Design of a Low Power Capacitive Sensor for a Micromachined Accelerometer

by

Daniel Good

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

Bachelor of Science in Electrical Engineering and Science

and

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2002

Author .....
Department of Electrical Engineering and Computer Science
May 24, 2002

Certified by .....
Rahul Sarpeshkar
MIT Thesis Supervisor

Certified by .....
Timothy Denison
VI-A Company Thesis Supervisor

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

Design of a Low Power Capacitive Sensor for a Micromachined Accelerometer

by

Daniel Good

Submitted to the Department of Electrical Engineering and Computer Science
on May 24, 2002, in partial fulfillment of the
requirements for the degrees of
Bachelor of Science in Electrical Engineering and Science
and
Master of Engineering in Electrical Engineering and Computer Science

Abstract

A new possible commercial application for a micromachined accelerometer is for use in handheld devices such as personal digital assistants and cellular phones, as an intuitive method of data entry which does not consume space on the ever-shrinking exterior. These devices are battery powered, which imposes stringent power consumption limitations on any hardware added. To make an accelerometer a viable addition to a handheld device, a low power version was designed, with the other device parameters, such as measurement range and noise performance, tailored to be suitable for use in a handheld. The final design measures ± 3 g's of acceleration with approximately 140 $\mu g/\sqrt{Hz}$ of noise, consuming only 200 μW of power.

Thesis Supervisor: Rahul Sarpeshkar
Title: Assistant Professor

Thesis Supervisor: Timothy Denison
Title: Design Engineer, Analog Devices

Acknowledgements

I want to thank Tim Denison for making this possible. There's no way I would have finished without knowing he was behind me all the way.

Endless technical help was provided by my coworkers at Analog Devices and by my fellow grad students in Prof. Sarpeshkar's lab. You guys are all excellent.

Finally, I would have assuredly gone insane had it not been for my family and friends. They provided me with love and support, and most importantly, made me stop working now and again.

Contents

List of Figures

Chapter 1

Introduction

Micromachined accelerometers have become a cheap and reliable way to measure acceleration. They contain a mechanical structure which reacts to an externally applied acceleration, and an electrical system which measures and processes the mechanical reaction. Several types of microelectromechanical structures have been used as accelerometers, but the most feasible for implementation on a single piece of silicon is a capacitive sensing structure. This consists of a set of tiny capacitors which vary in size in response to an acceleration. The electrical section measures and outputs the size of the capacitors.

One of the more appealing applications for a micromachined accelerometer is for use in handheld devices. The accelerometer's tiny size makes it easy to fit in with conventional electronics in a small package. However, when used in a handheld device, power consumption of the accelerometer becomes an important issue. Incorporation of an accelerometer into the design of a device such as a PDA or cell phone requires that the accelerometer not significantly reduce the battery lifetime of the device. The effective power consumption of a high power accelerometer can be reduced through techniques such as power cycling, but having a fundamentally low power device is significantly more useful.

It is difficult to design a low power device because power consumption is necessary to achieve a low noise measurement. The noise of the electronics sets the smallest signal which can be measured, so to get an accelerometer with high resolution, the noise level must be kept to a minimum. Reductions in noise can be accomplished through improvement in design, but past a certain point, the only way to improve

noise performance is through an increase in power consumption. To build a low power accelerometer, power consumption can be reduced in sections which do not contribute to noise performance, such as biasing, or power can be used more efficiently in sections which do contribute to noise performance.

This paper details the design and implementation of the capacitive sense electronics and electromechanical interface for a low power accelerometer. All of the electronics necessary to process the mechanical signal were designed, simulated, and laid out. In order to limit the scope of the project, the electronics were designed around a preexisting mechanical structure, rather than designing a custom structure for this particular accelerometer. To maximize performance for a commercial design, a custom structure would have to be built, but the research purposes of this project are not so demanding.

Chapter 2

Background

The idea of integrating mechanical structures alongside electrical circuits has been around for over 20 years. [1] The close interface of electrical and mechanical structures allows a broad range of sensing and signal processing technologies which were not previously possible, which are broadly known as microelectromechanical systems, or MEMS. MEMS technology has enabled fully integrated accelerometers and gyroscopes, instrumentation electrometers, high-Q mechanical filters, and mechanical energy storage and processing. Monolithic integration and tiny mechanical parts make possible great performance improvements and very small packages, at the expense of manufacturing and process difficulty. [2]

Microelectromechanical systems have their mechanical parts built from the silicon normally used to create transistors. A structure is etched into the silicon and released by one of several different methods. Tradeoffs between the methods include reliability, quality of mechanical structure, and ease of integration with electronics. The micromachining technique used in this project is silicon-on-insulator micromachining, where the structure is carved from the top layer of a wafer with a premade silicon oxide silicon sandwich. This process has the most electrical flexibility, allowing integration with modern silicon-on-insulator electronics, and allowing the fabrication of electrically isolated but mechanically coupled structures. However, it is the most difficult to fabricate, requiring several extra process steps. Additionally, the mechanical structures created with SOI MEMS are thinner than those created by other techniques, which degrades the quality of the mechanical signal. [3]

Currently, the only commercially viable micromachined products are accelerometers used as a car crash sensors. In the event of an accident, they sense the rapid deceleration of the car and trigger the airbag. The attraction of a MEMS accelerometer in this case is that it has a non-destructive self test function, allowing continual verification that the sensor will properly sense and respond to a crash. The power available in a car dwarfs that necessary to run an accelerometer, so in this case the power consumption of the accelerometer is not a concern.

Different features make MEMS accelerometers attractive for consumer applications. In applications not related to safety, the self test function of the accelerometer becomes less important, while the ability to fit the device into a convenient package becomes essential. A fully integrated accelerometer can be added to a printed circuit board just like any other chip, whereas the size of a conventionally manufactured accelerometer would overwhelm that of the package it would be incorporated into.

A desirable commercial application of an accelerometer is as a data entry device for a cellular phone or personal digital assistant. These are both very popular commercial devices, and have a continually increasing level of functionality incorporated into their tiny packages. An easy and intuitive method of data entry which does not consume space on the face of the device is highly desirable due to their ever-shrinking size. This goal can be accomplished by measuring tilt and shake applied to the device with an accelerometer.

Both cell phones and PDAs have strong restrictions on power consumption, due to their small size and great needs for power in their basic functionality. A cell phone has an RF transmitter/receiver and a PDA has a large screen, which both consume a lot of power. A requirement to any additional hardware added to these devices is that it not significantly shorten the already taxed battery life. In order to open up the large market of data entry to accelerometers, a very low power model is necessary. This project details the development of such a model of accelerometer.

A good design strategy for a low power device is to first find a limit on the power consumption to meet the given specifications, and then design to come as close to that limit as possible. In an analog system, the power consumption is fundamentally limited by the desired signal to noise ration (SNR) [4]. Increasing the SNR requires either increasing the signal or reducing the noise, both of which take power. The maximal signal level is set by the power supply, which is usually specified to the designer beforehand and is not

available as a design parameter. As a result, improving the SNR is usually a problem of reducing the system noise. By carefully analyzing the sources of the noise, and what it takes to reduce them to the desired level, a power consumption limit for a given SNR can be found. In the case of a MEMS device, the problem is complicated further than for a standard analog system, because there are noise sources in the mechanical section, which can be an important contribution to the total noise of the system.

Meeting a power consumption limit based on SNR is impossible, however, because there are other places in the system which consume power and do not contribute to improvement in SNR. These sections include biasing, generation of a temperature compensating reference, and generation of clocks if they are to be used. The power consumption of these sections can be reduced simply by designing with an eye towards low power, and so the power consumption limit is hopefully one which can be approached.

The target specifications for the accelerometer are as follows:

Input: ± 3 g of acceleration

Output: Analog output 0V - 3V at 0.5 V/g, biased at 1.5 V for 0 g

Output Load: 1 MΩ in parallel with 10 pF

Bandwidth: -3 dB point at 100 Hz

Power Consumption: 100 μA drawn from 3 V supply

Noise: 5 mg noise at 100 Hz bandwidth -> 500 $\mu g/\sqrt{Hz}$ white noise

Mechanical Specifications: Beam has nominal gap $g_0 = 1.5$ μm and resonant frequency $f_0 = 4$ kHz. Beam overlap area $l \cdot t = 10^{-9}$ m^2. There is a beam to substrate capacitance of 1.5 pF.

Chapter 3

Electromechanical System Design

3.1 Characteristics of the Mechanical Structure

The most important part of the design of a microelectromechanical system is the design of the mechanical section and analysis of how it fits into the electrical section. Once this has been completed, the electronics can be designed to properly interpret the mechanical signal. A representative drawing of the mechanical section is shown in Figure 3-1, and a mechanical schematic is shown in Figure 3-2. The system has mass m, spring constant k_c, damping R, resonant frequency $f_o = \frac{1}{2\pi}\sqrt{\frac{k_c}{m}}$, and quality factor $Q = 2\pi f_o m/R$.

The input to the mechanical system is an externally applied acceleration. Assuming the spring is massless,

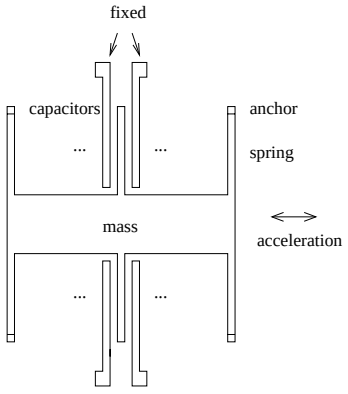

Figure 3-1: Representative Drawing of the Sensing Structure

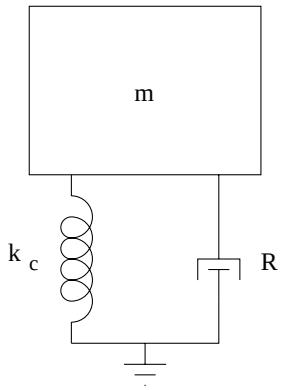

Figure 3-2: Mechanical Schematic

the acceleration appears to the mechanical system as a force on the proof mass, by Newton's Law $F = ma$. The measured output of the system is the compression of the spring, which is dictated by Hooke's Law $F = -k_c x$. By mechanical systems theory [5], the force transfer function is

$$G(f) = \frac{1}{\sqrt{(1 - (\frac{f}{f_0})^2)^2 + \frac{(\frac{f}{f_0})^2}{Q^2}}}. \tag{3.1}$$

The complete transfer function, from input acceleration to output position is

$$H(f) = \frac{m}{k_c} \frac{1}{\sqrt{(1 - (\frac{f}{f_0})^2)^2 + \frac{(\frac{f}{f_0})^2}{Q^2}}} = \frac{1}{(2\pi f_0)^2} \frac{1}{\sqrt{(1 - (\frac{f}{f_0})^2)^2 + \frac{(\frac{f}{f_0})^2}{Q^2}}}. \tag{3.2}$$

This structure is most useful in the frequency band where the amplitude of the transfer function is independent of frequency. This occurs when the frequency of interest is well below the resonant frequency, or $\frac{f}{f_o} \gg 1$. In this range, the transfer function becomes

$$dx(f) = \frac{a(f)}{(2\pi f_0)^2}, \tag{3.3}$$

where $dx(f)$ is in m, and $a(f)$ is in m/s^2.

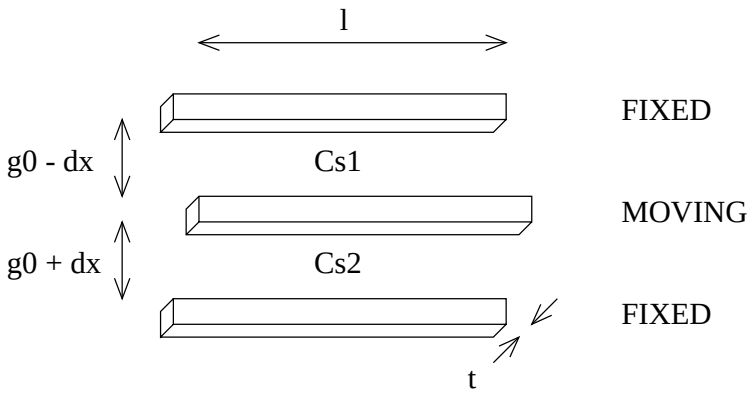

Figure 3-3: Sense Capacitor Geometry

3.2 Electromechanical Interface

The beam displacement is measured by a set of pairs of differential capacitors. Each pair consists of two capacitors with the same nominal value, which change differentially as the sensing structure moves. Their geometry is shown in Figure 3-3. The capacitance of the fingers spaced as shown is

$$C_{s1} = \frac{l \cdot t \cdot \epsilon}{g_0 - dx} \tag{3.4}$$

$$C_{s2} = \frac{l \cdot t \cdot \epsilon}{g_0 + dx} \tag{3.5}$$

thus

$$C_{s1} - C_{s2} = \frac{2 \cdot \epsilon \cdot l \cdot t \cdot dx}{(g_0 + dx)(g_0 - dx)}, \tag{3.6}$$

where ϵ is the permittivity of the gap, and dx is the change in position of the beam.

Extracting the difference of the two sense capacitors gives an output which has the total signal strength. If the difference is normalized by the sum of the sense capacitors, a voltage signal linear with respect to the external acceleration is created. A circuit to calculate the difference of the sense capacitors using switched capacitor techniques is shown in Fig. 3-4(a). Its output is held at virtual ground by a feedback amplifier, and the current coming from the capacitors is measured. Each cycle of the clock, charge is delivered to the output equal to

$$Q = V_{DD}C_{s1} - V_{DD}C_{s2} \tag{3.7}$$

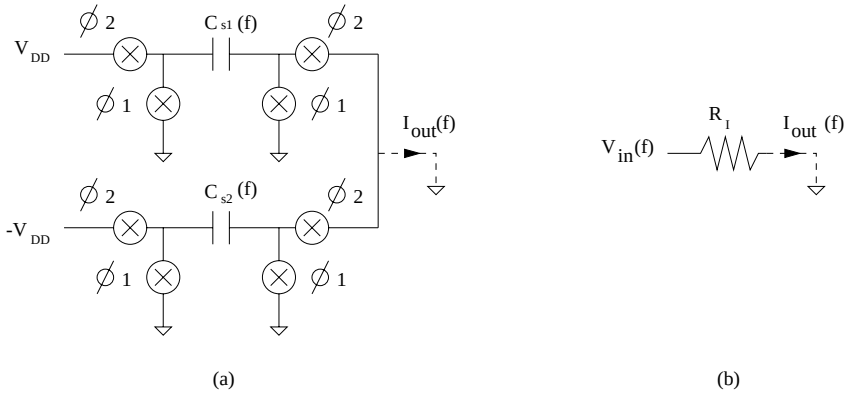

Figure 3-4: (a) Switched capacitor circuit to extract $C_{s1} - C_{s2}$. (b) Low frequency Thevenin equivalent circuit.

$$Q = \frac{2 \cdot V_{DD} \cdot \epsilon \cdot l \cdot t \cdot dx}{(g_0 + dx)(g_0 - dx)}, \tag{3.8}$$

so the total current is

$$I_{out}(f) = \frac{2 \cdot V_{DD} \cdot \epsilon \cdot l \cdot t \cdot f_{clk} \cdot dx}{(g_0 + dx)(g_0 - dx)}. \tag{3.9}$$

By equation 3.3,

$$I_{out}(f) = \frac{2 \cdot V_{DD} \cdot \epsilon \cdot l \cdot t \cdot f_{clk}}{(g_0 + dx)(g_0 - dx)} \cdot \frac{a(f)}{(2\pi f_0)^2}. \tag{3.10}$$

provided f is considerably below f_{clk}, which it will be for all circuits used in this paper. I_{out} can be synthesized by a Thevenin equivalent circuit, shown in Fig. 3-4(b). The Thevenin equivalent is useful because it transforms the sinusoidal input from a changing capacitor into a voltage source, permitting standard techniques of analysis to be used on the circuit. In this circuit, by equation 3.10 and Ohm's Law,

$$V_{in}(f) = V_{DD} \cdot \frac{a(f)}{(2\pi f_0)^2}, \tag{3.11}$$

$$R_I = \frac{(g_0 + dx)(g_0 - dx)}{2 \cdot \epsilon \cdot l \cdot t \cdot f_{clk}}. \tag{3.12}$$

The nonlinearity in R_I will be cancelled later, so it is not a concern for us. For the remainder of the paper, the Thevenin equivalent will be assumed to be the input to the circuit.

This equivalent resistance can also be derived using the standard formula for a switched capacitor resistor,

$$R = \frac{1}{f_{clk}C} \Rightarrow R_I = \frac{1}{f_{clk}(C_{s1} - C_{s2})}. \tag{3.13}$$

Chapter 4

Electrical Architecture Design

4.1 Top Level Architecture

The electrical architecture of this system must be chosen carefully to achieve low power while achieving the desired functionality and maintaining a reasonable level of noise. The architecture must demodulate the switched signal coming off of the mechanical beam and amplify it to a specified signal sensitivity.

The first design choice made in the electrical architecture was to use an integrator as the interface to the beam. This is advantageous over a frequency independent amplification for several reasons. Most importantly, putting an integrator into the system adds a state variable which holds the output value during the reset phase. This state performs the demodulation from a switched output to a continuous one, allowing the loop bandwidth to be well below the switching frequency.

Having the transfer function of the interface to the beam fall off immediately after the frequencies of interest is important for the filtering of high frequency noise. That is particularly important for accelerometer applications, as high frequency acceleration noise tends to be of high amplitude. If the approach taken is to filter later, noise can saturate the interface circuit, causing it to not be able to respond to signals of interest.

A final benefit of an integrating interface stage is that compensation of the feedback loop around the integrator is very simple. The pole location of the interface is set by the size of the feedback capacitor.

There are many possible integrator topologies, each of which has various tradeoffs. Several of these

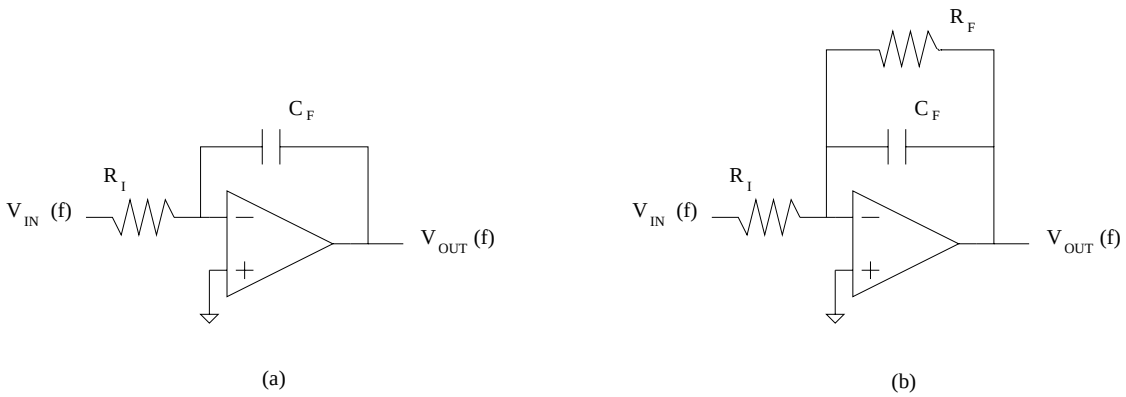

Figure 4-1: (a) Miller integrator. (b) Miller integrator with DC feedback.

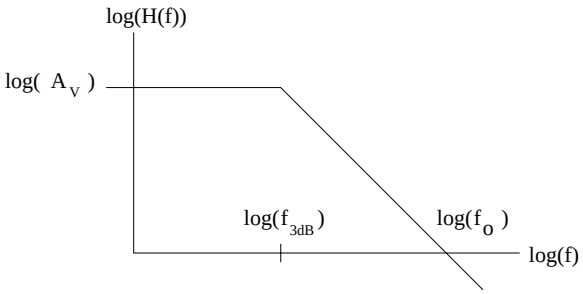

Figure 4-2: Frequency response of Miller Integrator with DC feedback.

are analyzed with respect to SNR to power consumption ratio in [6]. The most power efficient is a Miller integrator, due to its use of an active amplifier but no active transconductance element. All of the integrators function by performing a voltage to current conversion and integrating the current with a capacitor. The Miller integrator converts the input voltage to a current by putting the voltage across a resistor referenced to virtual ground, rather than by using an active element. This circuit is shown in Fig. 4-1(a).

While the integrating characteristic of this circuit is desired, it cannot be allowed to remain an integrator at all frequencies. At very low frequencies, the infinite gain of the integrator would cause the circuit to go unstable in the presence of any DC charge imbalance. To avoid this, a DC feedback path must be provided, ensuring stability and finite gain at low frequencies. The Miller integrator with DC feedback is shown in Fig. 4-1(b) and its frequency characteristics are shown in Fig. 4-2.

Assuming ideal amplifier characteristics, we can define the basic circuit performance in Fig. 4-2 [7].

