

Readout Circuits for Frequency-Modulated Gyroscopes

Igor Izyumin



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2015-214

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-214.html>

December 1, 2015

Copyright © 2015, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Readout Circuits for Frequency-Modulated Gyroscopes

by

Igor Igorevich Izyumin

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Engineering – Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Bernhard E. Boser, Chair

Professor Elad Alon

Professor Dorian Liepmann

Fall 2014

Readout Circuits for Frequency-Modulated Gyroscopes

Copyright 2014
by
Igor Igorevich Izyumin

Abstract

Readout Circuits for Frequency-Modulated Gyroscopes

by

Igor Igorevich Izyumin

Doctor of Philosophy in Engineering – Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Bernhard E. Boser, Chair

In recent years, MEMS gyroscopes have become nearly omnipresent. From their origins in automotive stability control systems, these sensors have migrated to a diverse range of applications, including image stabilization in cameras and motion tracking in video games and fitness monitors. A key application for MEMS gyroscopes is pedestrian navigation, which can help provide always-on location in portable devices while minimizing power consumption and infrastructure requirements. While modern smartphones have all of the necessary sensors for inertial navigation, their performance is not sufficient for this application. Consumer-grade MEMS gyroscopes are typically rate-grade devices, and generally have poor bias stability and scale factor accuracy. In addition, their power consumption tends to be too high to enable always-on operation without significantly impacting battery life.

This work describes the first frequency-output MEMS gyroscope to achieve <7 ppm scale factor accuracy and $< 6^\circ/\text{hr}$ bias stability with a 3.24 mm^2 transducer. By implementing continuous-time mode reversal in an FM gyro, the rate signal is modulated away from DC, making the system insensitive to the resonant frequency of the transducer. The scale factor is almost entirely ratiometric, depending primarily on the mechanical angular gain factor of the transducer and the accuracy of the timing reference. Scale factor sensitivity to variations in quality factor, electro-mechanical coupling coefficients, and circuit drift is significantly reduced compared to conventional open-loop and force-rebalance operating modes. Low-power frequency-to-digital converters enable a gyroscope with 91 dB of dynamic range and an estimated power consumption below $150\ \mu\text{W}$ per axis.

To my family.

Contents

Contents	ii
List of Figures	iv
List of Tables	vi
1 Introduction	1
1.1 Motivation	1
1.2 Performance objectives	3
2 Frequency-Modulated Gyroscopes	7
2.1 Background	7
2.2 Quadrature FM operation	8
2.3 Lissajous FM operation	9
2.4 LFM gyroscope error sources	11
2.5 Test setup	14
2.6 Sensor model	18
3 Frequency-to-Digital Conversion	19
3.1 Introduction	19
3.2 Performance requirements	20
3.3 FM demodulation methods	21
3.4 Comparator design	27
3.5 $\Sigma\Delta$ FDC Architecture	43
3.6 Double edge sampling	45
3.7 High-level design	47
3.8 Circuit implementation	50
3.9 Experimental characterization	60
4 LFM Signal Processing	63
4.1 Background	63
4.2 Filtering and resampling	63

4.3	Demodulation reference extraction	67
4.4	Synchronous demodulation	72
4.5	Quadrature extraction	73
4.6	Experimental results	73
5	Conclusion	78
	Bibliography	79

List of Figures

2.1	Lissajous FM gyroscope block diagram.	9
2.2	Layout of quad-mass gyroscope transducer.	14
2.3	Desired mode shapes of gyroscope transducer.	15
2.4	Scanning electron microscope photograph of fabricated structure.	15
2.5	Oscillator differential half-circuit.	16
2.6	Measured frequency noise spectral density for single oscillator channel.	16
3.1	Arctangent-based FM demodulator.	22
3.2	FM demodulation by period measurement.	24
3.3	Simplified equivalent comparator model.	29
3.4	Switching waveforms for comparator model.	30
3.5	Calculated and simulated comparator edge jitter vs. C with different levels of input noise.	31
3.6	Calculated comparator rise time vs. edge jitter due to input noise	32
3.7	Example comparator design.	34
3.8	AM to PM conversion in the presence of offset.	35
3.9	Comparator implementation.	40
3.10	Example comparator circuit schematic.	41
3.11	Example comparator switching waveforms.	41
3.12	Simulated comparator ARW vs. split frequency at the output of each stage.	42
3.13	$\Sigma\Delta$ FDC: (a) conceptual block diagram; (b) timing diagram.	43
3.14	Linearized equivalent model of $\Sigma\Delta$ FDC	44
3.15	Synthesis of $\Sigma\Delta$ FDC structure from second-order $\Sigma\Delta$ modulator	45
3.16	$\Sigma\Delta$ FDC configuration for double edge sampling.	46
3.17	AM rejection via double edge sampling.	47
3.18	ARW due to quantization noise as a function of reference clock frequency for 2nd order loop, $f_{split} = 200$ Hz, $BW = 100$ Hz.	49
3.19	Example fourth-order $\Sigma\Delta$ FDC structure with combined feedback path.	50
3.20	Simulated phase detector error histogram for 2nd, 3rd, and 4th order loops with combined feedback.	51

