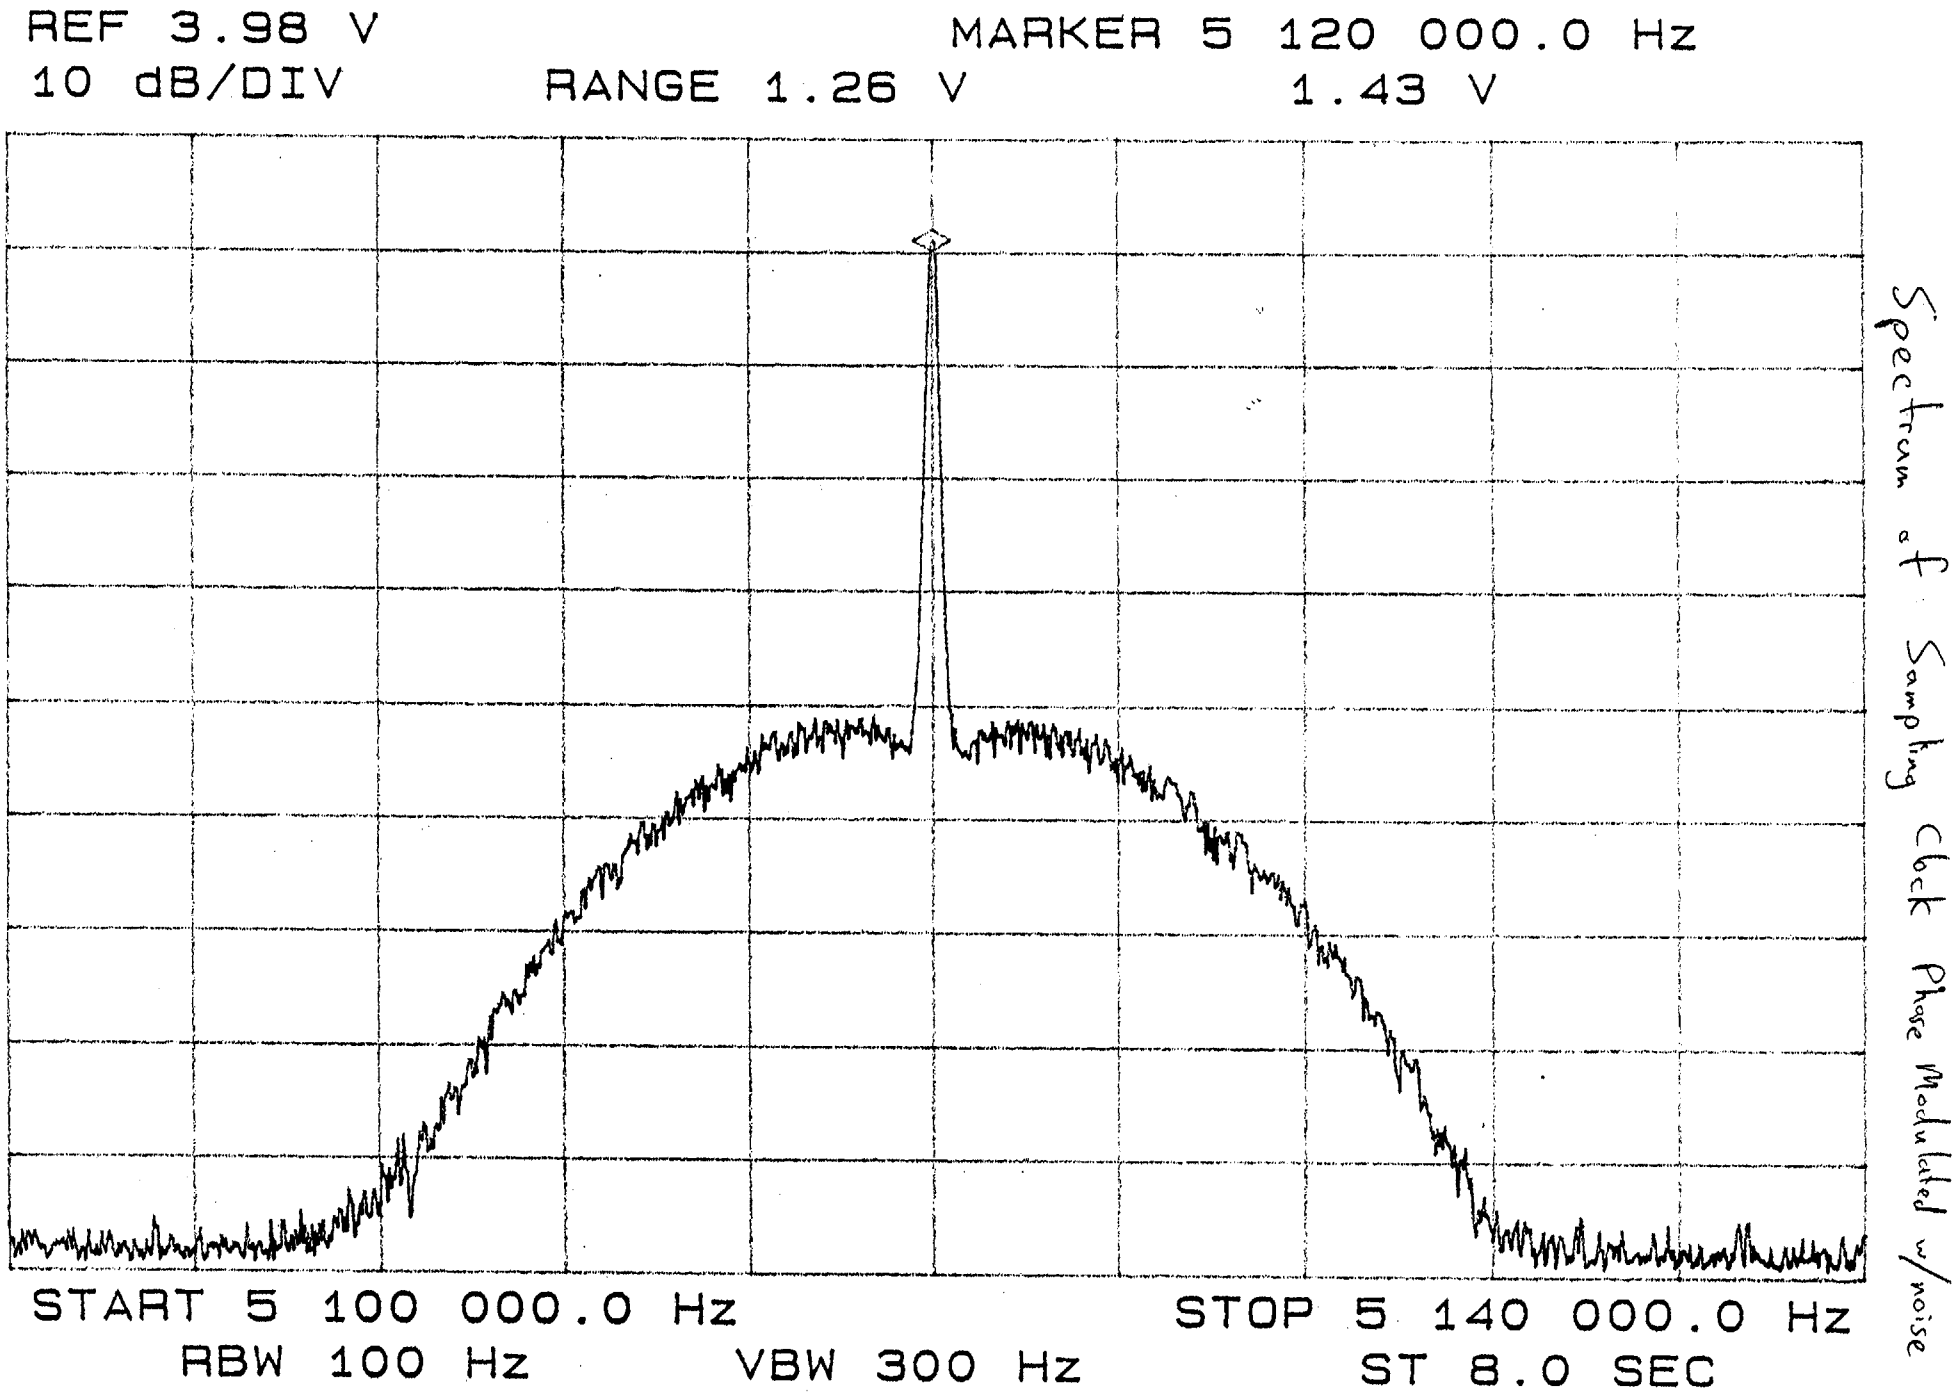

Spectrum of Sampling Clock Phase Modulated w/noise