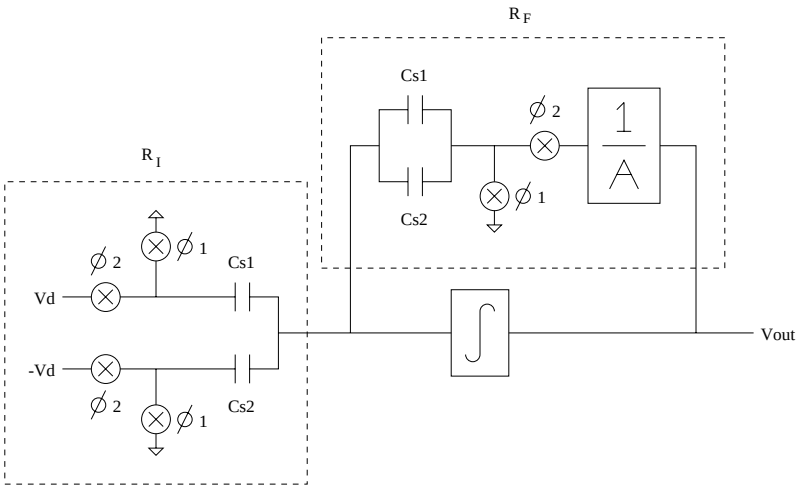

Figure 4-3: Accelerometer Top-Level Block Diagram

$$A_V = \frac{R_F}{R_I}, \tag{4.1}$$

$$f_{3db} = \frac{1}{2\pi}\frac{1}{R_F C_F}, \tag{4.2}$$

$$f_0 = \frac{1}{2\pi}\frac{1}{R_I C_F} = A_V \cdot f_{3db}. \tag{4.3}$$

4.2 Switched Capacitor Implementation

A block diagram of the switched capacitor implementation is shown in Fig. 4-3. This corresponds to Fig. 4-1(b), with R_I and R_F as marked.

In order to meet the desired sensitivity specification (V_{out}/a_{in}) in a well controlled manner, R_F and R_I need to be of approximately the same order of magnitude and fabricated in the same way. R_I is created as a switched capacitor resistor, from the capacitors in the sensing structure. These capacitors are tiny, on the order of 100 fF or smaller, and have a poorly controlled absolute value, due to etch tolerances. In order to fabricate R_F to match R_I in a reliable way, R_F should be also be built as a switched capacitor resistor

from a similar MEMS structure, such as the sense capacitors themselves. To use the sense capacitors in this way, they are driven in common mode, with the differential capacitors added to create a capacitor of approximately constant size. As the sense elements are used in common mode in the feedback path, this technique is known as common-mode feedback.

$$R_F = \frac{1}{f_{clk}(C_{s1} + C_{s2})}, \tag{4.4}$$

$$R_F = \frac{1}{f_{clk}(\frac{l \cdot t \cdot \epsilon}{g_0 - dx} + \frac{l \cdot t \cdot \epsilon}{g_0 + dx})}, \tag{4.5}$$

$$R_F = \frac{(g_0 + dx)(g_0 - dx)}{f_{clk} \cdot 2 \cdot \epsilon \cdot l \cdot t \cdot g_0}. \tag{4.6}$$

Repeating equation 3.12,

$$R_I = \frac{(g_0 + dx)(g_0 - dx)}{2 \cdot \epsilon \cdot l \cdot t \cdot f_{clk}} \tag{4.7}$$

thus,

$$A_V = \frac{R_F}{R_I} = \frac{1}{g_0}. \tag{4.8}$$

In order to keep A_V as a parameter available for design, a voltage attenuator is added in series with R_F. This effectively attenuates R_F by an equivalent factor, and allows A_V to be set to whatever is necessary for the design.

The output voltage of this configuration is:

$$V_{out} = A_V V_{in}, \tag{4.9}$$

$$A_V = \frac{A}{g_0}, \tag{4.10}$$

$$V_{out} = A V_{DD} \cdot \frac{a(f)}{g_0 (2\pi f_0)^2}. \tag{4.11}$$

Given the specified numbers for nominal gap and resonant frequency, and the desired sensitivity of 0.5 V/g, A can be computed to be $\frac{50}{3}$.

4.3 Common-Mode Feedback Design

There are several disadvantages associated with using sense and feedback capacitors as described in the architecture of Fig. 4-3. The first is that there are a limited number of fingers on the sensing structure, and using some of them as feedback reduces the number available for sensing. Fewer sets of sense capacitors reduces the signal strength, creating an effectively noisier sensor. As noise is a major issue in accelerometer design, any unnecessary degradation of the quality is to be avoided if possible.

An additional disadvantage is that the sense and feedback capacitors are being counted on to have identical geometrical properties, so that the nonlinearity cancels properly, and the linear result is achieved. If there is an offset between the capacitors, the nonlinear dx^2 term in R_I will not be cancelled, and the transfer function will differ significantly from that in equation 4.8.

It is therefore highly desirable to use the same set of micromachined capacitors for both the sense and feedback paths. Due to the similarities of the two paths, the voltage signals V_d and $\frac{V_{out}}{A}$ can be combined, and the principle of superposition will ensure that the result stays the same.

The maximum amplitude of the output voltage V_{out} is V_d, the power supply voltage. Thus, the maximum amplitude of the feedback signal is $\frac{V_{out}}{A}$. The sum of the drive voltage and the feedback voltage can only be driven up to V_d, so the drive voltage must be attenuated by a factor of $1 - \frac{1}{A}$, or $\frac{A-1}{A}$. The block diagram modified with the voltage attenuation and additions is shown in Fig. 4-4.

Reduction of the drive voltage by a factor of $\frac{A-1}{A}$ reduces the signal level by the same factor. This is undesirable because it creates an effectively noisier sensor, as described above. The way to get the maximal signal level is to use all of the sense fingers in the forward path, driven by the full supply voltage. However, this makes it impossible to use the sense structure in the feedback path, and prevents the linearization described above. Trading off noise performance for the linearization here is acceptable because the impact to the noise is nearly negligible for large gains, and the ease of design benefits of having a linear signal are great.

The final modification to the feedback architecture is to make it fully differential. The mechanical signal is differential, but becomes a single ended electrical signal when the difference of the sense capacitors is taken.

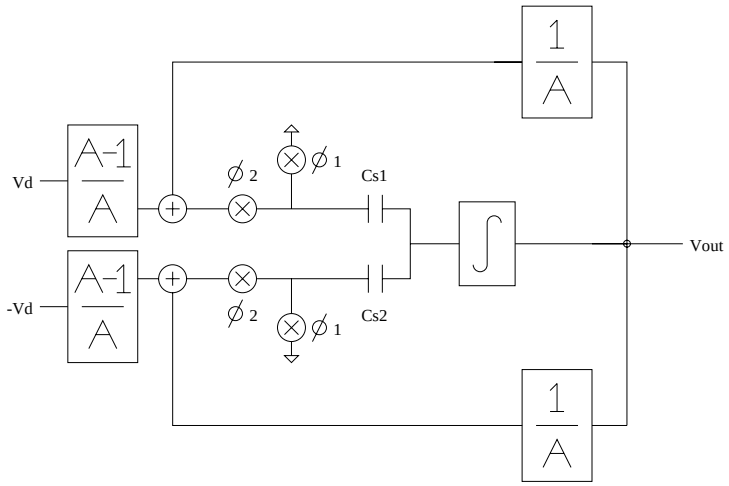

Figure 4-4: Accelerometer Top-Level Block Diagram with Combined Sense and Feedback Paths

The single ended signal path is prone to many accuracy problems, most importantly clock feedthrough and charge injection. If proper care is taken in design and layout of the fully differential path, these inaccuracies will appear as common mode to the integrator and be rejected.

The signal path is modified to be fully differential by essentially making two separate copies of Figure 4-4, one with the polarity reversed by switching V_D with $-V_D$ at the input. A fully differential integrator replaces the single-ended one, with the positive signal path connected to the positive input, and the reverse polarity part connected to the inverting input.

Creating two separate signal paths in this way requires two sets of sense capacitors which respond identically to the input acceleration but which are electrically isolated. This is difficult with standard bulk or surface micromachining, but can be done with SOI micromachining, and is in fact one of the main reasons to use the SOI process. One advantage of SOI electronics in general is that isolation trenches can be built, providing very good isolation between two separate electrical sections. These trenches can be built and filled with a structural dielectric, such as Silicon Nitride, to create multiple sections of the structure which are electrically isolated but mechanically coupled.

4.3.1 Electrostatic Spring Constant

A further benefit of common mode feedback is that the opposite sides of the sense capacitors are driven in such a way there is no net electrostatic force on the beam. If a constant drive voltage is put on the capacitor plates, an attractive force is imparted which is proportional to the gap between the plates. For small displacements, this can be modeled as an electrostatic spring constant, $F = k_e x$. Because it is an attractive force, it acts as a negative spring constant, and its magnitude must be kept below that of the mechanical spring constant to preserve the net repulsive force.

The electrostatic force on the capacitor plates is:

$$F = -\frac{\partial}{\partial x}\frac{1}{2}CV^2 = \frac{1}{2}\frac{\epsilon \cdot l \cdot t \cdot V^2}{x^2} = \frac{C^2V^2}{2 \cdot \epsilon \cdot l \cdot t} \tag{4.12}$$

Common-mode feedback causes the output voltage to drive the capacitors in such a way that the sum of the charge coming off of the two sense capacitors is equal to zero. $Q = C_1V_1 = -C_2V_2$, so the forces imparted on the beam from each side are equal and opposite, and cancel one another out. Since there is no force imparted onto the beam, there is no electrostatic spring constant, and there is no danger of the net spring constant going negative.

4.4 Integrator Architecture Design

The largest consumer of power in the architecture as specified so far is the integrator. The op-amp used in it needs to have high gain to minimize offset errors, and low input-referred noise to make a high resolution measurement. Improving these specifications costs a lot of power, so it is important to get the maximum possible benefit from the power spent on them.

Switched capacitor topologies for minimizing op-amp nonidealities have been discussed by Nagaraj [8] and Temes [9]. These use a correlated double sampling (CDS) approach to effectively attenuate offset and 1/f amplifier noise, and adjust for the finite gain error. The filter contains a zero at DC, which removes steady state amplifier offset. The topology used for this architecture is a slightly modified version of the one presented by Nagaraj to be compatible with the top-level switched capacitor topology while still achieving

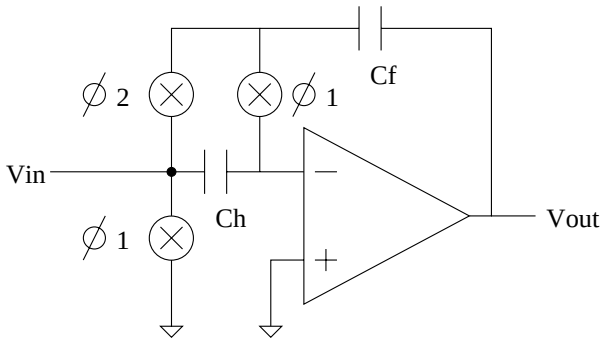

Figure 4-5: Correlated Double Sampling Integrator Architecture

the desired filtering.

The correlated double sampling integrator architecture is shown in Fig. 4-5. It contains an op-amp with a feedback storage capacitor, as is required for an integrator. Rather than connecting the feedback capacitor directly to the inverting terminal of the op-amp, it is switched between the sides of a error hold capacitor. On phase one, the input is connected to ground and the hold capacitor is connected from the inverting terminal of the op-amp to ground, charging it to the voltage at the inverting terminal. Ideally that voltage would be ground, but there are voltage contributions from amplifier offset, amplifier noise, and finite gain error. The feedback capacitor is connected to the inverting input and holds its voltage during phase one, keeping the output voltage constant while the input is disconnected. In phase two, the feedback capacitor is connected around the hold capacitor and the circuit integrates the charge from the input with the hold capacitor subtracting its stored voltage from the input voltage. Qualitatively, this cancels DC offset errors, cancels finite gain errors to the best ability of the amplifier, and removes any components of noise which are present in both phases.

4.4.1 Finite-Gain Error Cancellation

Feedback block diagrams illustrating the cancellation of finite-gain error are shown for the CDS architecture in Fig. 4-6. In phase one, the input is disconnected, and an output voltage V_{OUT} is produced. To create this voltage, there must be some voltage V_{CH} at the input to the amplifier. By Black's Formula [10], the voltage needed there is $V_{CH} = \frac{1}{1+A} V_{OUT}$. This voltage goes to zero as the amplifier gain goes to infinity; it is thus

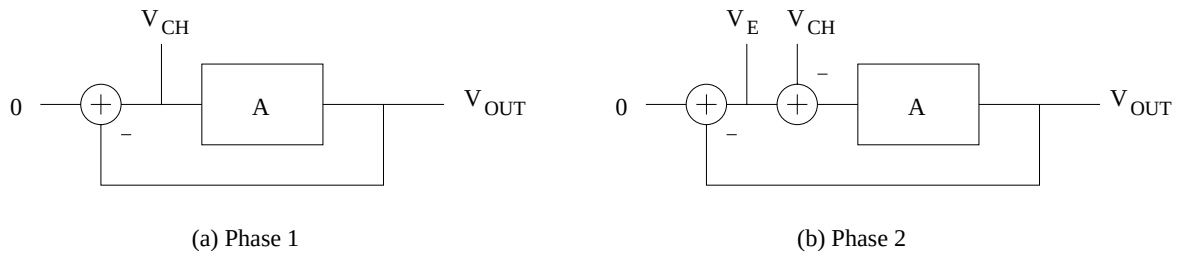

(a) Phase 1 (b) Phase 2

Figure 4-6: Block Diagrams of CDS Architecture in Phase 1 and 2

referred to as finite gain error.

In phase two, the hold capacitor is turned around in the circuit by the switches, and its voltage is subtracted from the error voltage V_E. By superposition of Black's Formula for the signals V_{CH} and V_{OUT},

$$V_E = \frac{1}{1+A}V_{OUT} - \frac{A}{1+A}V_{CH}, \tag{4.13}$$

$$V_E = \frac{1}{1+A}V_{OUT} - \frac{A}{1+A}\frac{1}{1+A}V_{OUT}, \tag{4.14}$$

$$V_E = \frac{1}{(1+A)^2}V_{OUT}. \tag{4.15}$$

The finite gain error for a standard amplifier goes down as A, but for this topology it decreases as A^2. Thus the amplifier gain only needs to be the square root of what would have been required to achieve a given error. The ability to use an amplifier with significantly reduced gain allows a much lower power amplifier to be designed.

4.4.2 Noise Shaping

The noise shaping properties of this circuit can be shown by looking at the frequency transfer function from the error voltage to the charge integration node. The error voltage appears at the inverting terminal to the op-amp, and includes voltage from amplifier offset and noise. The charge is summed on the integration capacitor. Given an error voltage at the inverting terminal e_n, during phase one, an error charge of $C_H e_n[\phi_1]$ is stored on the hold capacitor. In phase two, this charge is subtracted from the error charge $C_H e_n[\phi_2]$, and this difference is delivered to the summing node. So, the total error charge is:

$$Q_e = C_H(e_n[\phi_1] - e_n[\phi_2]). \tag{4.16}$$

Taking z-transforms,

$$Q_e(z^{-1}) = C_H e_n(1 - z^{-1}) = C_H e_n \frac{z^{1/2} - z^{-1/2}}{z^{1/2}}. \tag{4.17}$$

Letting $z^{-1} = e^{-j\omega}$,

$$Q_e(e^{-j\omega}) = C_H e_n \frac{e^{\frac{j\omega}{2}} - e^{-\frac{j\omega}{2}}}{e^{\frac{j\omega}{2}}}, \tag{4.18}$$

$$Q_e(e^{-j\omega}) = C_H e_n \sin(\frac{\omega}{2}) \cdot \frac{2j}{e^{\frac{j\omega}{2}}}. \tag{4.19}$$

For ω small,

$$Q_e(e^{-j\omega}) = C_H e_n \frac{\omega}{2} \cdot 2j = jC_H e_n \omega. \tag{4.20}$$

There are several consequences of this equation. The simplest is that there is a zero at DC in the transfer function, meaning that no DC error charge is integrated. Therefore, any amplifier offset will be cancelled.

The other advantage of the CDS architecture demonstrated by this derivation is the noise shaping properties. Low frequency $1/f$ noise will turn to flat noise near DC. The interface op-amp has to have a MOS input stage to avoid input currents, so $1/f$ noise becomes a significant problem. Usually the only solution to this is to raise the input transconductance by burning more power in the input stage, reducing the input-referred noise. The CDS architecture addresses this problem, once again without burning any significant power. The architecture, however, is not without its tradeoffs.

Looking at equation 4.19 over all frequencies, we can see that the magnitude of the noise is doubled for frequencies π, 3π, 5π, etc., where 2π is the sampling rate. This corresponds to frequencies of noise which have equal and opposite magnitude at the error node in the sample phase and the amplification phase, so that when the two are differenced, they add instead of cancel.

In addition to the noise gain performed by the noise shaping, there is noise gain caused by aliasing. Any sampled data system is going to suffer from this effect, where noise which is at higher frequencies than the sampling frequency is mixed down into the band below the sampling frequency. The entire noise spectrum

appears in the sampled output, effectively amplifying the noise in the band below the sampling frequency by the ratio of the noise bandwidth to the sampling bandwidth.

The power spectrum of the input-referred noise is flat out to the crossover frequency of the amplifier, where it has its pole. Since the amplifier gain is designed to fall of like a single pole near crossover, the effective brick-wall noise bandwidth is $\frac{\pi}{2} f_{0,opamp}$. The noise gain due to aliasing is thus $\frac{\pi}{2} \frac{f_{0,opamp}}{f_{clk}}$. The amplifier is designed to settle to $5\tau = \frac{5}{2\pi f_{0,opamp}}$ in each clock phase, to ensure that the values it is driving the switched capacitor circuits have settled accurately. Therefore, $f_{0,opamp} = \frac{5}{\pi} f_{clk}$, and the aliasing noise gain is a factor of 2.5.

The total noise charge summed on the integration capacitor can now be computed. Noise of different frequencies sums as the square root of the sum of squares, so the total noise is the square root of the integral of the noise squared. The bandwidth of the circuit is 100 Hz, compared with a sampling frequency of 100 kHz, so the frequency band where the noise is integrated is DC to $\frac{2\pi}{1000}$. These frequencies are very small, so the approximation made in Equation 4.20 is valid.

$$Q_{tot} = \sqrt{\int_0^{\frac{2\pi}{1000}} (5\pi C_H e_n \omega)^2 d\omega} = 5\pi C_H e_n * \sqrt{\frac{1}{3}(\frac{\pi}{500})^3} = 4.51 \times 10^{-3} * C_H e_n. \tag{4.21}$$

To see how much of an improvement in noise performance the CDS noise shaping gains us, we can compare this with the total noise for a system with no noise shaping or aliasing.

$$Q_{tot} = \sqrt{\int_0^{\frac{2\pi}{1000}} (C_H e_n)^2 d\omega} = C_H e_n \sqrt{\frac{\pi}{500}} = 79.2 \times 10^{-3} * C_H e_n. \tag{4.22}$$

This is a very significant improvement in noise performance, particularly considering the other gains the switching achieves. If switched-capacitor circuits were used to perform the beam charge sampling as they are in this circuit, but a simple switched-capacitor integrator was used instead of the noise shaping one, the noise would be a factor of the aliasing gain, 2.5, higher than it is computed to be in Equation 4.22. The noise shaping circuit has a reduction in total noise of a factor of 44 from that higher level.

Chapter 5

Circuit Design

The implementation of this architecture took a large effort of circuit design. In addition to the core signal path visible in Figure 4-4, consisting of an integrator, attenuators, and a summation, there are also several support blocks necessary to make the core circuits function according to specifications. To properly turn on the analog circuits and make them function from the single 3V supply, biasing circuitry and generation of a 1.5V "ground" reference is necessary. Clock generation and phasing is necessary to control all of the switched capacitor circuits. Finally, an output buffer is necessary to convert the fully differential internal signal into a single-ended output and drive it into an off-chip load.

5.1 Top-Level Design

The top-level architecture of Fig. 4-4 needs to implemented in such a way that it meets the closed loop gain and bandwidth specifications. The specified gain is 0.5 V/g, and 3 dB bandwidth is 100 Hz.

The circuit topology to implement the top level architecture is shown in Fig. 5-1.

As computed in section 4.2, the attenuation A needed to get the desired closed loop gain is $\frac{50}{3}$. However, due to headroom issues discussed in section 5.3, all of the gain cannot be in the feedback loop. A factor of 5 gain is moved to the output buffer, and the loop attenuation A is reduced to $\frac{10}{3}$.