3.21	PFD and charge pump circuits: (a) PFD schematic, (b) timing diagram, (c) charge pump and integrator schematic.	52
3.22	Counter topologies	53
3.23	Slow counter schematic.	54
3.24	Fast counter schematic.	55
3.25	Dual-stage counter block diagram.	56
3.26	Dual-stage counter overlap compensation example.	57
3.27	Total counter dynamic power as a function of the number of fast counter bits.	58
3.28	SAR ADC block diagram.	58
3.29	ADC comparator schematic.	60
3.30	Measured FDC noise density with 30kHz function generator input.	61
3.31	Measured FDC power spectrum with 100 Hz sinusoidal frequency modulation having 10 Hz deviation amplitude.	62
3.32	Noise floor with and without applied full-scale modulation.	62
4.1	Split frequency sensitivity as a function of decimation ratio and split frequency for a 3-stage CIC decimator.	65
4.2	Power spectral density before and after $24\times$, 3-stage CIC decimator with full-scale modulating signal.	66
4.3	Power spectral density after resampling.	66
4.4	Reference phase extraction and initial error	68
4.5	Reference phase extraction circuits.	68
4.6	Simulated phase estimate error as a function of time for the bang-bang phase detector phase recovery circuit	69
4.7	Simulated phase estimate error as a function of time for the counter-based phase recovery circuit	70
4.8	Bias Allan deviation.	74
4.9	Scale factor Allan deviation.	74
4.10	In-band rate noise power spectral density	76
4.11	Time-domain scale factor measurement	77

List of Tables

1.1	Comparison of commercially available consumer-grade low-power 3-axis MEMS gyroscopes commonly used in smartphones.	2
1.2	Comparison of this work to existing products.	4
3.1	Controller state table	57

Acknowledgments

I would like to thank my advisor, Professor Bernhard E. Boser, for his unwavering support, guidance, encouragement, and honest, constructive critical feedback. I would also like to thank my dissertation committee members, Prof. Elad Alon and Prof. Dorian Liepmann, for their time and valuable feedback.

During my time at Berkeley, it was a great privilege to work with such wonderful friends and colleagues as Mitchell Kline, Richard Przybyla, Octavian Florescu, Stefan Shelton, Karl Skucha, Mischa Megens, Simone Gambini, Rikky Muller, Behnam Behroozpour, Burak Eminoglu, Yu-Ching Yeh, Pramod Murali, Hao-Yen Tang, Mekhail Anwar, Efthymios Papageorgiou, Niels Quack, Matthew Spencer, Steven Callender, Milos Jorgovanovic, and many others. They have truly made my graduate school experience both educational and enjoyable.

I would like to acknowledge the staff of the EECS department, BSAC, and Swarm Lab, including Richard Lossing, Kim Ly, Alex Luna, Mila MacBain, Shelley Kim, and Shirley Salanio. They have all been extremely helpful in dealing with administrative matters.

I am forever indebted to my parents, Natalia and the late Igor Sr., and my brother Oleg for their love, countless sacrifices and endless support and encouragement.

I would like to thank TSMC, InvenSense, and the Thomas W. Kenny group at Stanford University for providing integrated circuit and MEMS fabrication. This work was supported in part by the Defense Advanced Research Projects Agency.

Chapter 1

Introduction

1.1 Motivation

In recent years, MEMS gyroscopes have become nearly omnipresent. From their origins in automotive stability control systems, these sensors have migrated to a diverse range of applications, including image stabilization in cameras and motion tracking in video games and fitness monitors. Every modern smartphone includes a 3-axis gyroscope.

Location data is a key feature enabling many smartphone applications. Location is used for navigation and maps; it can help provide geographically relevant suggestions and targeted advertising; it also provides high-resolution real-time traffic data. Accurate and ubiquitous indoor positioning enables applications such as building automation.