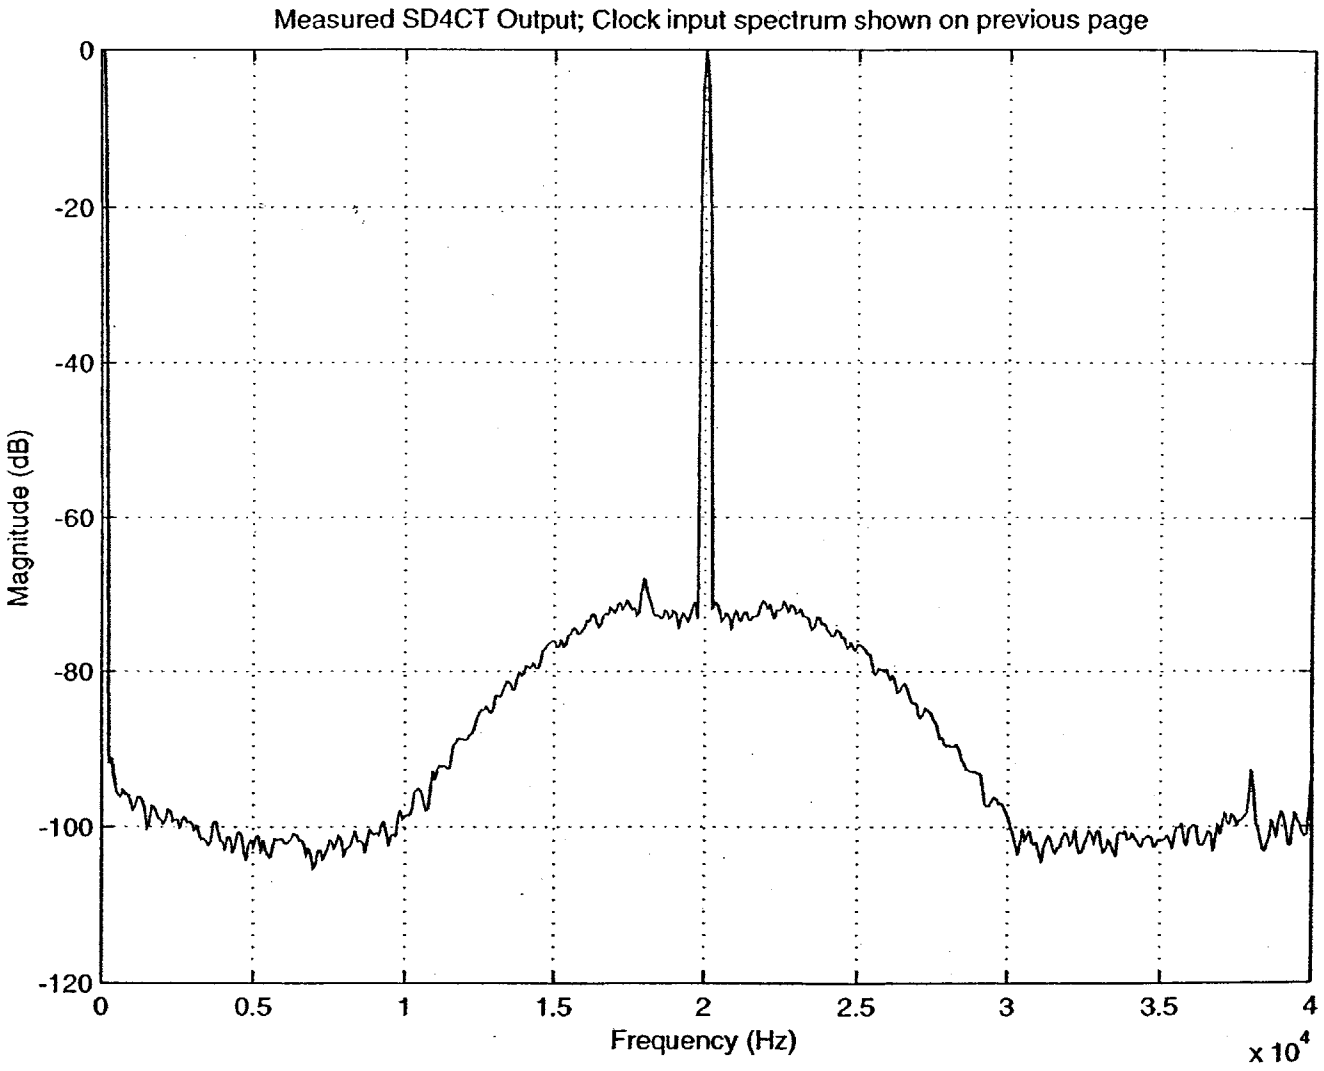
Measured SD4CT Output; Clock input spectrum shown on previous page
0
-20
-40
-60
-80
-100
-120
Magnitude (dB)
0
0.5
1
1.5
2
2.5
3
3.5
4
Frequency (Hz)
x 10^4