As this architecture is essentially the integrator shown in Fig. 4-1(b), its bandwidth can be computed

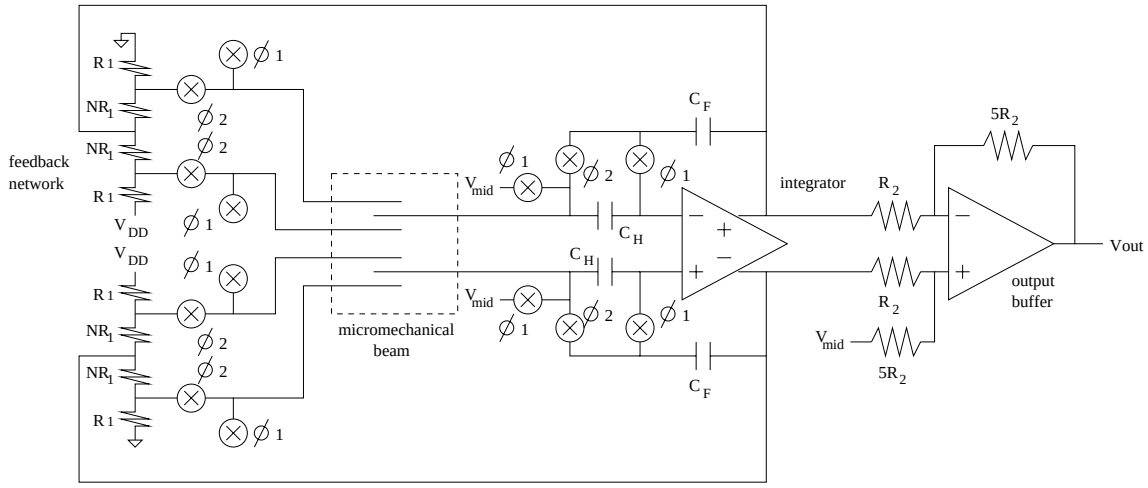

Figure 5-1: Top Level Circuit Schematic

from that given in Equation 4.2. The 3dB bandwidth of the Miller integrator with DC feedback is:

$$f_{3db} = \frac{1}{2\pi}\frac{1}{R_F C_F}. \tag{5.1}$$

From Equation 4.4, $R_F = \frac{1}{fclk(C_{s1}+C_{s2})}$, as it is created by switching the parallel combination of C_{s1} and C_{s2}. The bandwidth-setting capacitor, C_F, is the feedback capacitor used in the integrator. Its value is free to be chosen to whatever is necessary to set the top-level bandwidth appropriately. Its minimum size is limited by the size of the beam parasitic capacitance loading the input to the integrator, as will be explained when the minimum op-amp bandwidth is computed. This constraint is not an issue in this case, as the necessary capacitance for the bandwidth specification is well above the beam parasitic.

The architecture of Fig. 4-4 has an additional attenuation not present in the integrator of Fig. 4-1(b). This attenuation reduces the loop gain by a constant factor across frequency, and thus brings in the 3dB bandwidth by the same factor. So, the 3db bandwidth of the total loop is

$$f_{3db} = \frac{1}{2\pi}\frac{f_{clk}\cdot(C_{s1}+C_{s2})}{A\cdot C_{FB}} \tag{5.2}$$

The circuit was designed using f_{clk} = 100 kHz. Using the specified numbers for beam overlap area, $C_{s1} + C_{s2}$ can be computed to be 200 fF. From above, A = $\frac{10}{3}$. Therefore, for a target 3db bandwidth of

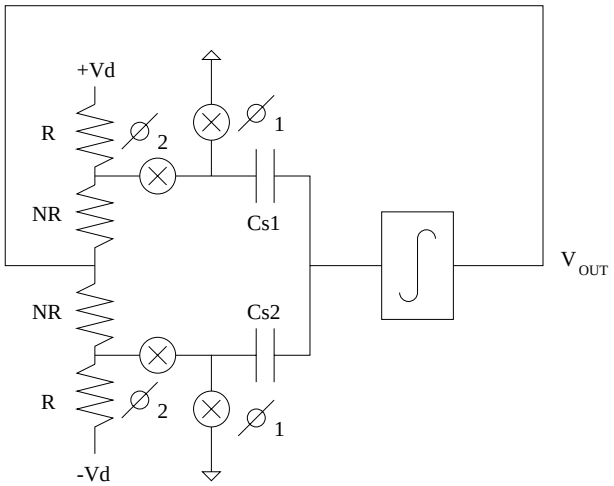

Figure 5-2: Implementation of Feedback Attenuation and Summation

100 Hz, C_{FB} should be set to 9.55 pF. This is a large capacitor to integrate on a chip, particularly as two need to be built for the fully differential architecture. However, it is not so large that it can't be built. If it were impossible to build, an op-amp output stage with wider swing would have to be built, so more of the gain would be in the loop, and A would be larger.

5.2 Feedback Network

The implementation of the feedback attenuation and summation is shown in Fig. 5-2. The attenuation factors of $\frac{A-1}{A}$ and $\frac{1}{A}$ are achieved through a voltage divider of resistors R and NR. The voltage divider attenuates the fed back voltage by a factor of $\frac{R}{R+NR}$, and the drive voltage by a factor of $\frac{NR}{R+NR}$, and sums the two attenuated voltages to the switched node. To achieve the desired attenuation factor, N is set to $A-1$.

This is a useful circuit architecture, performing several desired functions with one simple resistor network. However, it creates several problems for a low power design. The circuit contains a resistive path from $+V_d$ to $-V_d$ thus drawing a DC current equal to $\frac{2V_d}{2R(N+1)}$. To bring the current draw down to be comparable with the power specifications, a total resistance on the order of several hundred kilohms needs to be used. This is within the range that can reasonably be fabricated on a chip, but it is still very large for an integrated

design. It was decided that an acceptable tradeoff between die size and power consumption is to use a total resistance of $200k\Omega$, drawing $15\mu A$.

It was attempted to build the feedback resistors as switched capacitors, since there are already clocks present on the chip. This would allow the creation of much higher impedances in the attenuator, and cut down on the power consumption. However, this presented great difficulty, since the feedback attenuation needs to be very accurate to get the correct gain factor. The attenuated feedback signal is likely to be on the order of tens of millivolts, so small errors from effects like charge injection in the switched capacitor circuit are going to produce unacceptably high gain errors.

Another drawback of the resistive network is the loading effects on the integrator circuit. It creates a resistive load of $\frac{(N+1)R}{2}$ that the integrator needs to drive. For the values of the feedback network chosen, this resistance is $50k\Omega$. This necessitates a low impedance output stage on the integrator which can drive a comparatively large amount of power. Care must be taken to not waste power in this output stage, as a significant fraction of the allotted power for the chip is already being used for the attenuation network.

The power consumption of the feedback attenuator is 7.5 μA $\times$ 3 V = 22.5 μW. There are two feedback networks necessary, one for each side of the fully differential signal path, consuming a total 45 μW of power.

5.3 Integrator Op-Amp

The most important block of signal processing in the accelerometer is the op-amp used to perform the charge integration. It measures the trickle of charge coming off of the micromachined beam, integrates it, and amplifies it into a reasonable output level. It must have low noise so as not to corrupt the input signal, reasonable gain to avoid finite-gain errors, and enough bandwidth to settle in well under the period of the switched capacitor sections.

5.3.1 Specifications

Before designing the op-amp, it is necessary to determine a set of specifications it has to meet. The external specifications on the accelerometer and the design choices made thus far dictate a specific performance level

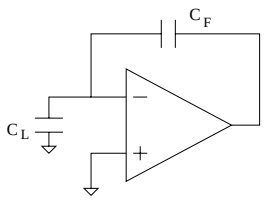

Figure 5-3: Op-Amp Charge Integration Feedback Loop

for the op-amp.

Bandwidth

It was decided to use 100 kHz clocks for the switched capacitor circuits. The system bandwidth is 100 Hz, and a factor of 1000 in frequency ensures that the switched capacitor approximations will be valid. Additionally, the clock frequency needs to be above that of any environmental noise or vibrations which might be picked up by the circuit, to prevent them from being aliased into baseband. The op-amp needs to settle once on each phase of the clock, and should settle to better than 99% of its final value in the clock phase. This is equivalent to settling to five time constants each clock period, or having a settling time constant a tenth of the total clock period. The crossover frequency of the op-amp should thus be a factor of $\frac{10}{2\pi}$ above that of the clock frequency.

At a time scale on the order of the switching period, the closed loop op-amp circuit looks an integrator with a load capacitance at the input equal to the 1.5 pF beam parasitic capacitance, shown in Fig. 5-3. These two impedances form a feedback attenuation equal to $f = \frac{C_F}{C_L+C_F}$. This reduces the loop gain at all frequencies by a factor of f, and thus moves the closed loop bandwidth in by a factor of f. This attenuation puts a limit on the minimum size of the feedback capacitor, as if it is significantly smaller than the load capacitance, the loop will be very slow. It has already been determined that the feedback capacitance is 9.55 pF. The specified parasitic beam capacitance is 1.5 pF, giving an attenuation of 0.86. This feedback attenuation and the settling time specification specify the minimum crossover frequency, at least $\frac{10}{2\pi \cdot 0.86} \approx 1.85$ times greater than the clock frequency, or 185 kHz.

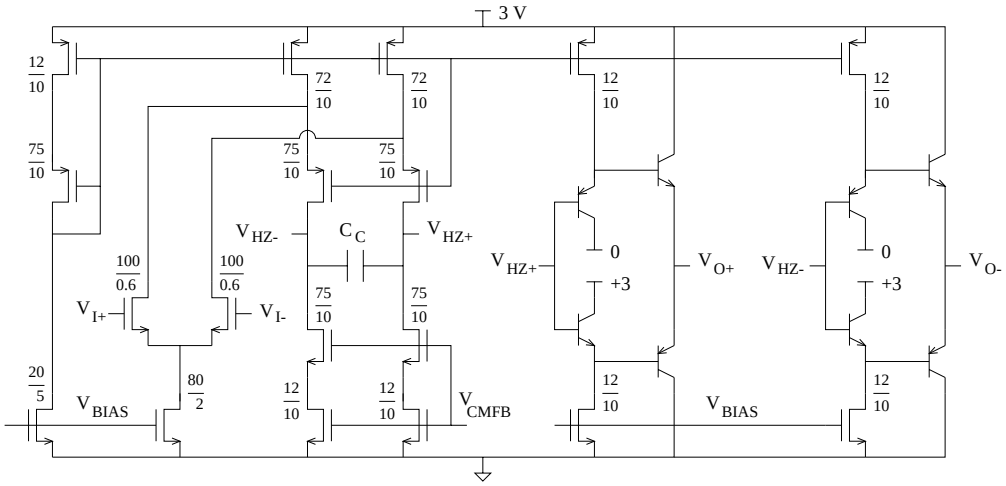

Figure 5-4: Schematic of Differential Signal Path (all dimensions in μm)

Open-Loop Gain

For the finite-gain error of the op-amp to be a negligible part of the total system error, it is desirable to have the open-loop gain greater than 120 dB. With the gain-squaring of the CDS circuit, this corresponds to an open-loop gain of 60 dB, or 1000. This is easily achievable in a single stage running at low power.

Noise Performance

The overall noise target of the chip is white noise of magnitude 500 $\mu g/\sqrt{Hz}$. The total noise will primarily come from the op-amp and from the switches. Since noise sums as square root of the sum of squares, if one noise source is greater than the other, its contribution will dominate the total. So to maximize the benefits of the noise tradeoffs, the circuit should be designed to have equal contributions from the amplifier and from the switches. If the circuit is exactly meeting the spec of 500 $\mu g/\sqrt{Hz}$, each section should generate 353 $\mu g/\sqrt{Hz}$.

By the calculations done in Chapter 7, to have an output noise less than 353 $\mu g/\sqrt{Hz}$, the input referred noise of the op-amp should be less than 126 $nV/\sqrt{Hz}$.

5.3.2 Differential Mode Design

A fully differential folded cascode topology was chosen for the differential signal path of the op-amp. It was chosen for its high gain in a single stage and ease of attaining adequate speed and compensation. The schematic of the differential signal path is shown in Fig. 5-4.

The input differential pair is biased at $10\mu A$ tail current, half going into each transistor. In order to achieve maximum transconductance for the power consumed, the transistors are sized so as to be in subthreshold. Calculations for the minimum size transistor to be in the subthreshold region for a given drain current can be found in Vittoz and Fellrath's seminal paper on subthreshold CMOS [11]. Alternatively, it can be found by examining a family of I_D vs. V_{DS} curves plotted for multiple values of $\frac{W}{L}$. It was decided that an acceptable value of $\frac{W}{L}$ for the input pair is $\frac{100\mu m}{0.6\mu m}$. The subthreshold transconductance of a MOS transistor is $\frac{I_D}{n}\frac{q}{kT}$, where q is the fundamental charge, k is Boltzmann's constant, T is absolute temperature, and n is the subthreshold slope factor of the transistor, $\frac{\partial V_{GS}}{\partial \Psi_S}$. The subthreshold slope factor varies between transistor processes in a range from approximately 1.2 to 2, and is consistent for a process from lot to lot. For this process, n is approximately 1.4, and so the g_m of each input transistor is approximately $145\mu S$ at room temperature.

The output high impedance nodes are created with cascoded current sources on each power supply. Subthreshold wide-swing cascodes are used, to provide a wide linear output range and avoid having to create additional bias voltages for the cascode transistors. The gates of the current sources and the cascode transistors are connected to the same potential, but the cascodes are sized much larger for their bias currents. The V_{DS} of the current source transistor is held at the difference in the gate-source voltages of the two transistors. It only takes about 100mV V_{DS} in subthreshold, independent of I_D, to keep the transistor in saturation, which is easily created with the transistor scaling. The current source transistors are sized at $\frac{W}{L} = 1.2/\mu A$ of drain current, and the cascodes are sized at $\frac{W}{L} = 7.5/\mu A$.

The current sources on the low rail for the input pair and the output stage are biased with an external voltage generated with a bandgap circuit. This voltage is the base voltage from a current mirror circuit scaled to produce a $1\mu A$ bias current from a $\frac{20}{5}$ NMOS transistor. The current sources biasing the low side of the high impedance output nodes are controlled by the common mode feedback circuitry to be described

in the next section.

The fully differential outputs of the folded cascode are buffered by a low impedance output stage, to be able to drive the $50k\Omega$ feedback load. A class AB push-pull stage is used, to minimize total power consumption. $15\mu A$ is always drawn by the feedback resistors, regardless of the output voltage driven by the op-amp. A class AB stage allows that current to be supplied by either the output stage or the power supply, depending on the output voltage. When the output stage is not supplying the current to the feedback network, it only consumes a small bias current. If a standard class A stage were used, it would have to consume $30\mu A$ all the time, regardless of whether that current was being supplied to the resistors.

The class AB stage is built with a bipolar diamond circuit, consisting of two emitter-followers to provide current gain and level shifting to the output transistors, and two more emitter-followers to deliver the output current. This provides a current gain of $\beta^2 = 10000$, limiting the input current to 1.5 nA, safely below the bias current in the high impedance node. Bipolar transistors are used in this stage for their superior ability to drive current in a push-pull configuration. The low transconductance and output impedance of the MOS transistors makes them much less suitable for applications involving current drive.

The disadvantage of the diamond circuit is its limited output swing. It can only act as a buffer for signals away from either rail by $V_{GS,sat} + V_{BE}$. This is at least 0.8 V, if the MOS transistor is in subthreshold, and could be as high as 1 V. This limits the maximum amplitude of the output signal to only 1 V, which is considerably less than we would like. The solution to this problem is to give the output differential to single-ended converter a gain of 5, reducing the maximum signal at this point in the signal path to 0.6 V, easily within the linear range of the buffer. The result of this change is that the feedback attenuation A in the loop is one fifth of what it was previously, and thus the loop crossover frequency is 5 times higher. This necessitates a feedback capacitor 5 times larger in the integrator to get the desired bandwidth, which is a great cost in die size.

Compensation of the differential signal path is relatively easy, as the amplifier topology consists entirely of low impedance nodes, with the exception of the two high impedance output nodes. These two nodes are used for dominant pole compensation of the amplifier.

The impedance at the output node is one half the output impedance of an MOS cascode, as there are

two cascodes in parallel. $R_o = \frac{1}{2}(r_{o1} + r_{o2} + g_{m1}r_{o1}r_{o2})$, where transistor 1 is the cascode and transistor 2 is the current source. The gain of the amplifier is simply $G_m R_o$, where G_m is the transconductance of the subthreshold input pair, $\frac{I_D}{nV_T}$. The dominant pole of the amplifier occurs at a frequency $\frac{1}{R_o C_C}$, where C_C is the compensation capacitance at the high impedance node. The dominant pole will cause the crossover frequency to be at $\frac{G_M}{C_C}$, unless another pole occurs earlier.

The output impedance of the op-amp is lowered by the switched-capacitor common mode feedback network, which will be discussed in the next section. This network looks like a parallel RC, and so the op-amp's output resistance is lowered and its output capacitance is increased. The capacitance of the common-mode feedback network is $\frac{1}{3}$ pF, and the resistance is 100 MΩ.

Using the device parameters generated by the simulator, $R_o = 70M\Omega$. This is in parallel with the 100 MΩ of the common mode network, giving a total output resistance of 41 MΩ. The transconductance of the input pair is 106 μS, giving a DC gain of $G_m R_o = 4346$, which meets the spec of 1000.

The compensation capacitor chosen was 2 pF. As shown in the schematic, it runs from one output leg to the other. This has the effect of compensating both output legs with a single capacitor and of miller multiplying the capacitance by a factor of two in differential mode. So the total capacitance needed is reduced by a factor of four. The compensation capacitance looks effectively like 4 pF from the Miller capacitor in parallel with $\frac{1}{3}$ pF from the common mode network, for a total of 4.33 pF. The output node pole is thus at 896 Hz, and crossover occurs at 3.9 MHz.

After the output pole, the second most relevant one is the one caused by the capacitance at the drains of the input transistors. It is spaced far from the dominant pole, so its location is approximately equal to the time constant of that capacitance. The resistance seen at that node is $\frac{1}{g_m}$ of the PMOS cascode transistor in parallel with the output resistances of the current source transistor and the input transistor, for a total of $R_D = 24$ kΩ. The total capacitance is $2C_{gd} + C_{gs}$, placing the pole at 4.8 MHz.

A Bode plot of the simulated differential mode gain and phase is shown in Figure 5-5. The low frequency open-loop gain is about 2500, meeting the specification of 1000. The dominant pole is at about 800 Hz, and dominates the frequency response up to the second pole at about 2.5 MHz. This gives a crossover frequency of 2 MHz, meeting the bandwidth specification. The second pole contributes about 35 degrees of phase shift

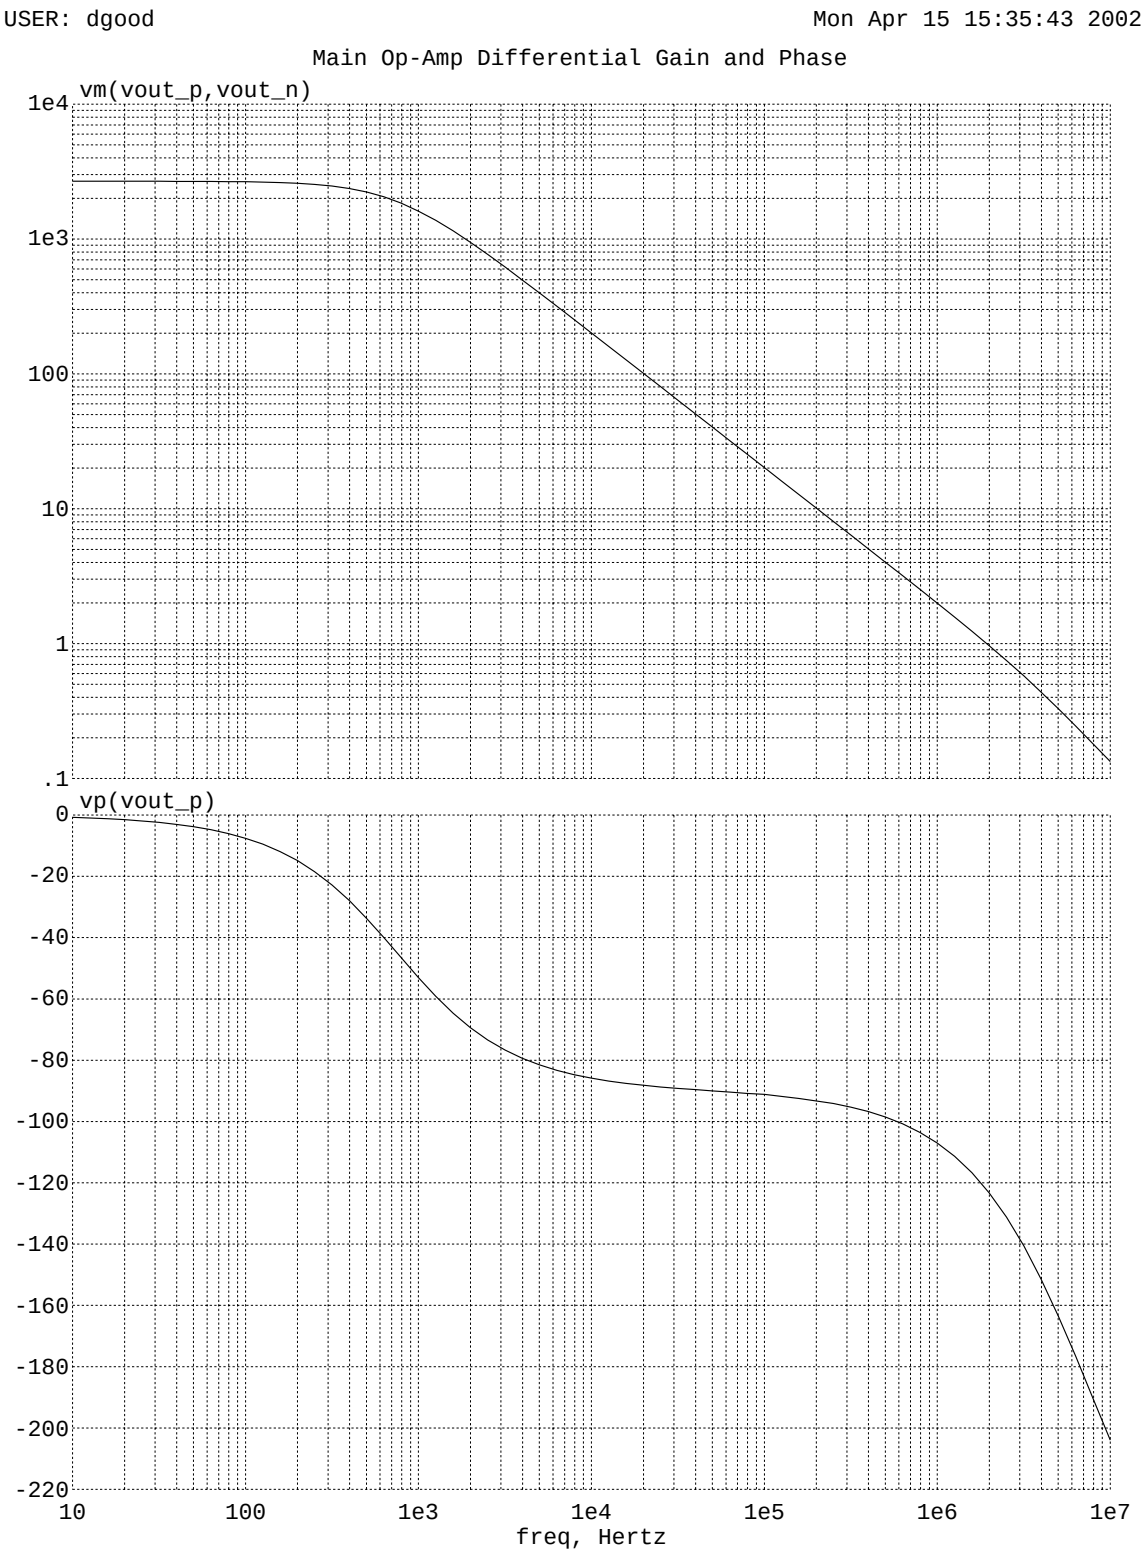

Figure 5-5: Open Loop Bode Plot of Integrator Op-Amp Differential Mode

at crossover, yielding a phase margin of 55 degrees. This is a fairly stable amplifier.