Smartphones use a variety of techniques to obtain position data. The primary location source is the Global Positioning System (GPS). GPS provides accurate and high-resolution position data; however, it does not work indoors and has high power consumption due to the required receiver sensitivity. A typical GPS chip consumes 9 mW during operation [1]. Wi-Fi positioning is a lower-power alternative that has the potential to work indoors; however, it has relatively low accuracy. Much work has been done on beacon-based indoor location systems [2]; however, their deployment requires significant capital investment and ongoing maintenance costs, and such systems are therefore unlikely to ever become truly ubiquitous.

An approach that can help provide always-on, high-resolution location data is inertial navigation, also known as dead reckoning. Such a system generally uses a 6-axis inertial measurement unit, consisting of orthogonally mounted accelerometers and gyroscopes that continuously integrate motion data to compute a location estimate. In the general case, inertial navigation requires extremely high-performance, high-cost, macro-scale

	InvenSense MPU-6050 [3]	Bosch BMG160 [4]	ST A3G250D [5]
Angle random walk	5 mdeg/s/ $\sqrt{\text{Hz}}$	14 mdeg/s/ $\sqrt{\text{Hz}}$	30 mdeg/s/ $\sqrt{\text{Hz}}$
Max rate (min. ARW)	250 deg/s	125 deg/s	245 deg/s
ZRO (over temp.)	± 20 deg/s	± 1.9 deg/s	± 3.8 deg/s
Scale factor (25°C)	$\pm 3\%$	$\pm 1\%$	—
Scale factor (over temp.)	$\pm 2\%$	$\pm 3.75\%$	$\pm 2\%$
Power	8.6 mW	12 mW	15 mW

Table 1.1: Comparison of commercially available consumer-grade low-power 3-axis MEMS gyroscopes commonly used in smartphones.

sensors to obtain even moderate accuracy. Constraining the problem to pedestrian dead reckoning over relatively short time periods makes it possible to obtain reasonable position accuracy with high-performance MEMS sensors [6]. While the long-term accuracy of dead reckoning is generally poor, such a system could be used to provide continuous location data between periodic GPS or beacon-based location updates, reducing power and improving the user experience.

While modern smartphones have all of the necessary sensors for inertial navigation, their performance is not sufficient for this application. Consumer-grade MEMS gyroscopes are typically rate-grade devices, and generally have poor bias stability and scale factor accuracy. In addition, their power consumption tends to be too high to enable always-on operation without significantly impacting battery life. Table 1.1 shows a comparison between several recent consumer-grade 3-axis MEMS gyroscope units. Bias stability is generally not specified for these parts; the zero rate output (ZRO) and scale factor vary significantly over the operating temperature range. While it is possible to calibrate the ZRO by auto-zeroing when no motion is detected, it is difficult to correct scale factor errors. Scale factor errors can cause significant position errors: for example, a 2% error in the gyroscope scale factor would cause a 3.6° heading error after a single 180° turn. This error, in turn, causes a 6 m position error for every 100 m traveled. Since pedestrian travel in a high-density urban environment may involve a relatively large number of turns, scale factor accuracy is a critical specification.

Frequency-modulated gyroscopes promise to address many of these issues. The foremost advantage of FM operation is the intrinsically accurate scale factor, which does not

significantly vary with temperature [7]. Another advantage is the large dynamic range available: while most of the gyros in table 1.1 are restricted to 250 deg/s or less to maintain the specified angle random walk, the FM gyroscope can avoid this trade-off. Finally, virtual mode reversal rejects errors due to cross-axis damping, improving bias stability.

1.2 Performance objectives

Bias stability

Gyroscope performance requirements for an inertial navigation system are difficult to define. The specific navigation system implementation largely determines how gyroscope errors translate into position errors. Nevertheless, it is useful to define a performance metric that quantifies errors introduced by the gyroscope. One such metric is the integrated heading error for a given integration time with zero rate input. This metric ignores scale factor errors, which are input-dependent and which need to be considered separately.

When the gyroscope output is integrated, multiple error sources corrupt the measurement. These error sources can generally be separated by their spectral characteristics into angle random walk (integrated angle error $\propto \sqrt{t}$), bias instability (integrated angle error $\propto t$), and rate random walk (integrated angle error $\propto t^2$). At short integration times, angle random walk (ARW) tends to dominate; the other sources generally dominate for longer integration times.

The most commonly accepted method of characterizing the stability of gyroscopes is the Allan deviation [8]. To compute the Allan deviation for a given time duration τ , a long rate measurement is sliced into pieces of length τ , and the mean is computed for each such piece. Subsequent values are then differenced; the standard deviation of the resulting differences is the Allan deviation. The Allan deviation thus represents the rms random drift for a given integration time. Multiplying the Allan deviation by the integration time provides an estimate of the rms integrated angle error for a given integration time.