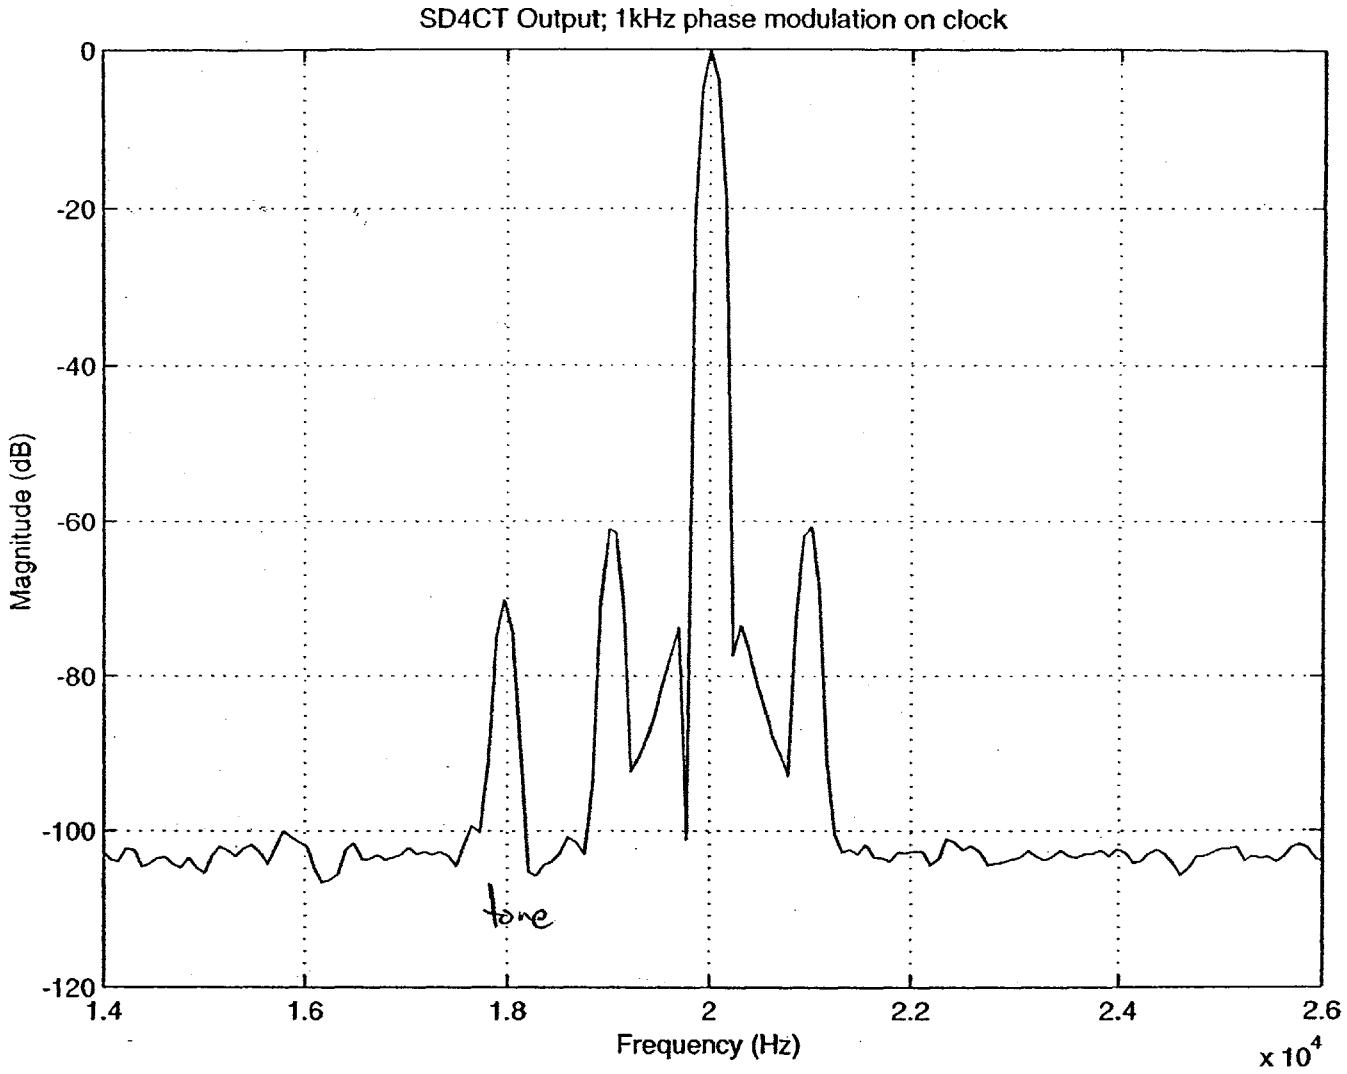
SD4CT Output; 1kHz phase modulation on clock
0
-20
-40
-60
-80
-100
-120
Magnitude (dB)
1.4
1.6
1.8
2
2.2
2.4
2.6
Frequency (Hz)
x 10^4
tone