The calculated and simulated results correspond reasonably, though there is some error. The error is most likely due to miscalculation of the output impedance of the amplifier, due to the complexities of the subthreshold wide-swing cascodes and the switched-capacitor feedback network. Additionally, there is some gain loss caused by the Early effect in the input pair, as the drain terminals are modulated. The small-signal gain to those terminals is $G_m R_D = 2.5$, with R_D the same as above. This is small, so it will not hurt the overall gain terribly.

The differential path consumes 10 μA in the input pair, 1 μA in each cascode leg, 2 μA bias current in each diamond circuit, and 1 μA in the biasing, for a total of 17 μA. Any additional current drawn from the supply by the output stage will be to drive the feedback network, and its total current has already been counted, so it should not be counted again here. The total power consumed by the differential path of the integrator is 51 μW.

5.3.3 Common Mode Design

The op-amp is built fully differentially, causing the two outputs to be referenced to one another, rather than to ground. The feedback loop enclosing the amplifier has a large differential mode loop gain, producing a well-defined differential mode output. However, the input to the amplifier intentionally has common mode rejection, to remove common mode offset from input signals. The loop thus has very low common mode gain, and the common mode voltage at the output will not be well controlled. This becomes a problem because the amplifier loses linear output range if the output signals are not biased around the middle of the supply.

Control of the common mode is achieved by wrapping it in a separate feedback loop. A circuit is used that measures the common mode of the output and alters the biasing of the amplifier output until it is biased around the middle of the supply. Because the common mode and differential mode signals are linearly independent, the common mode loop has its own dynamics, separate from those of the differential mode loop, and must be analyzed and compensated separately.

The circuit used to control the common mode loop is shown in Fig. 5-6. It takes as input the two output voltages, the target common mode output voltage, and a reference bias voltage. The circuit creates an

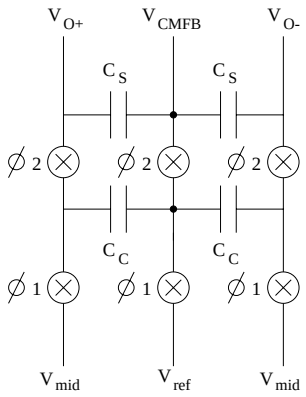

Figure 5-6: Common Mode Feedback Control Circuit

output voltage equal to the reference bias voltage plus the difference between the output common mode and the desired output common mode. If the common mode voltage is too high, the biasing current sources are driven harder, and they pull current out of the output node until the voltage moves down to where it should be. Likewise, if the voltage is too low, the biasing current sources are shut off, and the current sources on the top are allowed to push the output node up to where it should be.

The two capacitors C_S in the circuit are use to hold the output voltage during phase one, and are not switched. The capacitors C_C are switched, and are used to compute $V_{ref} + \frac{V_{O+}+V_{O-}}{2} - V_{mid}$ in steady state. Steady state occurs when the switches change from phase one to phase two without the output voltage moving. This means that no charge moves through the closed switches, and thus the sum of the charge on the capacitors C_C stays constant.

In phase one:

$$Q_{CC} = 2C_C(V_{ref} - V_{mid}). \tag{5.3}$$

In phase two, letting $V_{O+} = V_C + \frac{V_D}{2}$ and $V_{O-} = V_C - \frac{V_D}{2}$,

$$Q_{CC} = C_C(V_{cmfb} - V_C + \frac{V_D}{2}) + C_C(V_{cmfb} - V_C - \frac{V_D}{2}). \tag{5.4}$$

Therefore, setting the capacitor charge equal between the two phases,

$$V_{cmfb} = V_{ref} + V_C - V_{mid} = V_{ref} + \frac{V_{O+} + V_{O-}}{2} - V_{mid}. \tag{5.5}$$

Qualitatively, the feedback loop should push the output common mode to be equal to the V_{mid} reference voltage. Besides verifying the functionality of the feedback loop, it is necessary to make sure it is stable, and has appropriate bandwidth. Since the common mode feedback loop contains a switched capacitor filter, the loop should crossover well before the clock frequency, to ensure that there is very little gain where the filter stops behaving according to the continuous-time approximations. The filter has a resistance of 100 MΩ and a capacitance of $\frac{1}{3}$ pF, placing the pole at 4.77 kHz.

The closed loop differential mode Bode plot is shown in Figure 5-7. The closed loop pole appears at about 4 kHz, as predicted. The single pole response goes out until the frequency is no longer negligible compared to the clock frequency. At that point the equations used to describe the circuit earlier are no longer valid. The circuit was designed with the pole at low enough frequency that the gain near the clock frequency is well below -20 dB, and so the odd behavior at those frequencies is not a concern.

The common-mode circuit consists only of switched capacitor circuits running at low frequency. The power consumption of a switched capacitance circuit is $\frac{1}{2}CV^2f$. The small capacitances and low frequency clock make the power consumption of this circuit and all the switched capacitor circuits negligible compared with the DC currents, so they will be ignored.

5.4 Biasing

Two major circuits are necessary for the biasing of the accelerometer. The first is a temperature stable current reference, so that the properties of the circuit will be relatively stable across temperature. The reference should be stable enough so that the power consumption meets its specification across the entire temperature range, as power consumption is one of the circuit properties most variable with temperature.

The other necessary circuit block is a V_{mid} reference generator. This chip runs off of a single supply voltage, but needs to have an output voltage capable of swinging in both directions from its bias. A voltage equal to half of the supply is generated and used throughout the chip as “ground”, allowing the circuit to

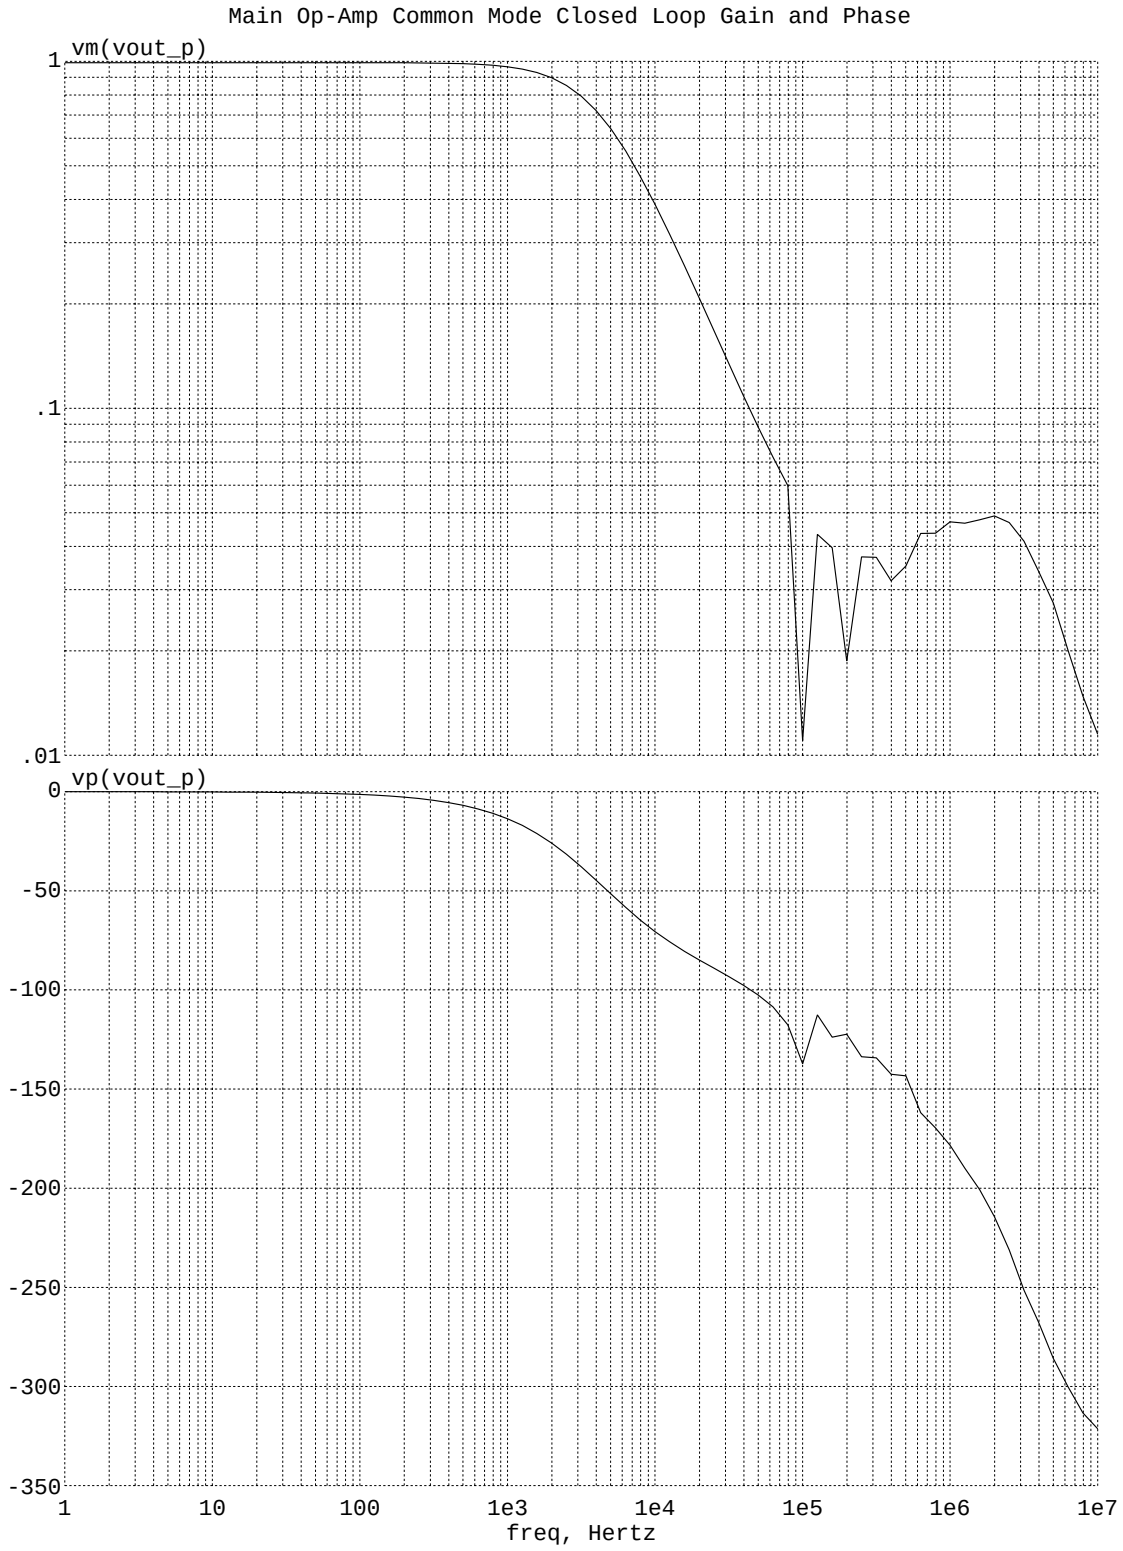

Figure 5-7: Closed Loop Bode Plot of Integrator Op-Amp Common Mode

be designed as though there are bipolar supplies present.

5.4.1 Bandgap Reference

The signal processing circuitry has many current sources which need to be biased. It is important that these bias currents remain constant with temperature. Since $\frac{\partial V_{GS}}{\partial T}$ is negative, if the bias voltage is constant, the current will increase with temperature. Similarly, if the voltage were generated with a current mirror biased with a resistor to the power supply, the current would increase with temperature. The current drawn in the biasing current sources controls the power consumption of the circuit, so the power consumption would increase with temperature for a simplistic biasing method like this. As this is primarily a low power device, it is unreasonable to let this happen. A more sophisticated biasing scheme needs to be used.

A bandgap reference is used to generate a temperature stable current reference. The relation between the temperature coefficient of V_{BE} of a bipolar transistor and the bandgap voltage of silicon is exploited to create a voltage which has no temperature coefficient at room temperature. This is done by taking the difference of the V_{BE}s of two transistors running at different current density, scaling it, and adding it to a V_{BE}, either of another transistor or one used in the differencing. This voltage is:

$$V_{out} = A(V_{BE1} - V_{BE2}) + V_{BE1}. \tag{5.6}$$

The simplified temperature coefficient of V_{BE} is:

$$\frac{\partial V_{BE}}{\partial T} = \frac{V_{BE} - V_{G0}}{T}, \tag{5.7}$$

where V_{G0} is the bandgap voltage of silicon. Thus, the temperature coefficient of V_{out} is:

$$\frac{\partial V_{out}}{\partial T} = A\frac{V_{BE1} - V_{BE2}}{T} + \frac{V_{BE1} - V_{G0}}{T} = \frac{V_{out} - V_{G0}}{T}. \tag{5.8}$$

This temperature coefficient becomes zero if $V_{out} = V_{G0}$. V_{G0} is itself a function of temperature, so the temperature coefficient is not zero everywhere, but it can be made very small over a wide range of

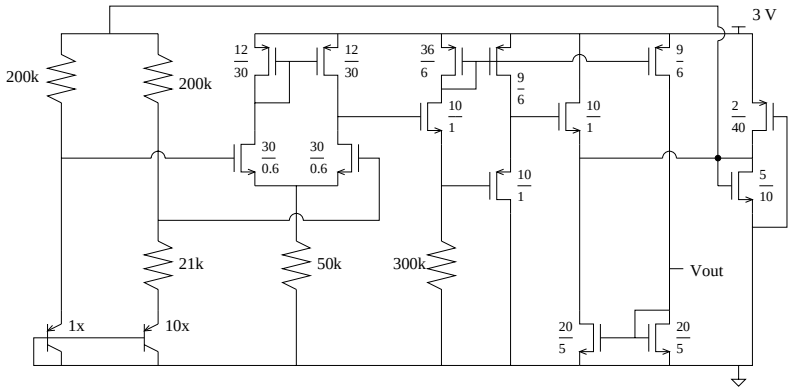

Figure 5-8: Bandgap Current Reference Circuit (all dimensions in μm)

temperature.

The circuit used to build this reference voltage and convert it into a temperature stable current is shown in Fig. 5-8. It uses a simple single stage op-amp and two bipolar transistors with area scaled by a factor of ten, run at the same collector current. The feedback of the op-amp guarantees that the its two input terminals will be at the same voltage, and so there will be the same current running into the two transistors. Therefore:

$$I = \frac{V_{out} - V_{BE1}}{R2} = \frac{V_{BE1} - V_{BE2}}{R1}, \tag{5.9}$$

$$V_{out} = \frac{R2}{R1}(V_{BE1} - V_{BE2}) + V_{BE1}. \tag{5.10}$$

This satisfies Equation 5.6. All that is left is to scale the resistors to set $V_{out} = V_{G0} \approx 1.2$ V, and to convert the voltage to a current.

It is unfortunately very difficult to convert the voltage to a current without introducing any additional temperature coefficient. The node V_{out} is fed back to draws a lot of current, so the output cannot just be put onto a resistor and the current drawn measured, because the current draw would be the sum of the resistor current and the fed back current. Instead, the output is buffered and fed back, and the output before the buffer is put across the resistor. The buffer draws no current, so the resistor current can be accurately

measured, and mirrored to the output.

The buffer used here simply consists of an NMOS source follower and a PMOS source follower, sized the same and run at the same drain current. This will nominally have a small offset, and its change in offset with temperature will be equal to the difference of the two V_{GS} temperature coefficients. These will not exactly cancel, but the difference should be much smaller than the nominal V_{GS} temperature coefficient. Additionally, the extra temperature coefficient can be cancelled by sizing R1 and R2 appropriately to give the bandgap circuit the opposite temperature coefficient at room temperature. I found the temperature coefficient with the simulator, using the temperature characterization data for the process, and adjusted the resistors until it was correct. Due to the complexities of MOS temperature coefficients, solving this accurately by hand is very difficult.

The resistor current is mirrored out to create an output voltage. The voltage is used to bias the gates of the current source transistors around the circuit, and is scaled such that it should generate a 1 μA current in a 20/5 μm transistor. Process differences in the resistors and transistors will cause this current to vary somewhat, so a trim was included. The temperature coefficient of the output current was fairly stable across process, but the absolute value of the current wasn't, so the trim was put in the voltage to current conversion resistor. Using laser blown fuses, the resistance can be varied from 230 kΩ to 390 kΩ, a wide enough range to account for all process corners.

A final important part of the bandgap circuit is a startup circuit. The circuit uses feedback to compute the bandgap voltage, which describe the output voltage as the set of simultaneous equations in Equation 5.9. Unfortunately, the equations also have a degenerate solution, with all currents and voltages equal to zero. When the circuit is turned on, it is likely that the circuit will simply sit at the zero state unless something is done to push it away. This is accomplished with a PMOS transistor, connected so that $V_{SG} = V_{DD}$. Whenever power is connected, this transistor will want to pull current. It is scaled very long and narrow, so that its current level will be fairly low. Its current is driven into an NMOS diode, whose voltage rises with the current. The voltage at the top of the diode is tied to the output voltage of the op-amp, so that it will be pushed away from zero and find the other solution to the feedback equation. The diode is scaled so that its current level will be low when $V_{GS} = V_{G0}$, and any additional current sourced by the startup transistor

will be delivered to the feedback resistors.

The simulated bandgap voltage and output current versus temperature for nominal process and fast and slow process corners before and after trim are shown in Figures 5-9 and 5-10. The process variation corners are provided by the ADI Limerick fabrication plant. Before trim, the output current varies less than 10% across a wide temperature range, but by about ±25% across process. There are eight steps available in the trim resistance to compensate for the variation of 50% around the nominal output current. This corresponds to a point every 6.25% in the process variations where it is cancelled exactly by the trim. The maximum bias current error is thus ±3.125% off of the nominal 1 μA. Figure 5-10 shows the output bias current trimmed to within 1% of the nominal value for the worst case process errors in each direction.

The bandgap circuit consumes 2.5 μA in each resistor leg, 0.5 μA in the differential pair, 4 μA to perform the voltage to current conversion with the 300 KΩ resistor, 1 μA in each of the three follower and biasing legs, and 3 μA in the startup circuit, for a total of 15.5 μA, or 46.5 μW.

5.4.2 V_{mid} Reference

The schematic for the V_{mid} reference generator is shown in Fig. 5-11.

The reference voltage equal $V_{DD}/2$ is created with four PMOS diodes running in series from V_{DD} to ground. Each one is configured with $V_{GD} = 0$ and $V_{BS} = 0$, such that they will all have identical V_{DS} when run at the same bias current. This produces a voltage at the middle node of the string equal to exactly half the supply. On the specified 3V supply, each diode has $V_{GS} = 0.75V$, which is well below the threshold voltage. Thus the transistors will conduct only nanoamps of bias current, and their power consumption will be entirely negligible. There is some worry that on a higher supply voltage, this circuit will conduct too much current and be destroyed, due to the direct control of the transistor's V_{GS}. On a 5V supply, which the circuit is not specified for, this will begin to conduct microamps of bias current, and its power consumption will begin to be significant compared with the power consumption of the rest of the chip. It takes more than 5V to destroy the devices in this configuration, and so this ought to be an entirely safe circuit on the specified power supply.