A rough estimate of the required gyroscope accuracy can be obtained by considering the requirements of a pedestrian navigation system. Such systems typically estimate the distance traveled using a step counting approach, and use the gyro to keep track of heading [6]. For this calculation, it will be assumed that the pedestrian is traveling along a straight line at a constant speed, and the speed is accurately known. The maximum integrated angle error $\Delta\theta$ over the integration time τ for a given distance error Δp is then

	InvenSense MPU-6050 [3]	This work
Angle random walk	5 mdeg/s/ $\sqrt{\text{Hz}}$	14 mdeg/s/ $\sqrt{\text{Hz}}$
Max rate (min. ARW)	250 deg/s	> 1000 deg/s (estimated)
Bandwidth	≥ 250 Hz	50 Hz (estimated)
Bias stability @ 1 hr	—	5.9 deg/hr
Scale factor accuracy	$\pm 3\%$	± 10 ppm
Power/axis	2.9 mW	< 150 μW (estimated)

Table 1.2: Comparison of this work to existing products.

given by the following expression, where V is the average velocity:

$$\Delta\theta = \frac{1}{\tau} \sin^{-1} \frac{\Delta p}{\tau V} \quad (1.1)$$

If the system goal is to achieve an rms error of 10 m over a 10 minute integration time, and the average walking speed is assumed to be 4 km/h, the maximum allowable Allan deviation is 5.2 deg/hr for $\tau = 10$ minutes.

Scale factor accuracy

Scale factor accuracy is another critical specification for a pedestrian navigation system. Scale factor errors cause a static heading error after every change of direction; such heading errors can result in significant position errors as the pedestrian travels. Because these errors are input-dependent, it is not possible to estimate their magnitude unless some statistics of the input signal are known or assumed.

A rough estimate of the required scale factor accuracy may be obtained by assuming the trip starts with a turn (for example, 180°) and calculating the distance error at the end of the dead reckoning navigation time. In the following equation θ_i is the initial turn angle, Δp is the required distance error, V is the velocity, and τ is the dead reckoning time.

$$SF = \frac{1}{\theta_i} \sin^{-1} \frac{\Delta p}{\tau V} \quad (1.2)$$

For example, if the walking speed V is 4 km/h and the dead reckoning time τ is 10 minutes, achieving $\Delta p = 1$ m distance error after an $\theta_i = 180^\circ$ initial turn requires scale factor accuracy of 480 ppm or less.

Gyroscope bandwidth

Insufficient bandwidth results in a gain error: some energy at higher frequencies is removed by the bandlimiting process. The gyroscope bandwidth should be set to limit the error to some acceptable bound for the fastest possible input signal; in order to do this, the characteristics of the input signal must first be established. In the case of pedestrian navigation, the input signal depends on the biomechanics of the human body. Considerable work exists in this field. In [9], head motion characteristics are measured during various activities, including the measurement of angular rates and accelerations.

A quantitative bandwidth estimate may be obtained by using the measured maximum angular rate and acceleration and assuming some shape for the rate pulse. A convenient pulse shape is a Gaussian pulse, where T is the pulse width:

$$\dot{\theta}(t) = \dot{\theta}_{pk} \exp \left[-\pi \left(\frac{t}{T} \right)^2 \right] \quad (1.3)$$

The angular acceleration is then:

$$\ddot{\theta}(t) = \frac{-2\pi\dot{\theta}_{pk}t}{T^2} \exp \left[\frac{-\pi t^2}{T^2} \right] \quad (1.4)$$

The peak angular acceleration occurs at $t = \pm T/\sqrt{2\pi}$; substituting this into (1.4) and rearranging the resulting expression yields

$$T = \sqrt{\frac{2\pi\dot{\theta}_{pk}}{e\ddot{\theta}_{pk}}} \quad (1.5)$$

The Fourier transform of (1.3) is

$$\dot{\Theta}(f) = \dot{\theta}_{pk}T \exp[-\pi f^2 T^2] \quad (1.6)$$

The energy in a given bandwidth is therefore

$$E(B) = \int_{-B}^B |\dot{\Theta}(f)|^2 df = \frac{\dot{\theta}_{pk}^2 T}{\sqrt{2}} \operatorname{erf}(BT\sqrt{2\pi}) \quad (1.7)$$

The gain error as a function of bandwidth is

$$\varepsilon = 1 - \sqrt{\frac{E(B)}{E(\infty)}} = 1 - \sqrt{\operatorname{erf}(BT\sqrt{2\pi})} \quad (1.8)$$

From [9] we obtain $\dot{\theta}_{pk} = 1.85 \text{ rad/s}$ and $\ddot{\theta}_{pk} = 151 \text{ rad/s}^2$ for the activity with the maximum $\ddot{\theta}_{pk}/\dot{\theta}_{pk}$ ratio, corresponding to the minimum pulse width T and thus maximum bandwidth. For a 50 ppm gain error, the estimated minimum bandwidth is 60 Hz.