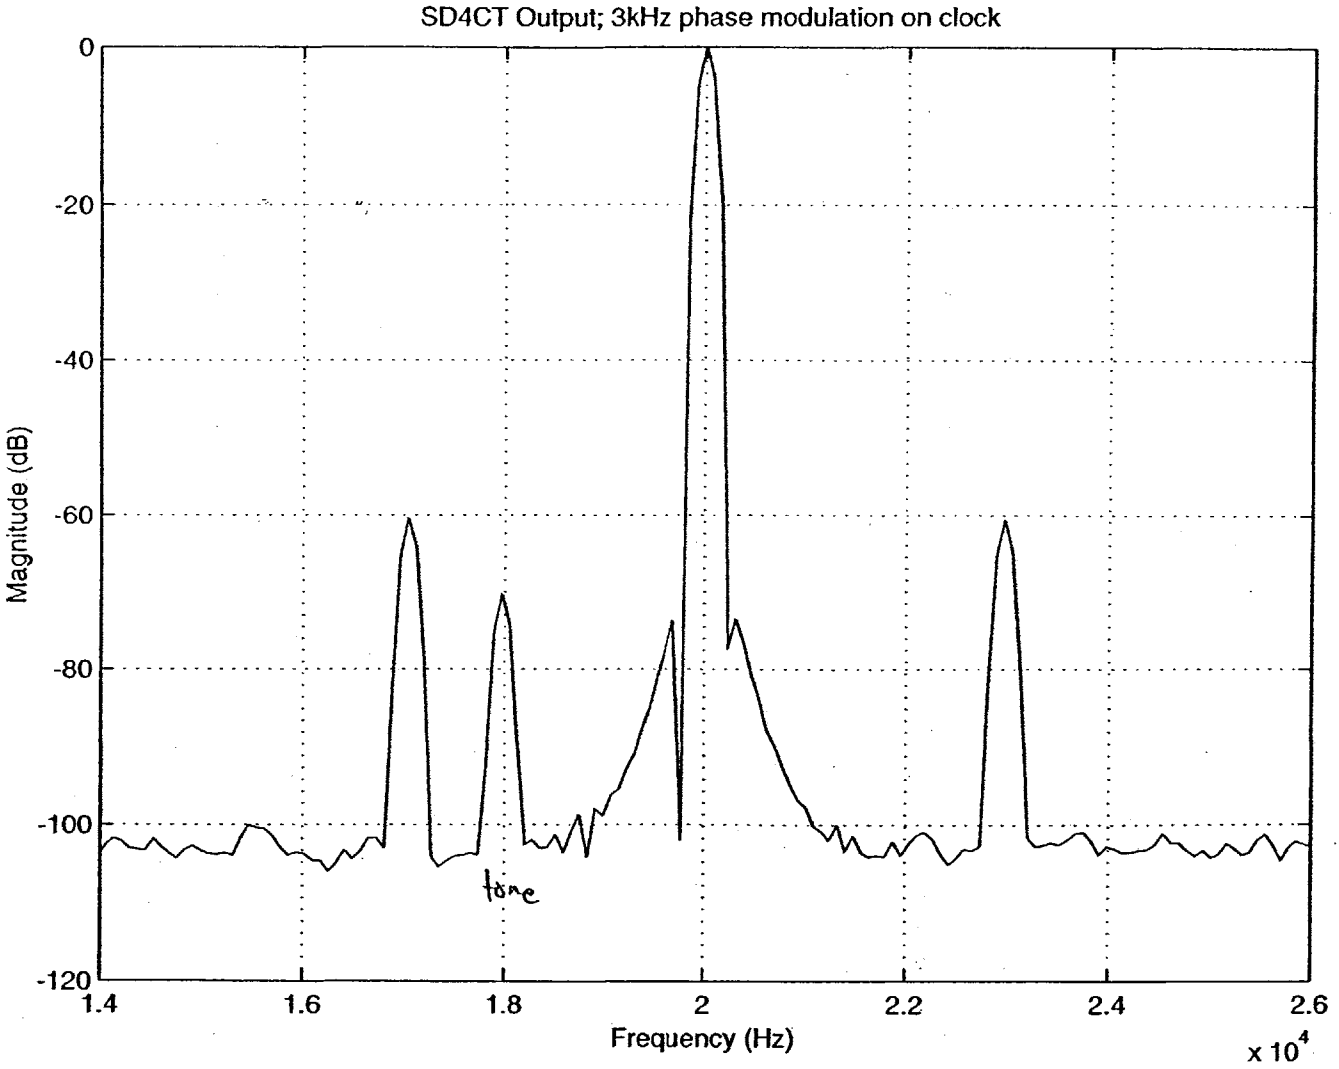
SD4CT Output; 3kHz phase modulation on clock
Magnitude (dB)
0
-20
-40
-60
-80
-100
-120
1.4
1.6
1.8
2
2.2
2.4
2.6
Frequency (Hz)
x 10^4
tone