Diodes were used here instead of resistors, which would be the natural choice, to minimize power con-

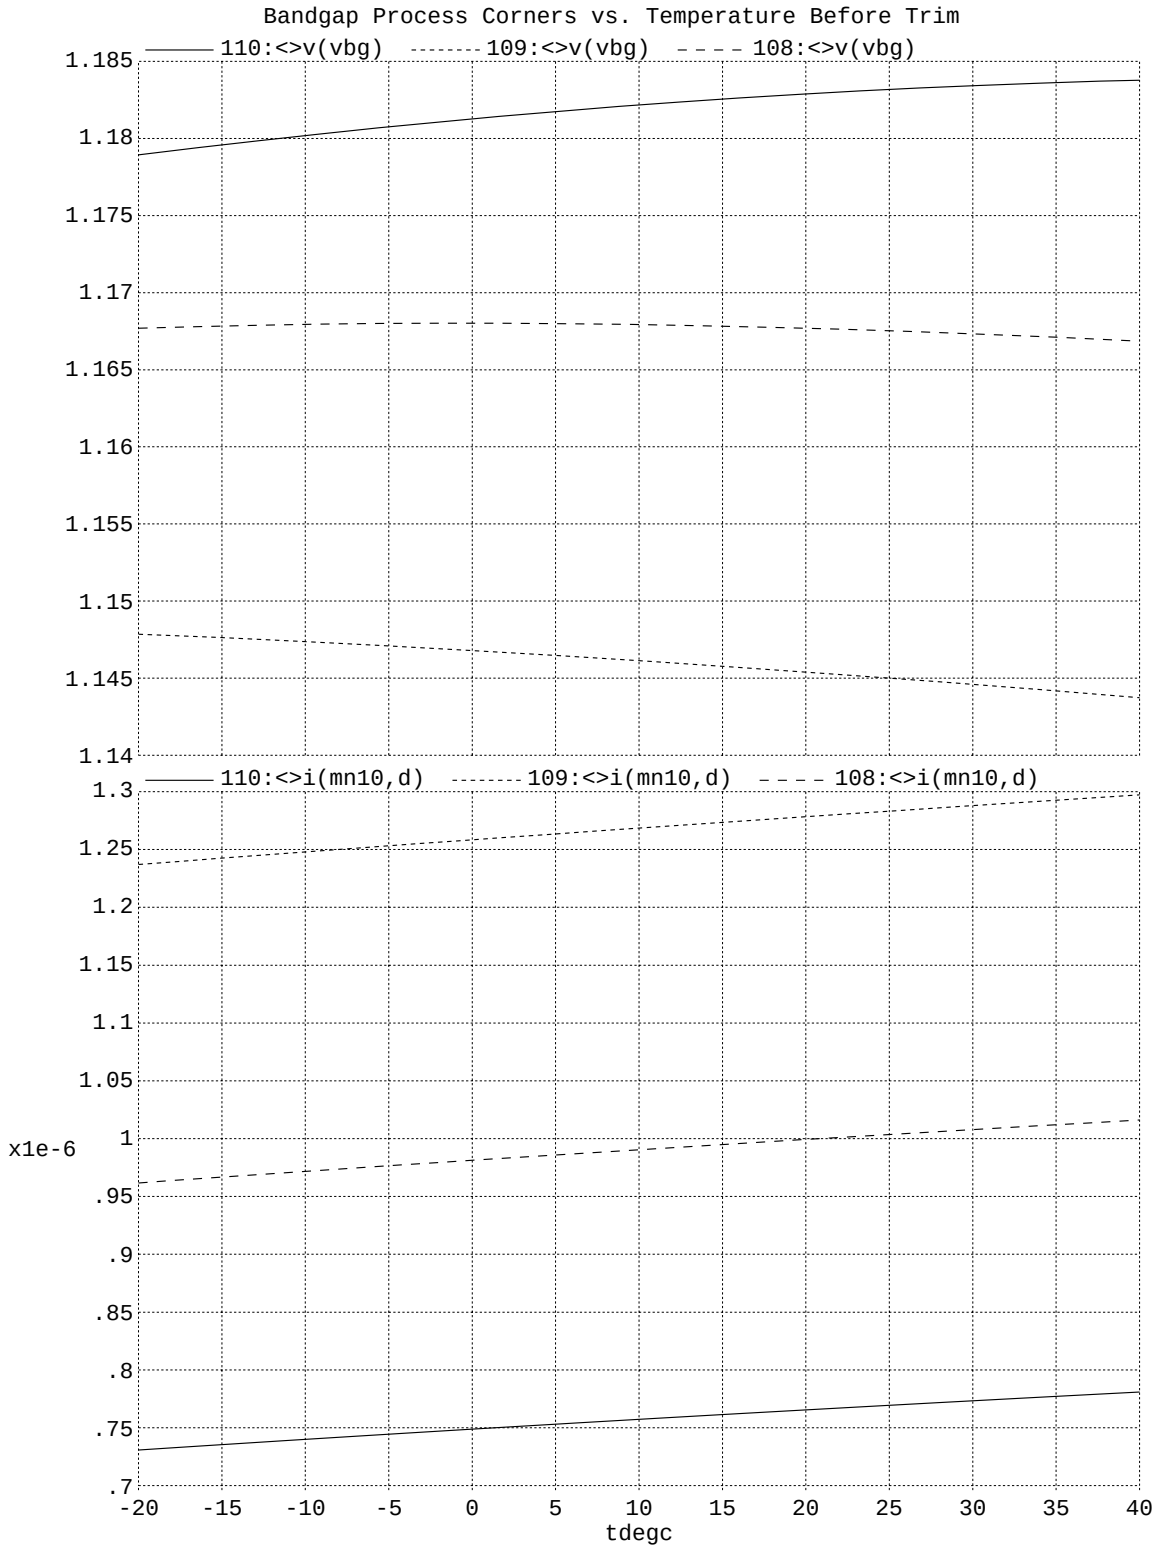

Figure 5-9: Bandgap Voltage and Output Current vs. Temperature, Before Trim

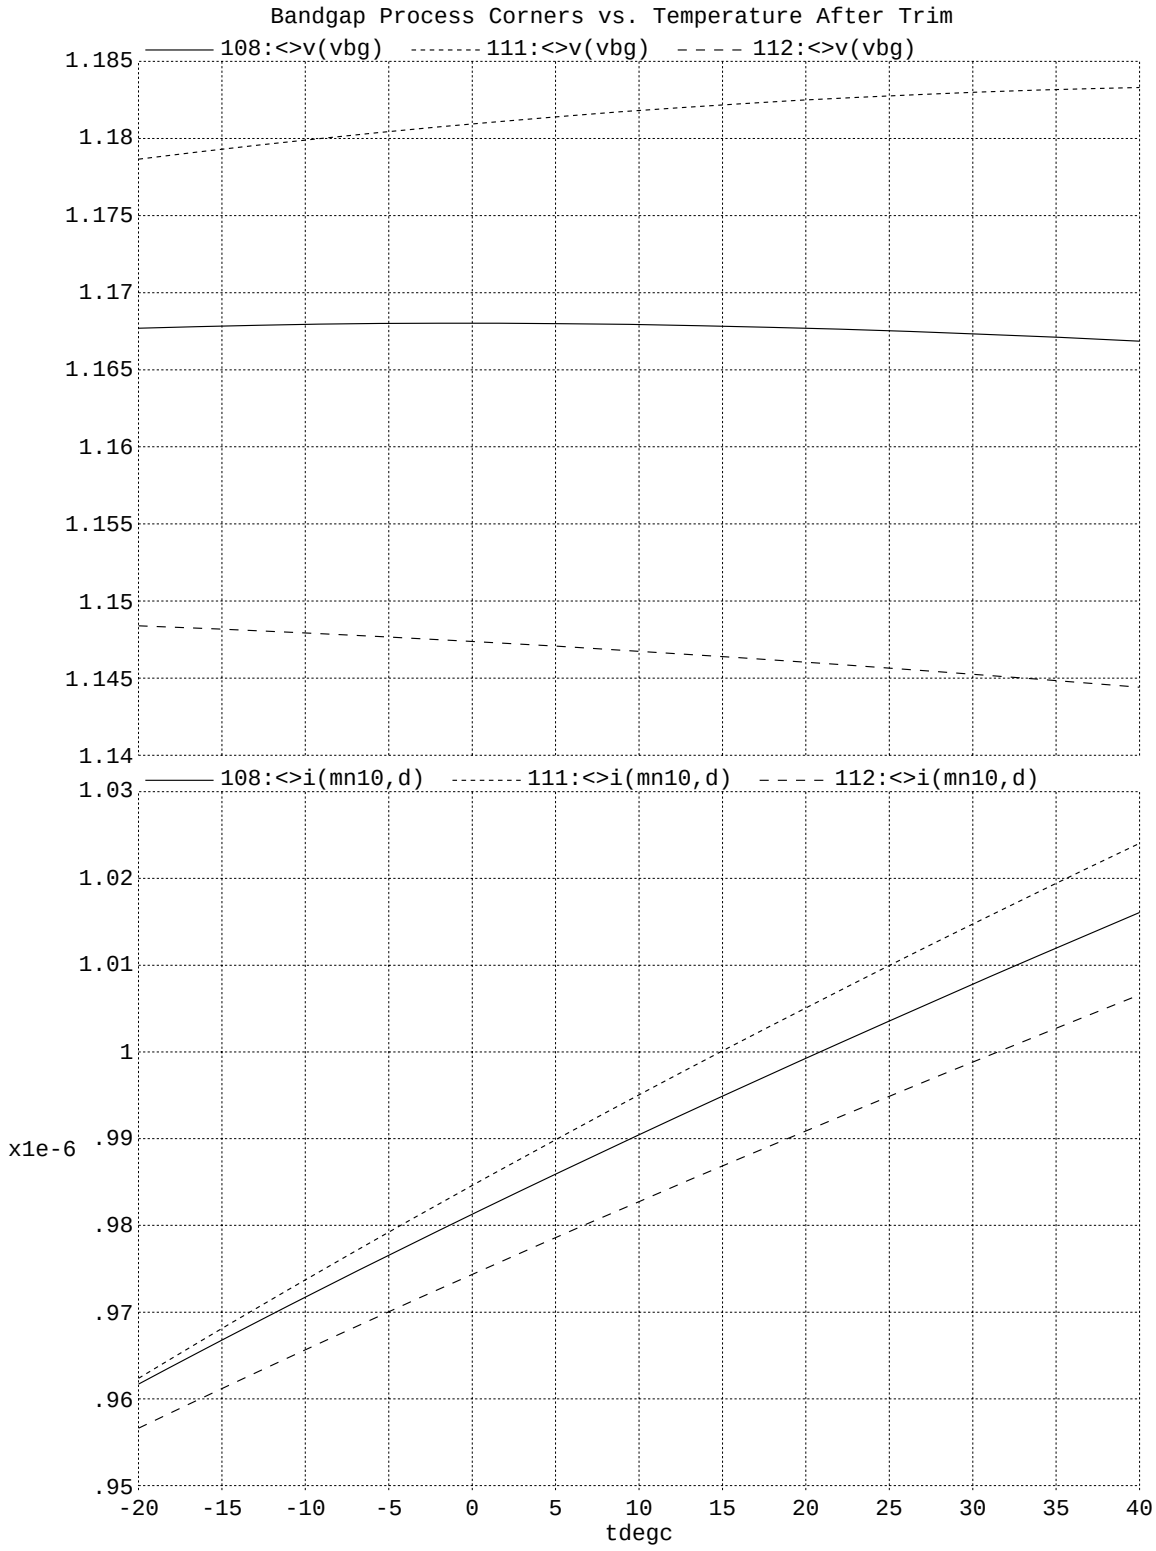

Figure 5-10: Bandgap Voltage and Output Current vs. Temperature, After Trim

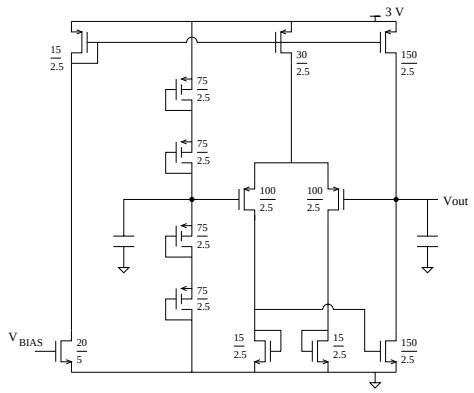

Figure 5-11: V_{mid} Buffer Circuit (all dimensions in μm)

sumption without using an unwieldy amount of die area. The diodes use less power and less die area, at the expense of specifying that the chip must be run on 3 V and not on 5 V or higher. For a product to be manufactured and sold, this would probably be an unreasonable requirement, and would likely end up with uninformed users destroying the parts on a high supply voltage. A different choice for that biasing would have to be made. A good one would likely be to use switched capacitor resistors, which would fabricate a very large impedance in a reasonable amount of space. This would cause a number of additional concerns, such as interdependence of the various biasing cells, and filtering the clock ripple from the ground supply. It is a good option, though.

Once the $V_{DD}/2$ reference is generated, it needs to be buffered, to isolate the reference generator from the circuits being driven. A simple two stage op-amp is used, configured as a unity gain follower. It consists of a diode loaded differential pair, followed by a current source loaded common source amplifier for high gain. The output of the circuit is the drains of the common source transistor and the current source, which looks like a high impedance. This could also conceivably be a problem, as the voltage would vary if it had to drive any current into a low impedance node. Fortunately, V_{mid} is only used for references into high impedance inputs to op-amps and switched capacitor circuits around the chip.

The advantage to building the buffer with a high impedance output is that the pole created by the output of the op-amp with the capacitance it has to drive is the dominant pole of the op-amp. If there were a low impedance output stage, this pole would be in addition to the internal dominant pole, and if the load

capacitance were great enough, the second pole could cause a significant loss in phase margin. As this buffer creates the ground reference for the circuit, it is important that it is very stable. By ensuring that the load capacitance affects the dominant pole, any changes in the load will affect the bandwidth of the amplifier but can only improve the stability.

Crossover of the amplifier's loop transmission, and thus the pole of its closed loop response, occurs at $\frac{G_m}{2\pi C_C}$. The DC gain of the amplifier is $G_m R_o$. The G_m of the input pair is 40 μS, r_o of the NMOS output transistor is 5 MΩ, and the G_m of the PMOS output transistor is 4 MΩ, giving a gain of 88. The total output capacitance is about 4 pF, corresponding to a pole at 18 kHz and crossover at 1.57 MHz.

Open loop Bode plots of the op-amp are shown in Figure 5-12. The open loop gain is 80, corresponding well with calculations. With no load capacitance, the dominant pole occurs at about 25 kHz, putting crossover at about 2 MHz. The second pole is slightly after 2 MHz, contributing about 30 degrees of phase by crossover, and yielding a phase margin of about 60 degrees. Any load capacitance added will move the dominant pole to a lower frequency. Since it is already contributing its full 90 degrees of phase, it will move crossover to a lower frequency without affecting the phase. As crossover moves lower, the second pole contributes less phase, and the phase margin actually improves.

The calculations for the single pole are reasonably accurate. There is some error in the bandwidth calculation due to the strange behavior of the nwell capacitor. This is the gate capacitance of an NMOS transistor, used to build a high density capacitor. Its value is not particularly well controlled, and varies strongly with the bias voltage. This inaccuracy is fine, however, as the capacitance is simply needed to slow the buffer down enough to be stable with no load capacitance. Having an accurate time constant at that node is not important.

Figure 5-13 shows the V_{mid} buffer output voltage across process and temperature. Across all conditions it varies by at most 5.2 mV out of 1.5 V, or slightly more than $\frac{1}{3}$ of 1%. This guarantees that the reference voltage will be consistent.

The buffer circuit consumes a few nanoamps in the voltage divider, 1 μA in the biasing, 2 μA in the differential pair, and 10 μA in the common source output stage, for a total of 13 μA, or 39 μW of power.