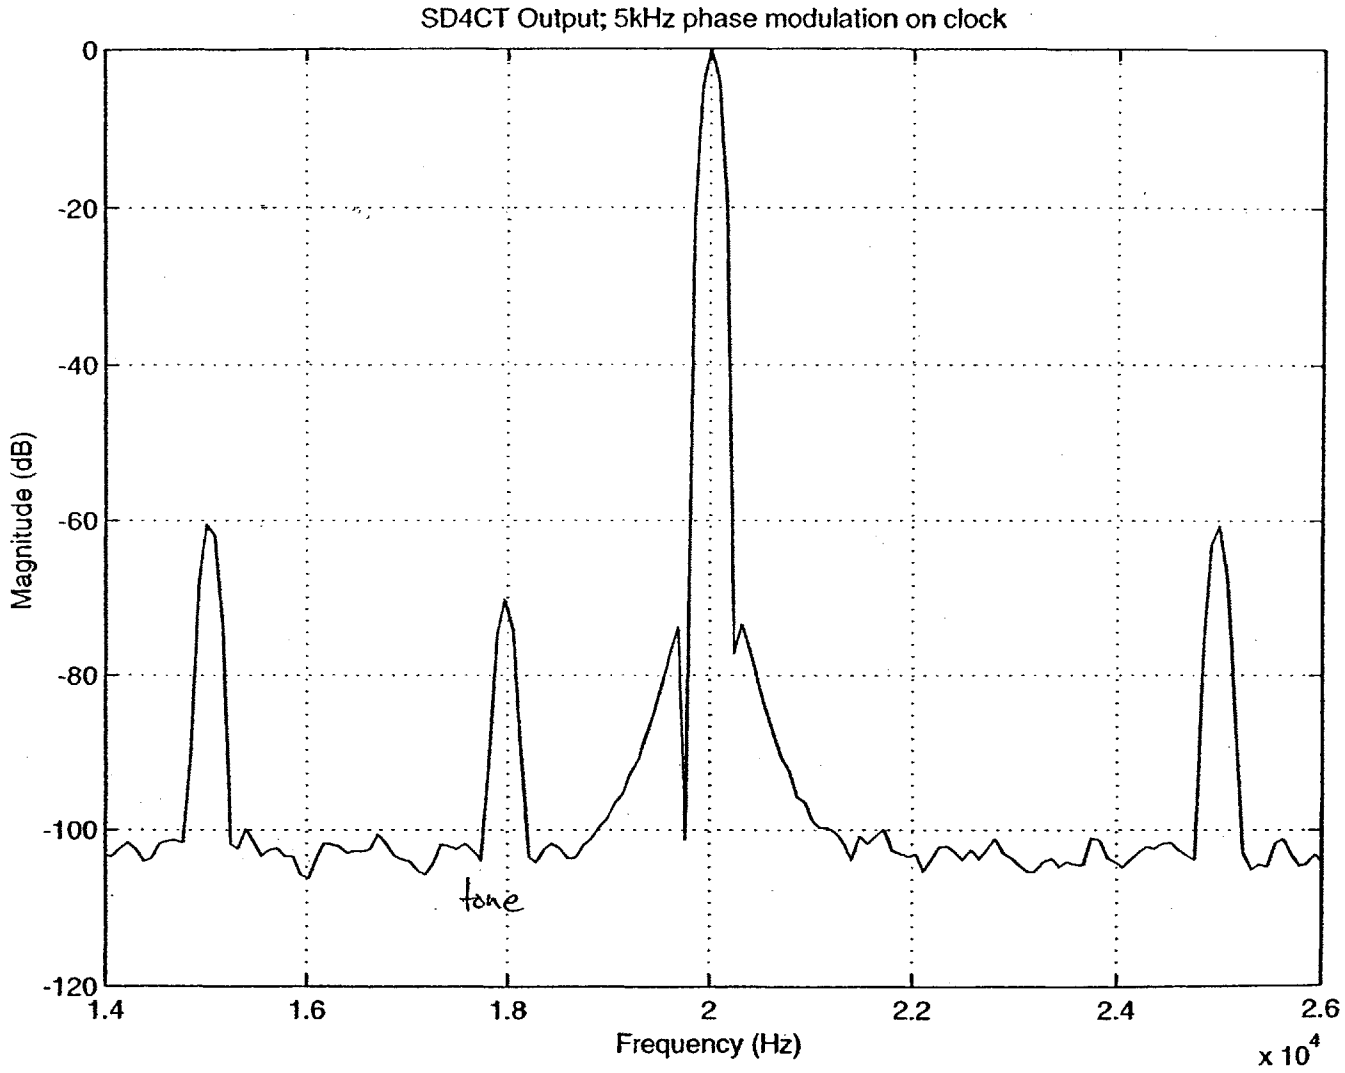
SD4CT Output; 5kHz phase modulation on clock
0
-20
-40
-60
-80
-100
-120
Magnitude (dB)
1.4
1.6
1.8
2
2.2
2.4
2.6
Frequency (Hz)
x 10^4
tone