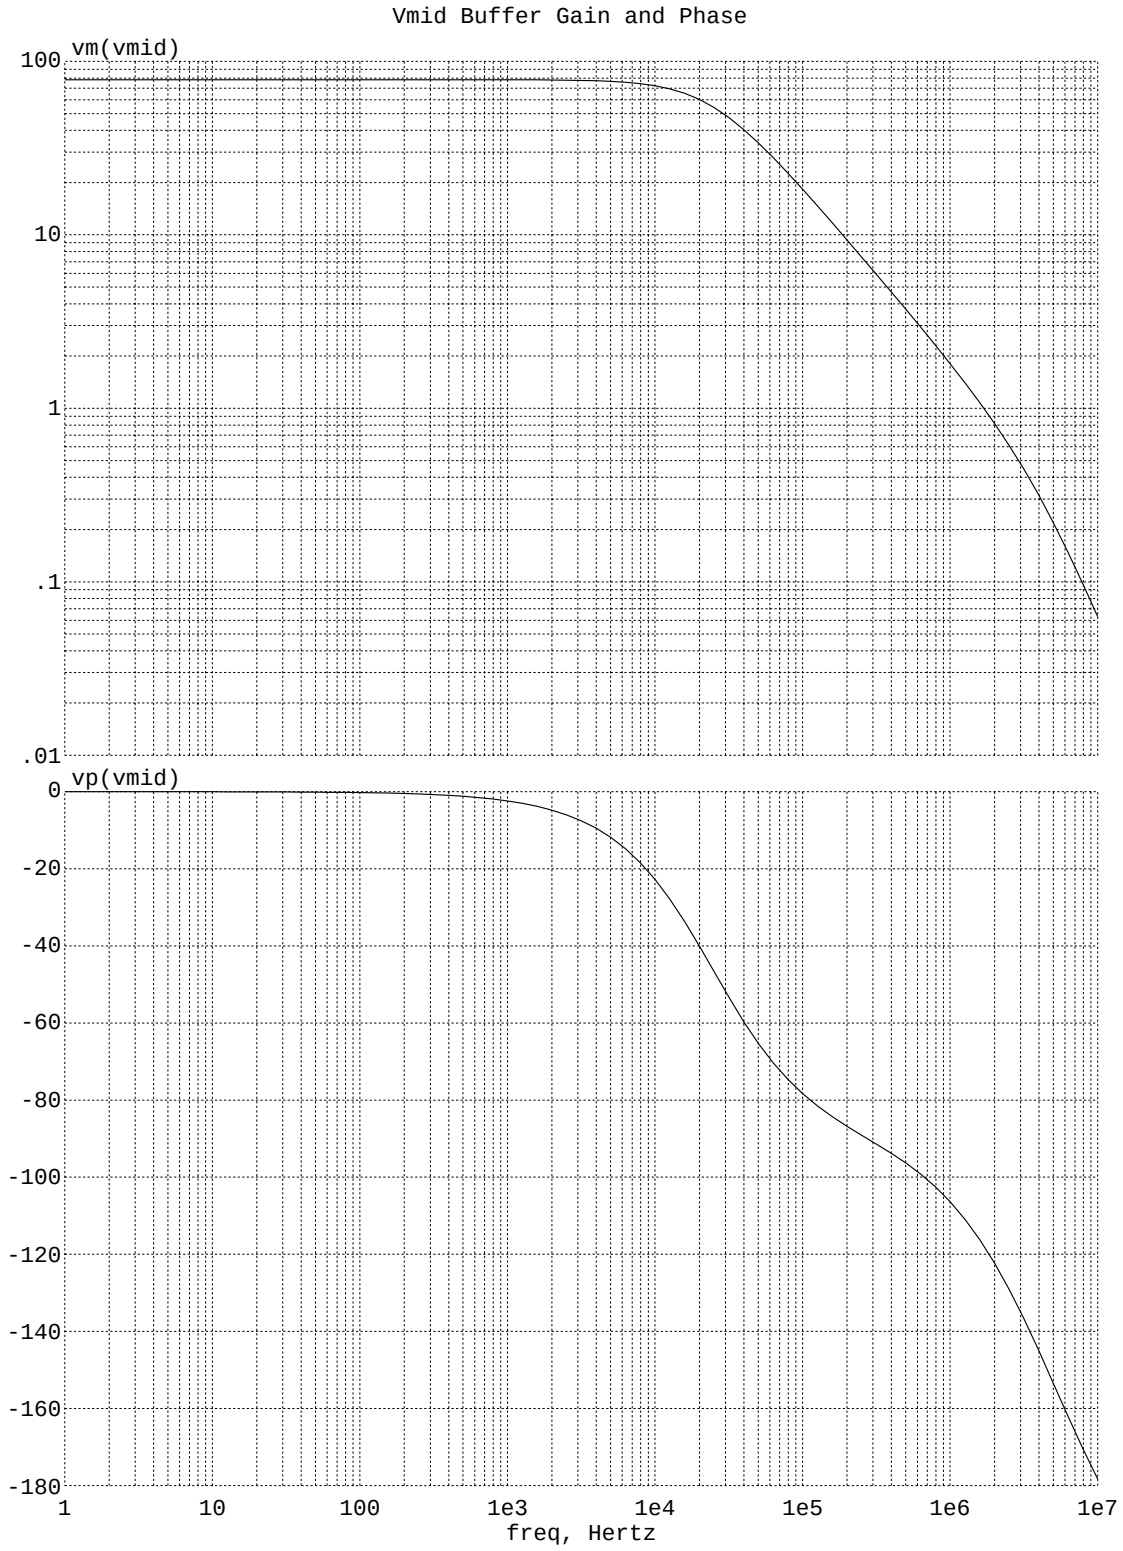

Figure 5-12: Open Loop Plots of V_{mid} Buffer Gain and Phase

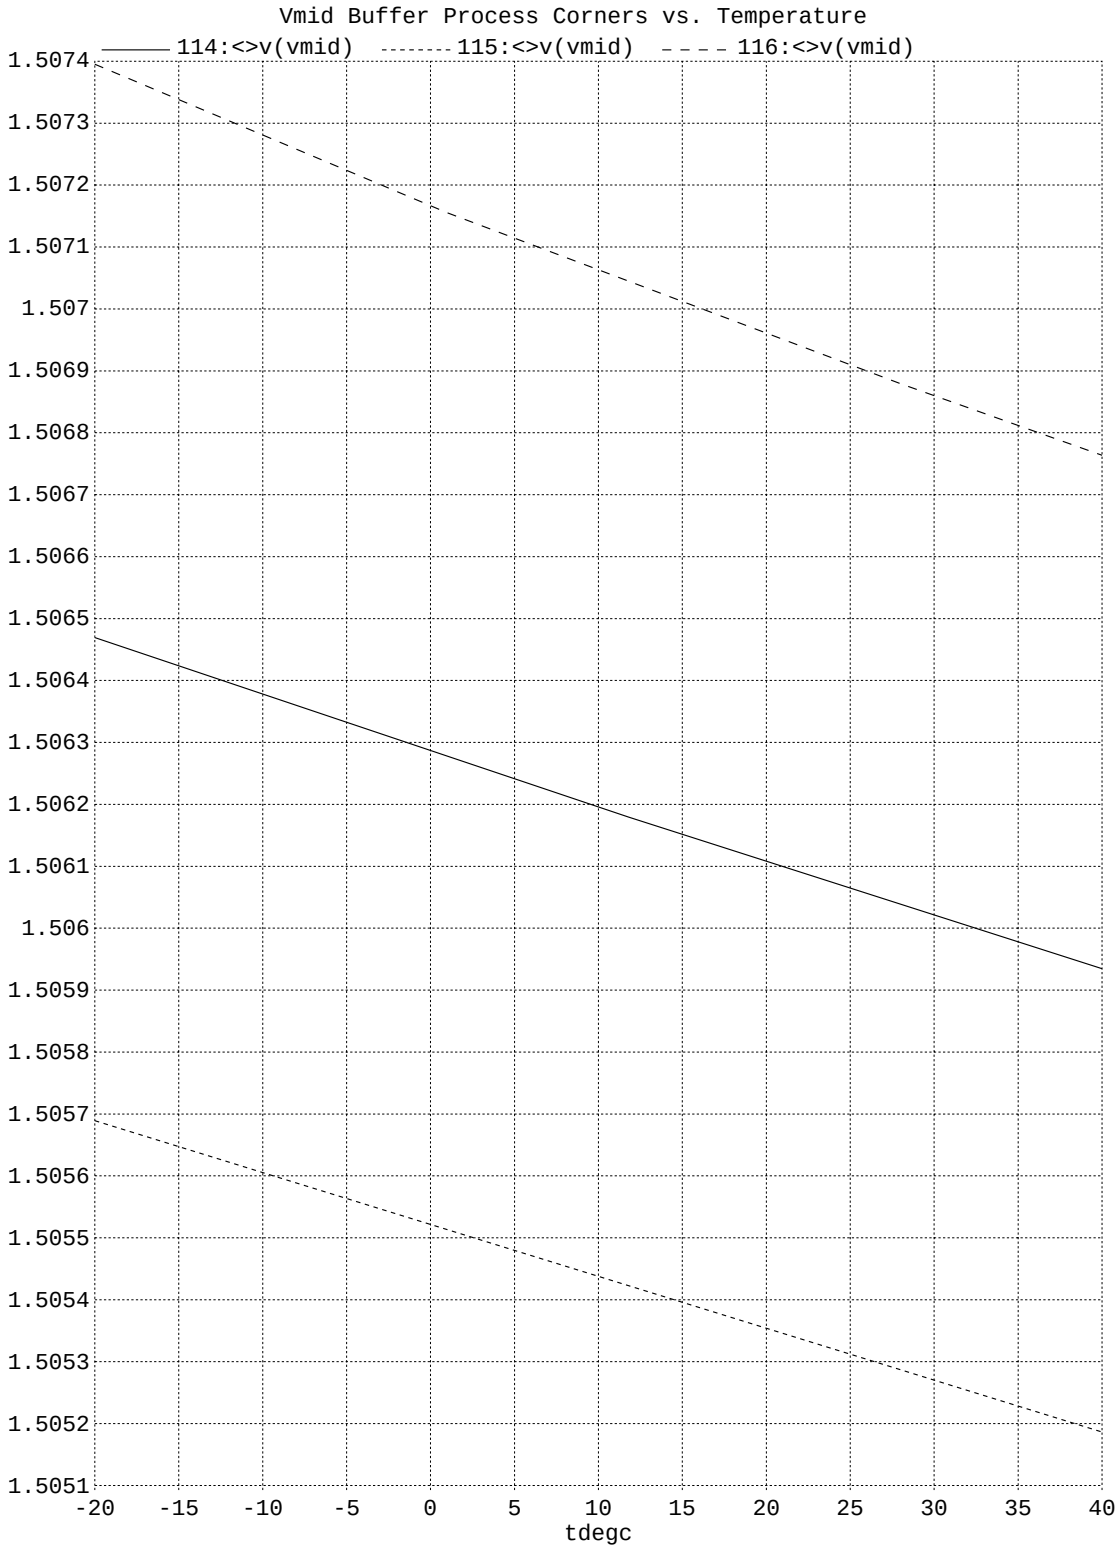

Figure 5-13: V_{mid} Output Voltage Across Process and Temperature

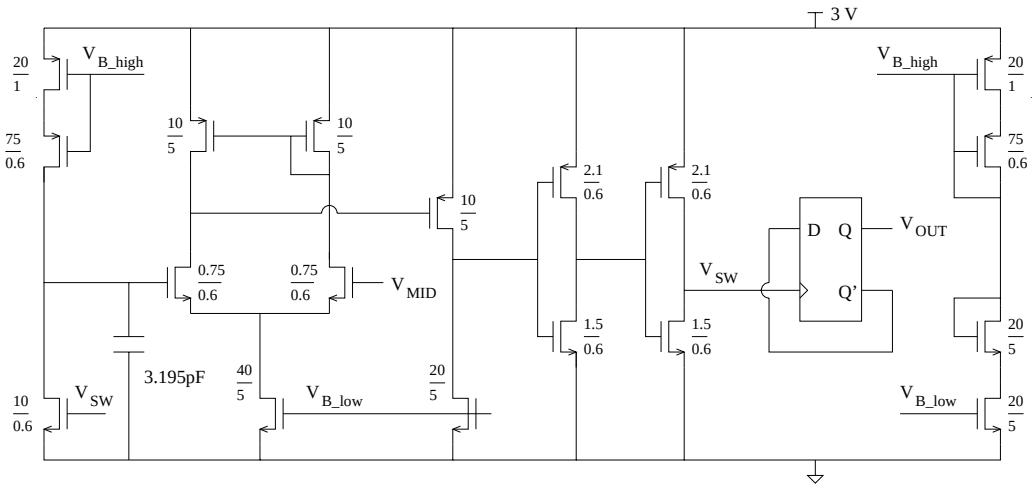

Figure 5-14: 100 kHz Clock Generation Circuit (all dimensions in μm)

5.5 Clock Generation

All of the switched capacitor signal processing elements rely on accurate clocks to control the switches. These clock signals are generated on chip and distributed to the various switched capacitor blocks. While having a clock frequency that is very accurate is not crucial to the operation of the circuit, it should be reasonably stable across process and temperature. However, the relative phasing of the clocks is very important to the function of the circuit, and care will have to be taken to generate the clocks in such a way that guarantees appropriate phasing.

The schematic of the clock generation circuit is shown in Figure 5-14. A voltage ramp is created by charging a capacitor with a current source. This current source is referenced to the bandgap current reference, so it will be reasonably stable with temperature. When the capacitor voltage exceeds the reference voltage, the comparator is tripped and the discharge transistor is turned on, pulling the capacitor voltage back down to ground, creating a triangle wave.

A 1 μA charging current is used with a 3.195 pF capacitor, resetting at 1.5 V. This gives a clock frequency of 208 kHz. This is slowed down by the propagation delay of the reset loop, and divided by two at the output, yielding a clock frequency of approximately 100 kHz. This is not a particularly accurate way to generate the clock pulse, but it is not necessary for this application to have a very accurate clock. The clock will be near

100 kHz, and that is what matters.

A simple CMOS buffer is used to clean up the comparator output into a reasonable waveform and provide some delay around the loop. If the reset pulse were provided simply by the comparator output, there is a chance that as the reset transistor is pulling the capacitor voltage down, the comparator will trip back and cause the capacitor to begin charging again from an intermediate point in the sawtooth. The propagation delay of the buffer is enough to ensure that the capacitor is completely reset before the discharge transistor is turned off.

A simple two stage differential amplifier is used as a comparator. Speed is a more important parameter than gain in the input stage, so the input transistors are minimum size to keep down input capacitance. The common source stage provides enough gain to make the amplifier function well as a comparator, and the output is taken from its drain. Since the comparator should never be in a continuous time closed loop, and speed is a primary concern, it is not compensated. When the amplifier is laid out, care must be taken to avoid creation of parasitic feedback paths closing a loop and causing the amplifier to oscillate.

The switch pulse is connected to the clock input of a DQ latch, with $\bar{Q}$ fed back to D, dividing the frequency down by a factor of 2. This reduces the capacitance needed for a given clock frequency by a factor of 2. If space were of a greater concern, additional divisions could be performed. The Q output of the latch becomes the system clock, driving the clock phasing circuitry.

The clock generator consumes 1 μA biasing, 1 μA capacitor charging current, 2 μA in the differential pair, and 1 μA in the common source gain stage, for a total of 5 μA, or 15 μW.

5.6 Clock Phasing

Several different phases of clocks are needed to correctly control the switched capacitor circuits. Most importantly, the main clock and its inverse need to be guaranteed to be nonoverlapping. Several of the switched capacitor circuits have two different drive potentials separated by switches of opposite phase in series. If the two clocks ever overlap, the charge balance equations used to design the switched capacitor circuits will be destroyed, and the output will be invalid.

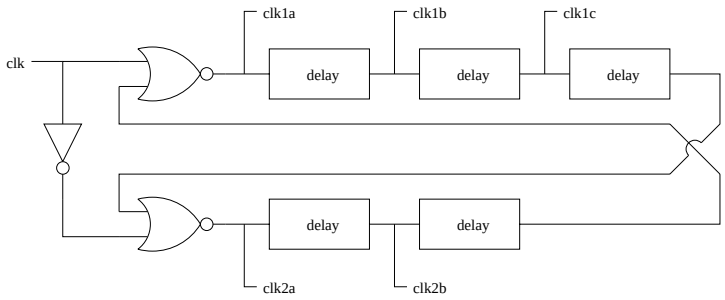

Figure 5-15: Nonoverlapping Clock Phasing Circuit

Secondly, several different fall times of each phase are utilized to minimize charge injection effects in the switched capacitor circuits. These must also be created.

A standard nonoverlapping clock generator is used to create the two main clock phases, and extra delays are inserted to create the different fall times. It is shown in Figure 5-15. When the input is high, clk1a will be low and clk2a will be high, and the loop will be stable. When the input falls, clk2a will fall immediately, but clk1a will remain low. Once the low signal has propagated through the delay cells, it will cause the top NOR gate to switch and clk1a to rise. Identical logic applies to when clk1a falls and clk2a rises. The two clock signals can thus never both be high at the same time.

The multiple phasings are created with extra delay cells in series in the loop. When clk1a falls, its edge will propagate through the delay and cause clk1b to fall, and then clk1c to fall. It is not guaranteed that clk1a is fully settled before clk1b falls, but these clock phasings are not necessary to prevent circuit failure, simply to improve performance. Any delay time between the two is a benefit.

The delay cells are implemented simply as four CMOS inverters in series. They are needed to generate a delay which is large compared to the rise or fall time of a single gate, so this is sufficient.

A plot of the clock phasing when clk1 transitions from high to low is shown in Figure 5-16. The nonoverlapping characteristic of clk1a and clk2a can be seen, as well as the various delays.

Since the clock phasing circuit consists entirely of digital logic, its power consumption is negligible compared with the rest of the circuit.

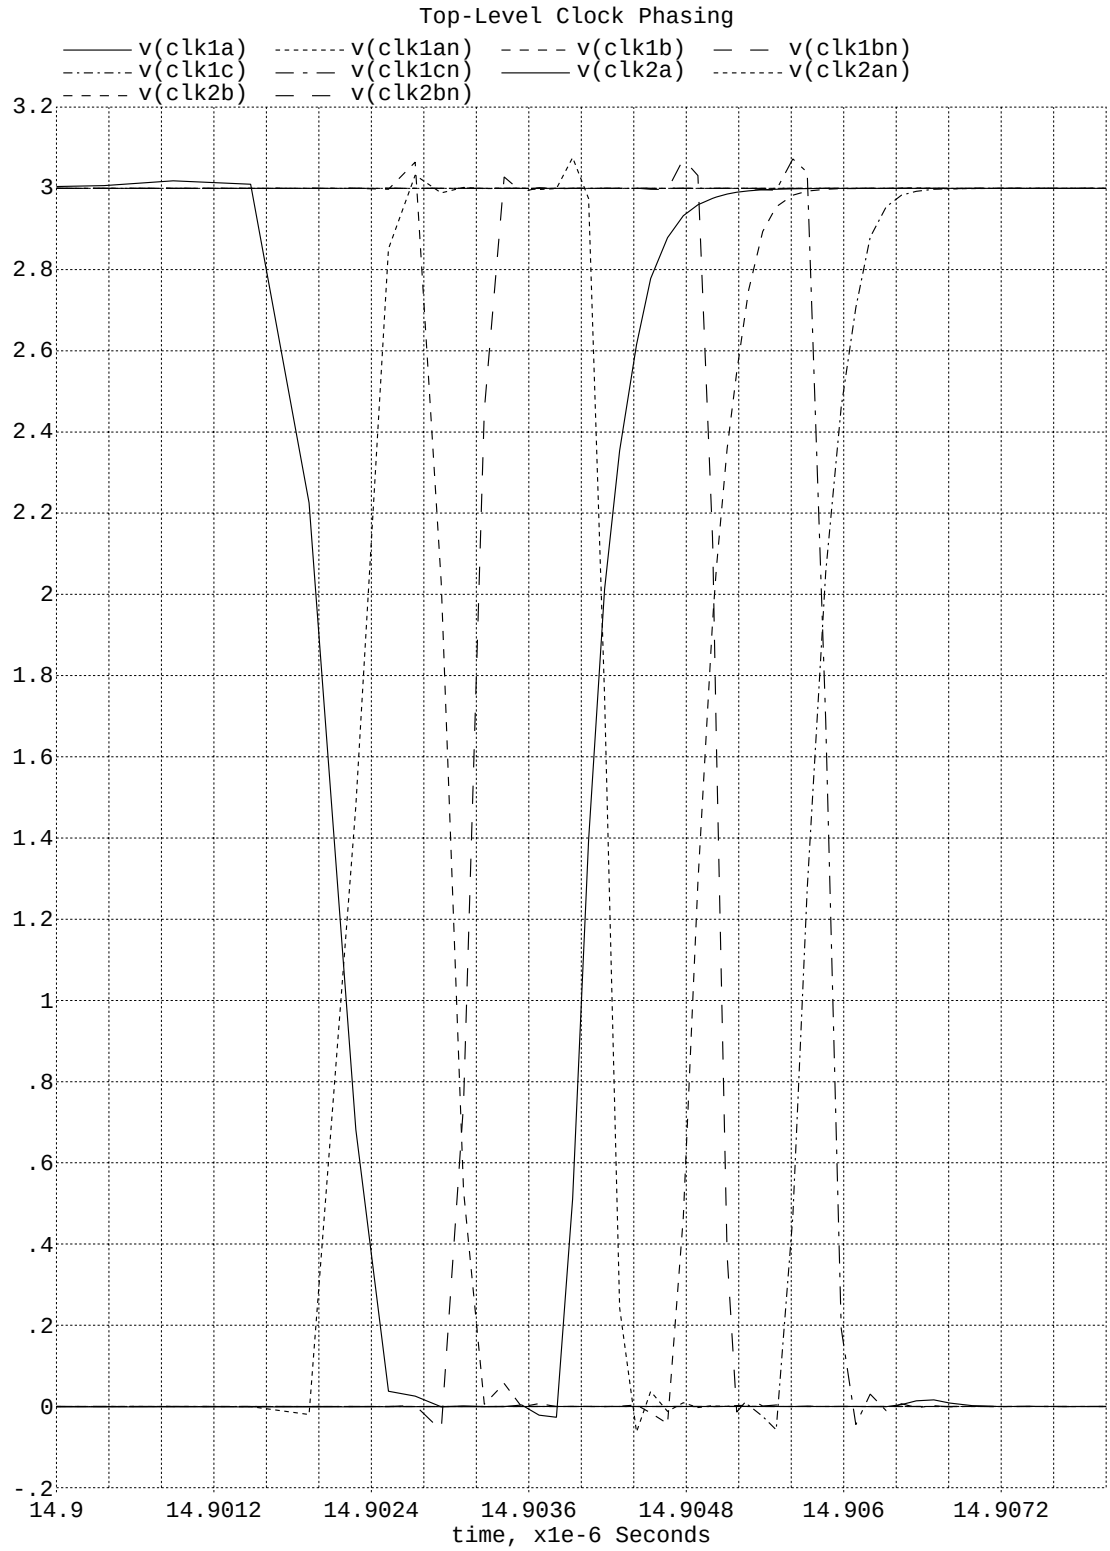

Figure 5-16: Clock Phasing on clk1 High to Low Transition

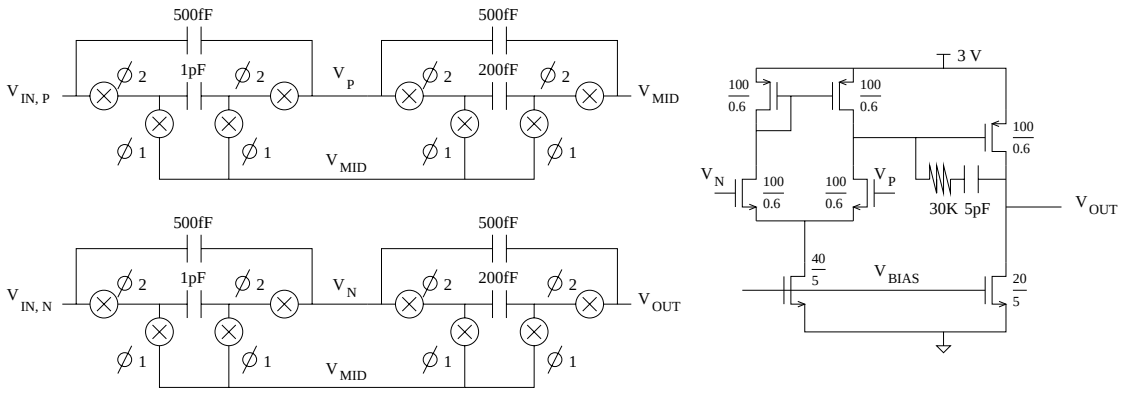

Figure 5-17: Output Buffer Schematic (all dimensions in μm)

5.7 Output Buffer

The final step in the signal processing is to convert the differential signal to a single-ended one, and give it gain to bring it up to the appropriate sensitivity. This is done with a simple op-amp subtracter, configured for a gain of 5. The resistors used to give the attenuations were built with switched capacitors, so that the op-amp would not have to drive a low impedance load.

The external load impedance that the chip can drive is specified at 1 MΩ, so that the power driven into the load does not dominate the power dissipated in the chip. It might seem more practical to be able to drive a lower impedance load, but any application which requires a low power device cannot afford to burn power in this way, and will want to provide the accelerometer with a very high impedance load.

The op-amp used is a high gain two stage amplifier. All of the transistors are running in subthreshold at 1 μA drain current for a $\frac{100\mu m}{0.6\mu m}$ device. As was done in the integrator op-amp, this is to maximize the transconductance to power consumption ratio. A common source amplifier is used to provide a high gain second stage. The current mirror transistors are sized the same as the common source transistor to provide correct biasing. Since they are very short, they have fairly low output impedance, which hurts the gain of the amplifier. However, the alternative would be to make the common source amplifier longer, which would hurt the gain more.

Miller compensation is used to give the amplifier good phase margin with a reasonable size compensation capacitor. A resistor is inserted in series with the miller compensation capacitor to compensate for the

right half plane zero in the common source amplifier. The resistor moves the zero to higher frequency, and eventually into the left half plane, so that it adds to the phase margin rather than reducing it.

The zero is at frequency:

$$\omega_z = \frac{1}{C_C(\frac{1}{g_m} + R_z)} \tag{5.11}$$

Without the compensation resistor, the zero is at $\frac{g_m}{C_C}$. Since the transconductance of MOS transistors running at low current is fairly low, the zero appears at low enough frequency to hurt the stability. A large enough value of R_z will fix this problem.

Open loop Bode plots of the output buffer gain and phase are shown in Figure 5-18. It has a crossover frequency of about 1 MHz, with about 60 degrees of phase margin. The right half plane zero can be seen to have moved near to the second pole at about 10 MHz, but remaining in the right half plane, from the bump in magnitude and sharp falloff in phase. Since the zero is well past crossover, it does the frequency response no harm.