Bibliography

[1] James A. Cherry and W. Martin Snelgrove. *Continuous-Time Delta-Sigma Modulators for High-speed A/D Conversion.* Kluwer Acadmic Publishers, 2000.

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

[3] Ivan H. H. Jorgensen. A 3rd order low power si sigma-delta a/d converter for voice-band applications. *Proc. IEEE International Symposium on Circuits and Systems*, 1, June 1997.

[4] Shujaat Nadeem. *A 16-Channel Oversampled Analog-to-Digital Converter for Multichannel Applications.* PhD thesis, Massachusetts Institute of Technology, 1993.

[5] Steven R. Norsworthy, Richard Schreier, and Gabor C. Temes. *Delta-Sigma Data Converters.* IEEE Press, 1997.

[6] O. Oliaei and H. Aboushady. Jitter effects in continuous-time sigma-delta modulators with delayed return-to-zero feedback. *IEEE International Conference on Electronics, Circuits, and Systems*, 1, September 1998.

[7] Alan V. Oppenheim and Alan S. Wilsky. *Signals and Systems.* Prentice Hall, 2nd edition, 1997.

[8] Tapani Ritoniemi, Tepp Karema, and Hannu Tenhunen. Design of stable high order 1-bit sigma-delta modulators. *Proc. IEEE International Symposium on Circuits and Systems*, 4, May 1990.

[9] James K. Roberge. *Operational Amplifiers.* John Wiley and Sons, 1975.

[10] Keith Santarelli. A low power digitizer for precision carrier band measurements. Master's thesis, Massachusetts Institute of Technology, 2000.