The output buffer consumes 2 μA in the differential pair and 1 μA in the common source output stage, for a total of 3 μA, or 9 μW.

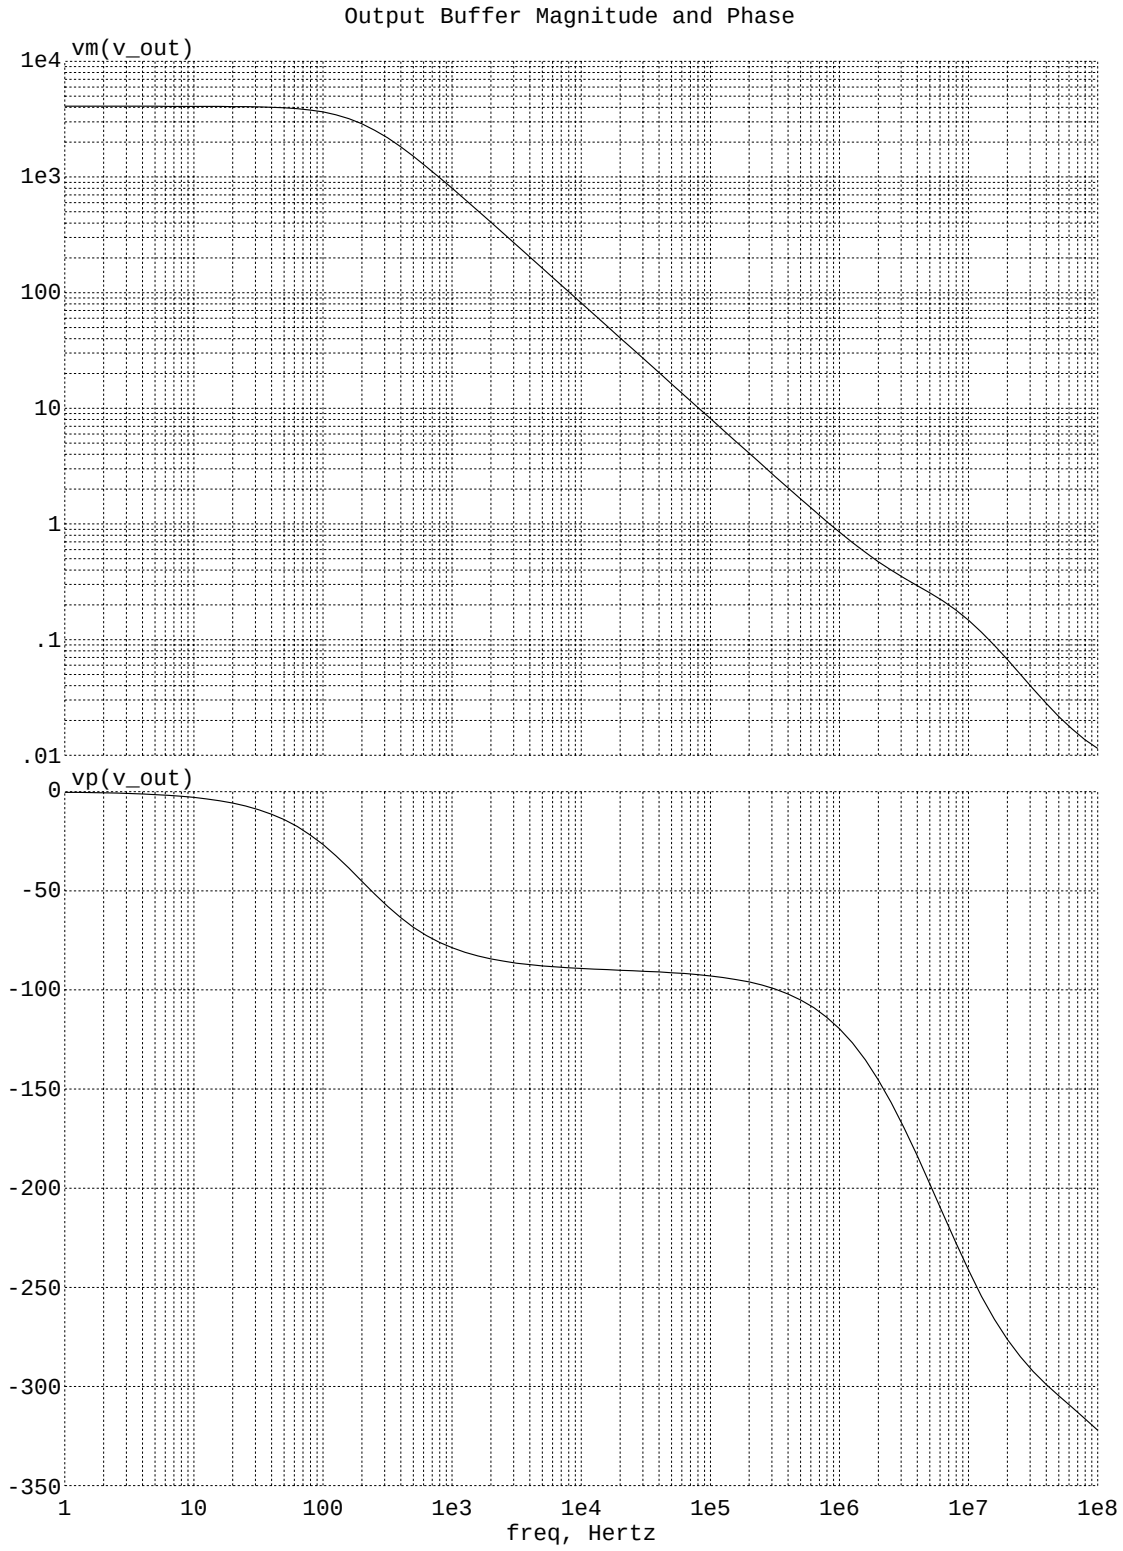

Figure 5-18: Open Loop Output Buffer Gain and Phase

Chapter 6

Layout Concerns

The performance of the completed circuit is very sensitive to layout nonidealities, and great care was taken to ensure that the performance of the fabricated chip would not be significantly hurt by those nonidealities.

The most crucial element of the layout is ensuring that the beam nodes are not allowed to couple in with any switching nodes. The node coming off of the beam and interfacing with the integrator is purely capacitive and very high impedance, and is thus highly sensitive to charge coupling in from the switching nodes. A DC charge will be removed by the switched capacitor DC feedback, but charge dumped every clock cycle cannot be. Avoiding overlap between the beam node and any switching node was also very difficult, due to the complexity of the switching topology interfacing with the beam node. The layout was designed to be fully differential, so any switching charge which couples onto the sensitive node will couple on as common mode, and will nominally be cancelled.

A similar problem was the case in each of the switched capacitor circuits. Switched capacitor circuits contain high impedance nodes which charge is allowed to flow onto and off of, and are thus highly sensitive to charge coupling. Overlaps were avoided wherever possible, and designed to be in common mode where it was not possible. The complexity of the six transistor switching cell is such that an overlap is necessary between the analog and digital waveforms the switch interfaces. The switch was designed so that the clock overlap is the minimum size possible, and there is an overlap of the clock antiphase a minimum distance away from it. This will hopefully cancel much of the charge injection. The overlap occurs on both terminals

of every switch, so any charge injection will also be nominally cancelled by the common mode rejection.

In order for any common mode to be cancelled by the fully differential integrator, the halves of the signal processing need to match perfectly. Process gradients down the chip will cause components fabricated in different spots to have slightly different values, reducing the effectiveness of the fully differential architecture. Using an approximation that the process gradients are linear, they can be cancelled by arranging the elements which are desired to match such that their geometrical centroids are in the same spot. This is done by breaking the each element into two parts, and putting the parts in a two by two grid. The parts in quadrants one and three are connect to add, as are the parts in quadrants two and four. The centroid of each appears at the origin. This technique is called cross-quad arrangement, and was done with the input pair transistors to the op-amps, the feedback capacitors, and the attenuation resistors.

The output buffer tries to accomplish a very accurate gain of five through ratioing capacitors. The capacitance of the parallel plate capacitor is proportional to the area of the capacitor. Any error from the nominal capacitance occurs by over-etching the sides of the capacitor, and causes a capacitance error proportional to the perimeter of the capacitor. Two squares with an area ratio of a factor of 5 will have a perimeter ratio of a factor of $4\sqrt{5}$, and thus any etching error will cause the two capacitors to no longer be scaled correctly. This problem can be solved by selecting the width and length of the capacitors such that the ratio of perimeter to area is the same for all which must maintain consistent scaling. In the case of any etching error, the capacitors will then be off from their nominal value, but the ratio will stay consistent. This was done for the gain-setting capacitors in the output buffer. The disadvantage of this strategy is that very awkward capacitor sizes have to be used, especially for large gains, but it is worth the extra space used for the accuracy gained.

A final layout strategy used to minimize the effects of process variations is a trim, to manually correct errors. It is very important to set the bias current accurately, as it sets gains and time constants in nearly every circuit on the chip. The resistor which performs the voltage to current conversion in the bandgap reference is built with a nominal resistance in series with 8 resistor links, each in parallel with a laser blown fuse. If the current is too large, fuses can be blown, increasing the resistance and decreasing the current to its desired value. The nominal resistance was picked with the simulator to be the smallest value necessary

across process variations necessary to produce the desired bandgap output current of 1 μA.

Chapter 7

Noise Analysis

Noise in the circuit is generated in every device on chip which dissipates energy. The noise from each device is processed through the interface and appears at the output, setting the floor for the minimum detectable signal. The two major contributors to output noise were determined and their levels computed by hand. It is very difficult to calculate by hand the total output noise, however, due to the sheer number of noise sources and the complexity of their transfer functions to the output. To get an accurate prediction of the total output noise, simulation tools were used. The simulated and calculated noise levels for the two major noise sources were matched to approximately 30%, verifying that they are reasonable results.

7.1 Noise Calculations

One of the two most significant noise sources in the circuit is the switch connecting the integration capacitor to the op-amp inverting input. When a switch is closed across a capacitor of value C, a voltage noise of $V^2 = \frac{kT}{C}$ is stored across the capacitor. This can be found by integrating the noise generated in the switch resistance across the total bandwidth of the filter created by the resistance and the capacitor being switched.

If the switch has a resistance R, the voltage noise generated by it is flat across frequency, and equal to $4kTR\ V^2/Hz$, which can be modeled as a voltage source in series with the ideal noiseless resistor. The RC filter has a single pole at $f = \frac{1}{2\pi RC}$, and its magnitude falls of as 20 db/decade thereafter. Since the voltage

noise has a constant amplitude across frequency, the filter can be modeled as a brick-wall filter whose cutoff is at $\frac{\pi}{2}\frac{1}{2\pi RC}$, and has magnitude 0 for all frequencies above cutoff. The product of the noise magnitude and the filter magnitude are then both independent of frequency, and the total integrated noise is the product of the magnitude and the filter cutoff frequency.

This product is equal to $4kTR * \frac{\pi}{2}\frac{1}{2\pi RC} = \frac{kT}{C}$. When the switch is opened, this total noise is stored. The fact that this expression does not depend on the switch resistance is very useful, because the switch resistance is usually not very well known or controlled.

When a switch is opened and closed with a frequency f_c, its noise becomes distributed across frequency. Making the continuous time switched capacitor approximation that the frequencies of interest are well below the switching frequency, its noise magnitude becomes $\frac{kT}{f_c C}\ V^2/Hz$. The noise is distributed evenly up to the switching frequency.

The voltage noise generated in the switch connecting the integration capacitor to the inverting input of the op-amp appears primarily on capacitor C_h of Figure 4-5. C_h = 1 pF and f_c = 100kHz, so voltage noise of 200 $nV/\sqrt{Hz}$ is generated on C_h. As one terminal of C_h is held at virtual ground by the op-amp, this voltage appears directly at the beam, charging the beam parasitic capacitance. The beam parasitic is specified as 1.5 pF, for a total input noise charge of $300 \times 10^{-21}\ C/\sqrt{Hz}$. The equivalent output noise for this charge is found by determining the equivalent output voltage to deliver that charge to the input node.

The output voltage is attenuated by the feedback resistors and presented to the sum of the sense capacitors to deliver charge to the input node. To get $300 \times 10^{-21}\ C/\sqrt{Hz}$ at the input, there needs to be 1.5 $\mu V/\sqrt{Hz}$ on the sense capacitors, so the op-amp output voltage will be 5 $\mu V/\sqrt{Hz}$. This is amplified by gain of 5 of the output buffer, and the total noise at the output is 25 $\mu V/\sqrt{Hz}$. The interface has a sensitivity of 0.5 V/g, so the switch creates an input-referred noise of 50 $\mu g/\sqrt{Hz}$.

The other most significant noise source is the input-referred noise of the op-amp, generated primarily by the input pair of transistors.

The correlated double sampling noise shaping has the effect of multiplying noise by a factor of two at harmonics of the switching frequency. To get a conservative estimate for the noise from the op-amp, its white noise magnitude will be multiplied by two across the entire band.

An input-referred voltage noise e_n $V/\sqrt{Hz}$ is first multiplied by two by the noise shaping, to get $2e_n$ at the input to the interface circuit. It is also multiplied by the aliasing gain of the switching section, which is equal to the ratio of the brick-wall bandwidth of the op-amp to the switching frequency. This gain is $\sqrt{\frac{\pi}{2}\frac{f_{BW}}{f_{clk}}}$. The op-amp bandwidth is 2 MHz and the clock frequency is 100 kHz, so the gain is 10π. This ratio is $\frac{Hz}{Hz}$, and all of the noise calculations are in $V/\sqrt{Hz}$, so the square root of this gain, $\sqrt{10\pi}$ is what should be used. The noise at the input of the sensor is thus $e_n \times 2\sqrt{10\pi}$ $V/\sqrt{Hz}$.

This voltage charges the parasitic beam capacitance. The voltage noise goes to the output of the circuit in the same way as the noise of the switch, scaling by the ratio of the beam parasitic to the sum of the sense capacitances, and by the gain of the interface loop and the output buffer. This total gain is $\frac{1.5pF}{200fF} \times \frac{10}{3} \times 5 = 125$. Multiplied by the gain from the input of the op-amp to the beam, the total output noise is $250\sqrt{10\pi} \times e_n$. The input-referred noise of the op-amp input pair, e_n, is $\sqrt{2\frac{4kT}{G_m}} = 17.6$ $nV/\sqrt{Hz}$, and this appears at the output as 25 $\mu V/\sqrt{Hz}$, or 50 $\mu g/\sqrt{Hz}$.

7.2 Noise Simulation

The noise in the complete accelerometer circuit was also simulated. This provides an estimate of how all of the noise sources sum together, as that would be intractable to calculate by hand.

To include the effects of noise aliasing, a periodic steady state circuit solver needs to be used. This kind of solver is primarily used for RF or switched capacitor circuits, as they have a periodically changing operating point. Rather than linearizing devices around a fixed bias point, a periodically changing bias point is found first, and the bias points of the devices are simulated to be changing periodically, at the clock frequency or carrier frequency. The tool used for these simulations was SpectreRF, from Cadence.

Two plots of the simulated noise are shown. Figure 7-1 includes the magnitude of the total output noise and that of the noise generated in each of the three correlated double sampling switches. Each switch is bipolar, switching with an NMOS transistor and a PMOS transistor. The two curves labeled <i12> make the switch connecting the integration capacitor to the op-amp. The two curves labeled <i11> are the switch connecting the integration capacitor to the beam, and the two curves labeled <i10> are the switch connecting

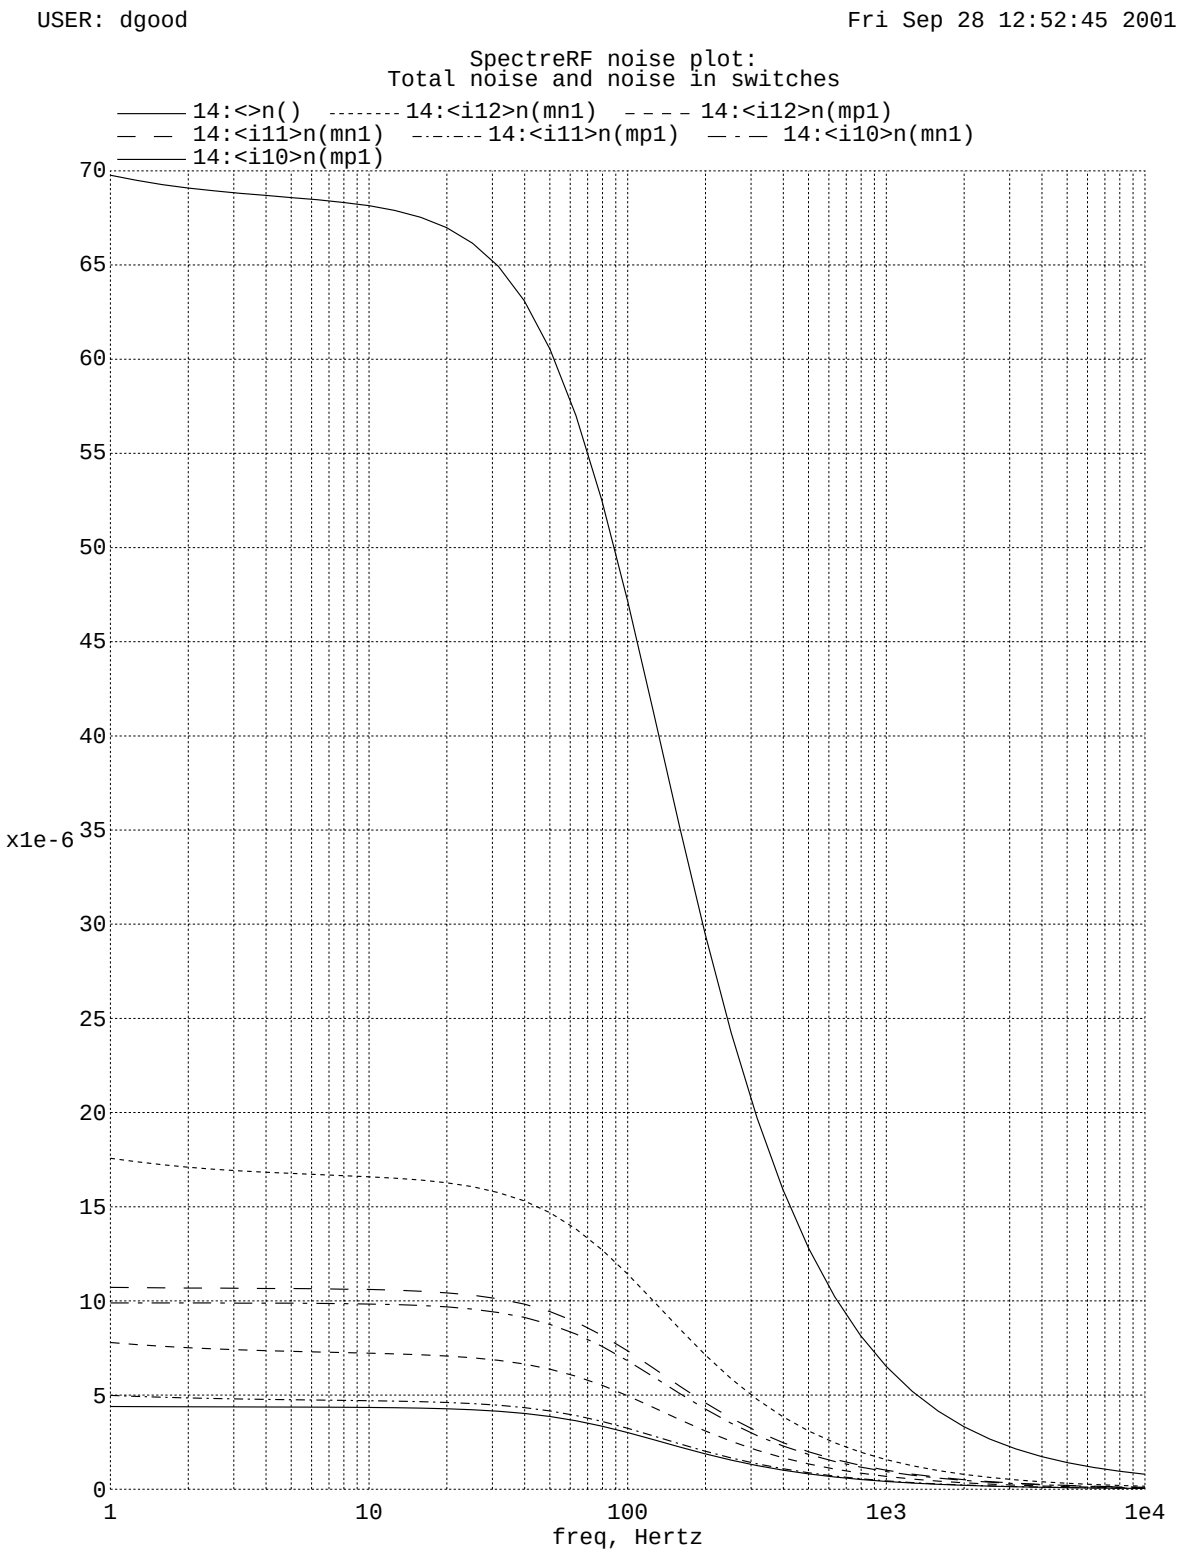

Figure 7-1: Output Referred Switching Noise

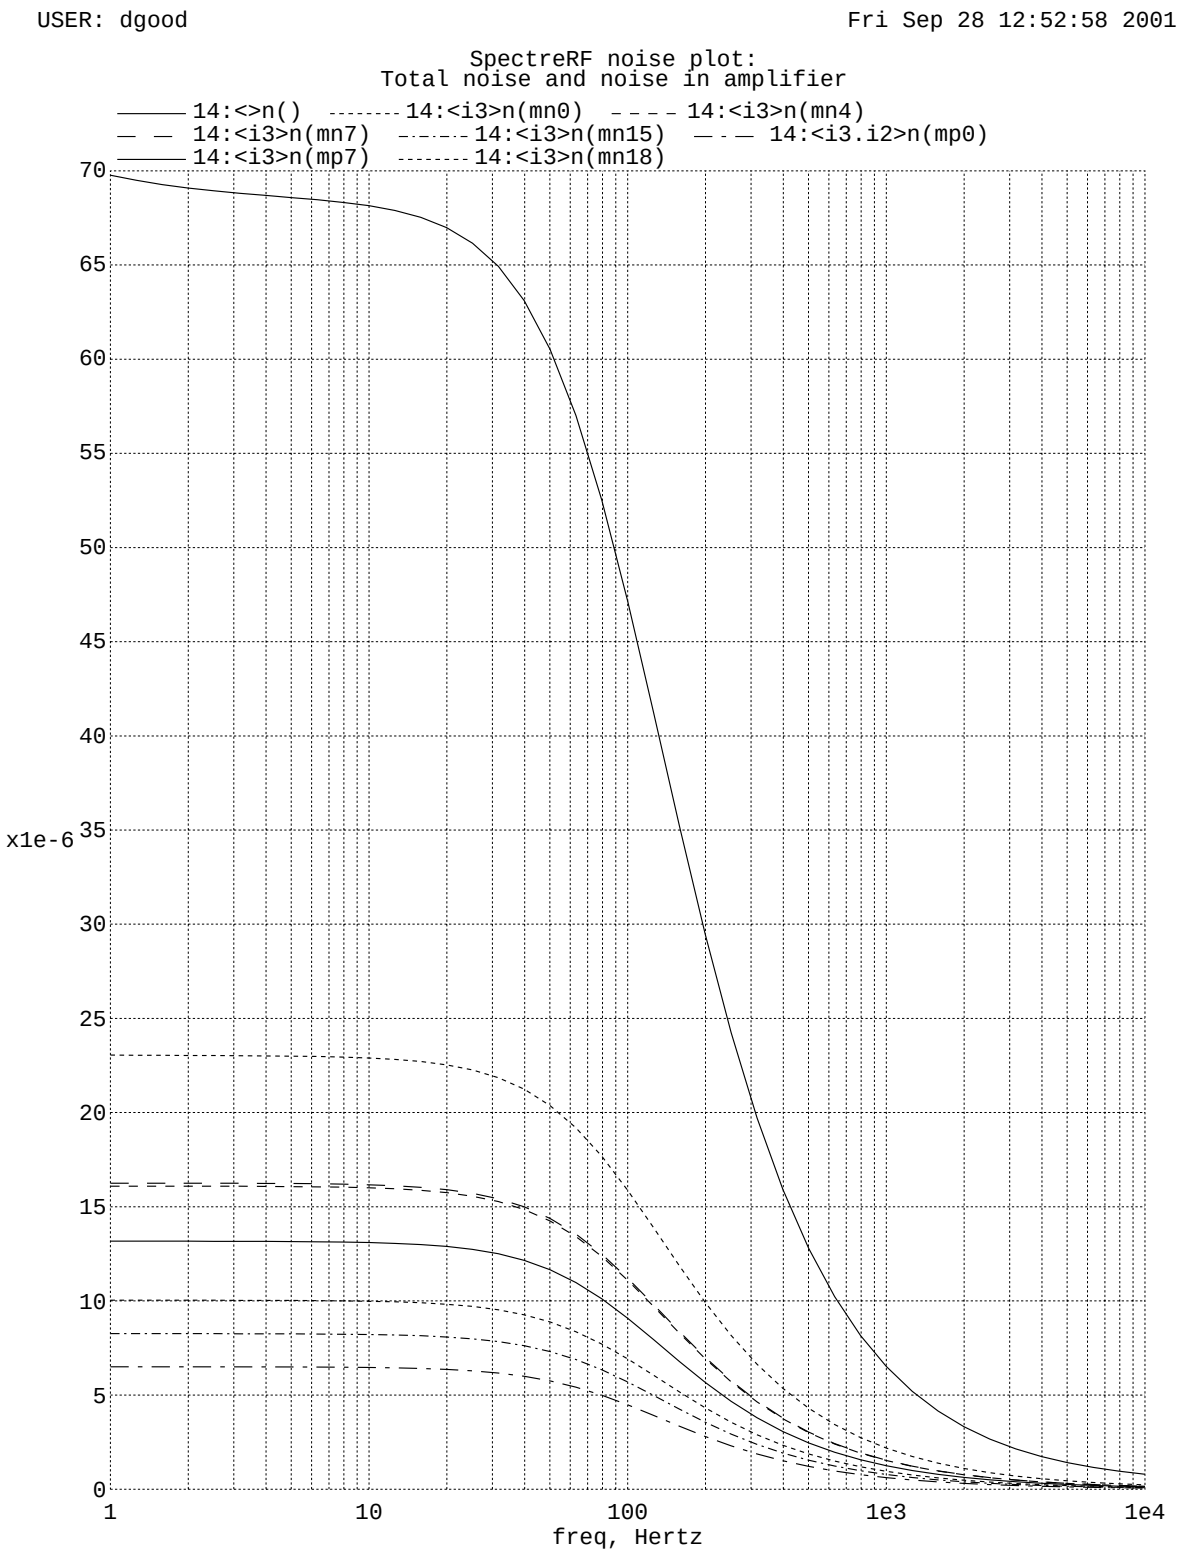

Figure 7-2: Output Referred Amplifier Noise

the input to ground. The summed noise for the switch analyzed by hand is 18.7 $\mu V/\sqrt{Hz}$. This is slightly more than 20% below the calculated value, but is still easily within reason.

Figure 7-2 shows the total output noise and the output noise generated by various transistors in the op-amp. The transistor mn0 is one of the two input transistors, and the remaining curves are the current mirror and cascode transistors. Summing the noise for the two input transistors, they cause about 32 $\mu V/\sqrt{Hz}$ noise at the output. This is about 30% more than was calculated, but is still reasonable.

Visible on the total output noise curve is a component of 1/f noise. This is from the input transistors of the output buffer, which have a significant 1/f noise component, and do not see the noise shaping of the correlated double sampling. However, since they are at the end of the signal path, their input referred noise is small, seen through the gain of the integrator loop. The 1/f noise corner of the circuit is at only a few Hertz, and so the 1/f noise is not a significant contributor to the total noise of the circuit.

The reasonable matching of the simulated and calculated noise results lends confidence to the accuracy of the simulated total noise. The simulated total noise is about 67 $\mu V/\sqrt{Hz}$, or 134 $\mu g/\sqrt{Hz}$. This is well within the desired noise level.

Chapter 8

Conclusions

The response of the accelerometer to a 1g step is shown in Figure 8-1, and a zoom in on the settled value is shown in Figure 8-2. The clock glitches on the output are fairly significant. However, they are at 100 kHz and 200 kHz, and the signal bandwidth is 100 Hz. So, they are easily filtered out, either with an off-chip network, or with fairly large on-chip components. On-chip components would be preferrable, though they would increase die size significantly.

Summing the power consumption of all the cells, the total power of the chip is 206 μW. This easily meets the power target of 300 μW. A white noise level of 67 $\mu V/\sqrt{Hz}$ gives a total noise of 0.83 mV with 3 dB bandwidth of 100 Hz. Compared with a full-scale signal level of 1.5 V, the circuit has 65 dB of dynamic range. This corresponds to a minimum detectable displacement of 25 $\times 10^{-12}$ m, a capacitance change of 1 aF, or an acceleration of 1.6 mg.

It has been empirically shown that an accelerometer can be designed with much lower power consumption than current commercial accelerometers, while still meeting all of the desired specifications. Analog Devices' lowest power accelerometer, the ADXL202E, consumes at least 1.8 mW. Bringing this design from a research project to a commercial product would possibly require more power to be spent, but further revisions could also potentially use power in some sections more efficiently. It is likely that a commercial design would use no more power.

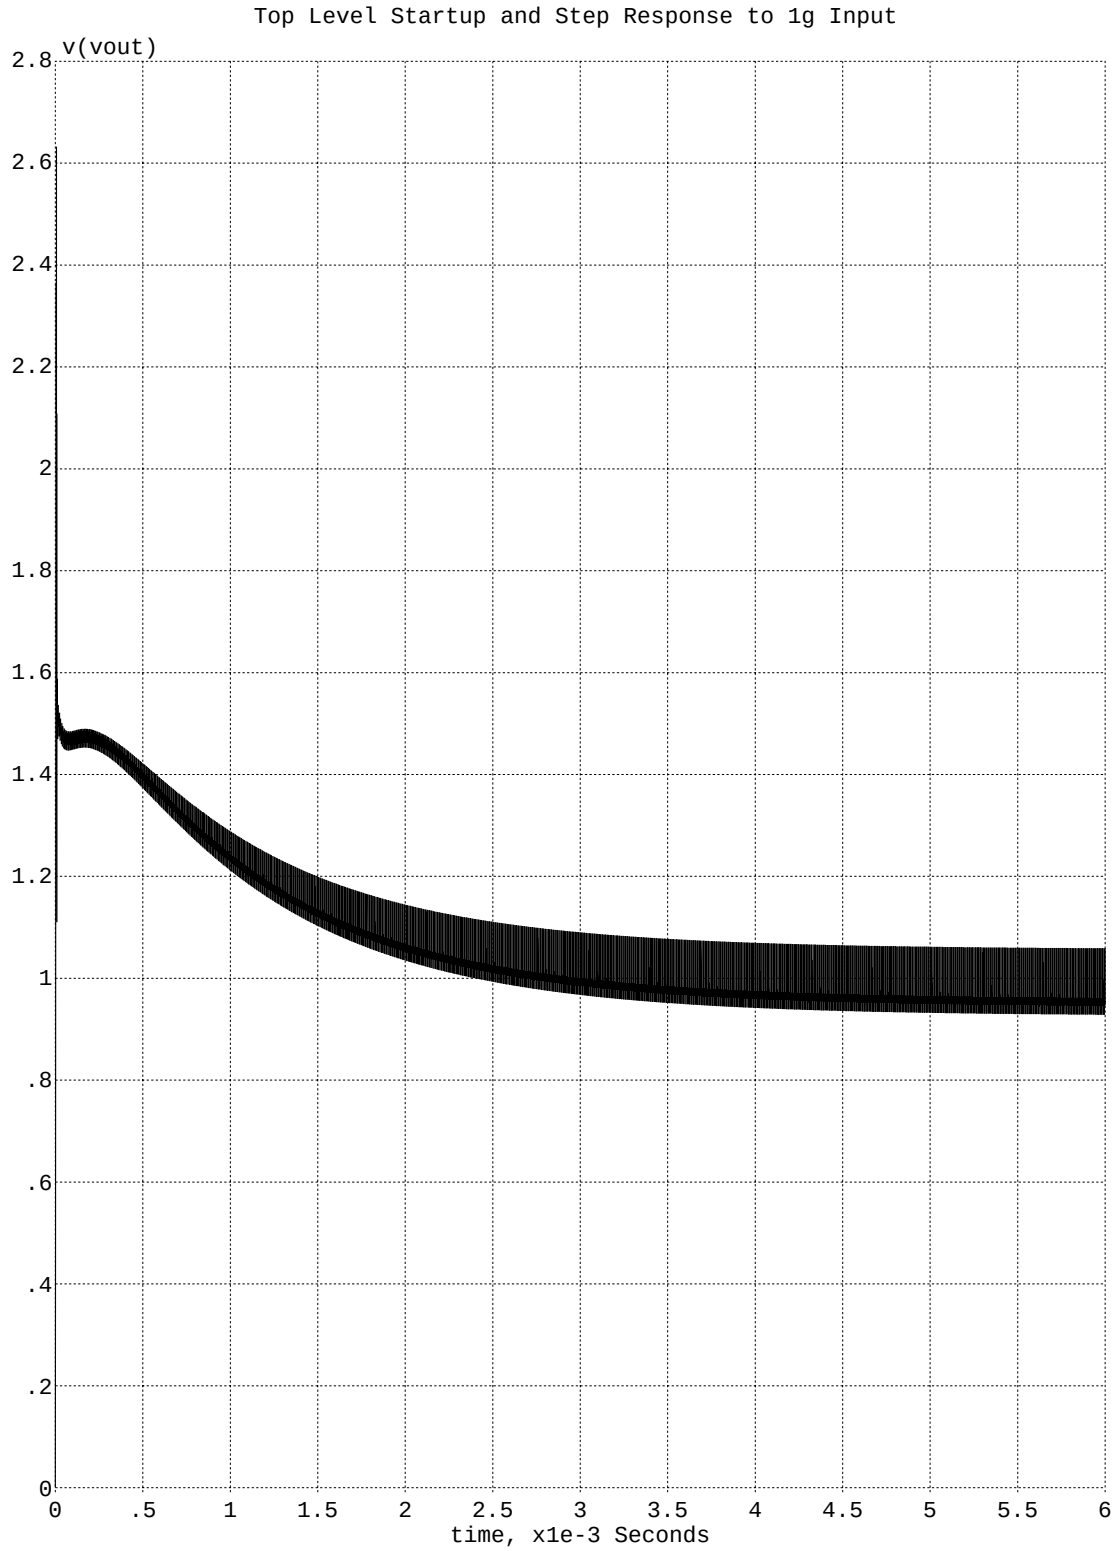

Figure 8-1: Response of Accelerometer to 1g step

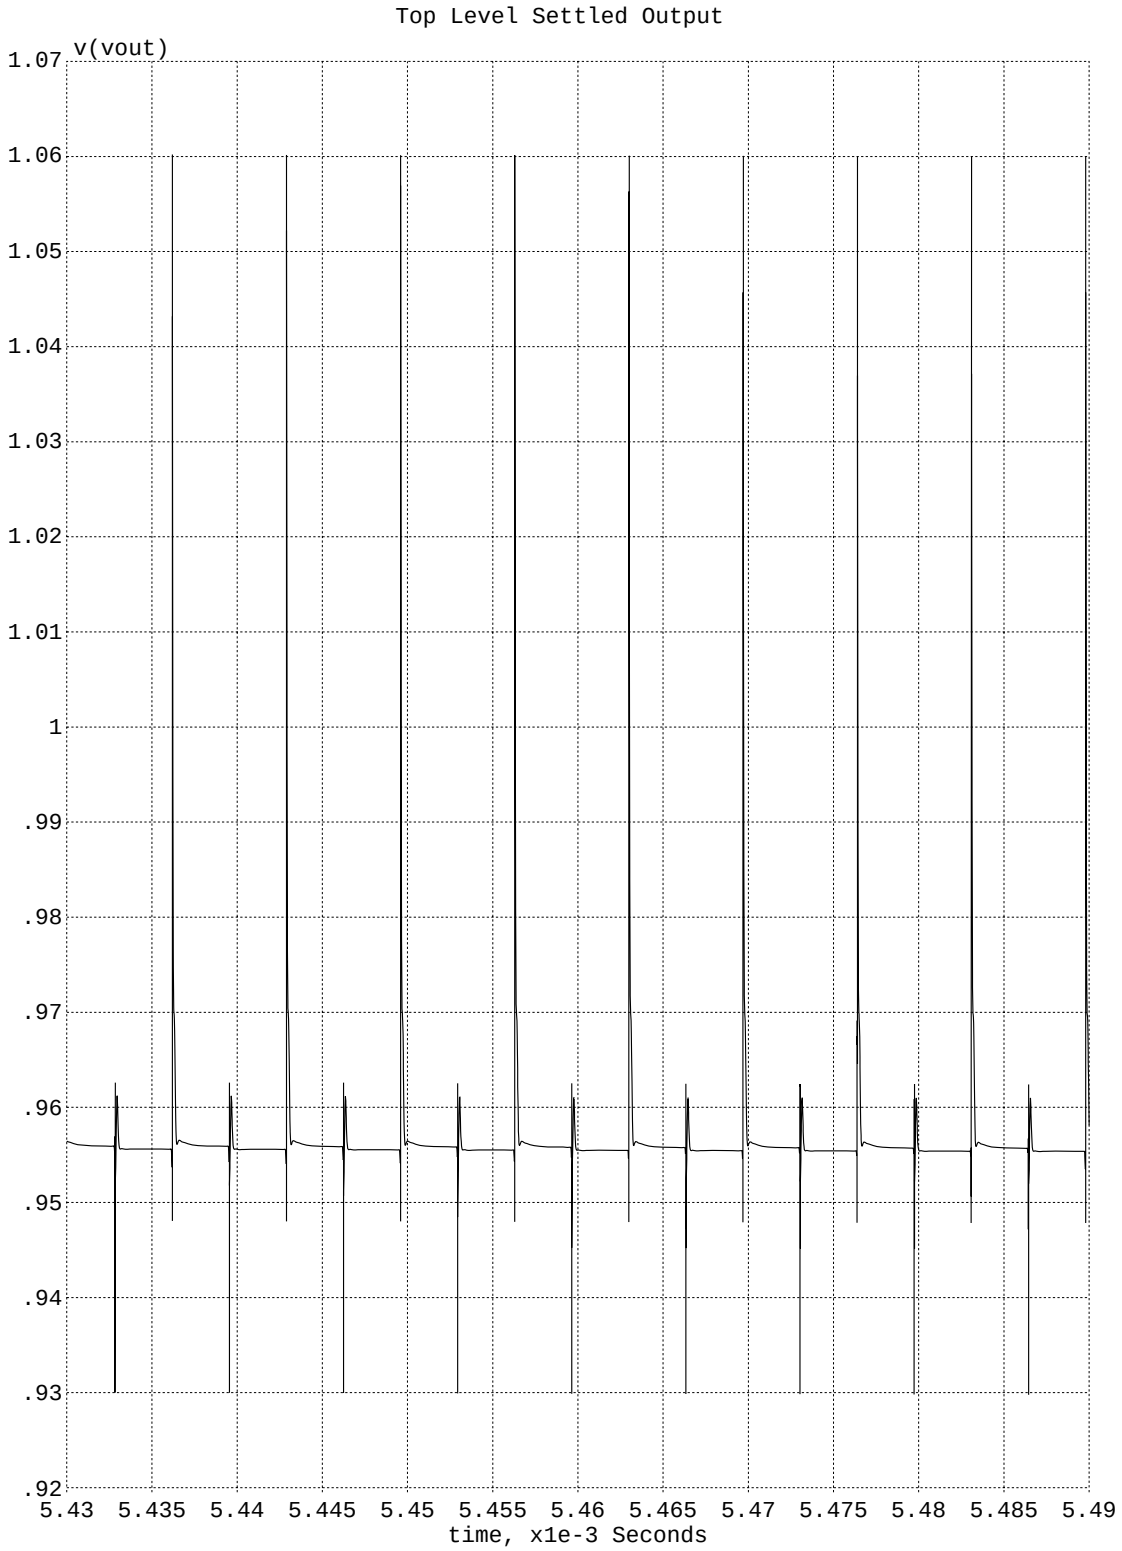

Figure 8-2: Zoom In To Settled Accelerometer Response

Bibliography

[1] Kurt Peterson. Silicon as a mechanical material. *Proceedings of the IEEE*, 70(5), 1982.

[2] Srinivas Tadigadapa. Integration of micromachined devices and microelectronic circuits: Techniques and challenges. *Proceedings of the 43rd IEEE Midwest Symposium on Circuits and Systems*, 2000.

[3] Timothy Brosnihan, James Bustillo, Albert Pisano, and Roger Howe. Embedded interconnect and electrical isolation for high-aspect-ratio, SOI inertial instruments. *International Conference on Solid State Sensors and Actuators*, 1997.

[4] Eric Vittoz. Low power design: Ways to approach the limits. *IEEE International Solid-State Circuits Conference*, 1994.

[5] Thomas B. Gabrielson. Mechanical-thermal noise in micromachined acoustic and vibration sensors. *IEEE Transactions on Electron Devices*, 40(5), May 1993.

[6] Jose Pedro Moreira and Manuel Medeiros Silva. Limits to the dynamic range of low-power continuous-time integrators. *IEEE Transactions on Circuits and Systems I*, 48(7), July 2001.

[7] James K. Roberge. *Operational Amplifiers: Theory and Pratice*. John Wiley, 1975.

[8] K. Nagaraj, K. Singhal, T. R. Viswanathan, and J. Vlatch. Switched capacitor circuits with reduced sensitivity to finite amplifier gain. *Proceedings of the IEEE International Conference on Circuits and Systems*, 1986.

[9] Jorge A. Grilo and Gabor C. Temes. Predictive correlated double sampling switched-capacitor integrators. *1998 IEEE International Conference on Electronics, Circuits and Systems*, Sept. 1998.

[10] William Siebert. *Circuits, Signals, and Systems*. McGraw-Hill, 1986.

[11] Eric Vittoz and Jean Fellrath. CMOS analog integrated circuits based on weak inversion operation. *IEEE Journal of Solid State Circuits*, SC-12(3), 1977.