Power consumption

Gyroscopes generally have high power consumption. As seen from table 1.1, low-power commercial gyroscopes consume 3 to 5 mW per axis. A typical smartphone is rated for up to 300 hours of standby time with a 9 Wh battery [10], corresponding to an average standby power consumption of about 30 mW. An always-on gyroscope consuming 10 mW would decrease battery life by at least 30%.

A reasonable upper bound for the power consumption of an always-on gyroscope is 2.5% of the total standby power, corresponding to 250 μ W per axis. As will be shown, frequency-modulated operation can readily achieve this goal.

Chapter 2

Frequency-Modulated Gyroscopes

2.1 Background

MEMS gyroscopes measure angular rate by employing the Coriolis effect in a microfabricated vibratory transducer. The most basic transducer comprises a proof mass, sensing and actuation electrodes, and a spring system that allows the proof mass to move along two independent, orthogonal axes. In the conventional amplitude-modulated operating mode, one of the axes of the transducer is driven at its natural frequency. When angular rate is applied, the direction of vibration is maintained relative to an inertial reference frame. From the rotating reference frame of the gyroscope, this appears as a transfer of vibrational energy from the drive axis to the sense axis, with the magnitude of this effect proportional to the applied angular rate.

The most common readout scheme used with MEMS transducers is the open-loop rate mode. In this scheme, the drive axis oscillation is maintained at a constant amplitude. As the gyroscope is rotated, some of this energy is again transferred to the sense axis, where it is dissipated by the sense axis damper. At the same time, the amplitude control loop replenishes the drive axis energy to maintain a constant amplitude. As a result, the sense axis displacement is proportional to the applied angular rate. This mode has the advantage of being simple to implement, and is commonly used with consumer-grade gyroscopes.

A related mode is force-rebalance (or force-feedback); it is commonly used in higher-performance parts [11]. The force-rebalance mode uses a feedback loop to null out the sense mode oscillation, and measures the force required to do so. This force is proportional to the angular rate. Force-rebalance operation eliminates scale factor dependence on quality factor and frequency split, allowing mode-matched operation and thus re-

ducing power consumption [12]. However, the inherent complexity and the difficulty of ensuring loop stability with the presence of parasitic modes makes this approach relatively unattractive for low-power, inexpensive consumer-grade parts. Furthermore, achieving accurate scale factor still requires precise control of the feedback forces, which is quite difficult.

The design of the transducer structure is generally dictated by the operating mode to be used. Open-loop gyroscopes use a highly-asymmetric transducer, with a relatively large split between the natural frequencies of the drive and sense axes (typically 1-2 kHz). This significantly increases the power consumption of the gyroscope, since the Coriolis force drives the sense axis far from its natural resonance frequency, and the signal is therefore greatly attenuated. However, the transducer has a relatively flat frequency response in this region, and scale factor sensitivity to the frequency split and the quality factor is reduced. Effectively, a large frequency split allows acceptable scale factor stability at the cost of higher power consumption. Frequency-modulated operation allows this trade-off to be completely avoided.

2.2 Quadrature FM operation

The most basic FM gyroscope operating mode is known as quadrature FM (QFM) [7,13]. In this mode, both axes are driven by oscillator loops at their natural frequency. A tuning loop adjusts the resonant frequency of one or both axes to ensure a 90° phase relationship between their displacements. In the case where the displacements of the two axes are equal, the proof mass orbits in a circle.

In the presence of angular rate, the proof mass continues to orbit undisturbed at its natural frequency (in an inertial frame of reference). However, since the pickoff electrodes are rotating relative to the orbit, the observed frequency shifts by the angular rate with a unity scale factor. In effect, this operating mode is similar to whole-angle readout, where the scale factor is also unity. However, unlike the whole-angle operating mode, the proof mass displacement envelope is rotationally symmetric, obviating the need for a complex controller that must keep track of the proof mass oscillation angle.

In a practical MEMS gyroscope, the scale factor is not quite unity because of the angular gain factor. This factor occurs because some vibrating parts of the gyroscope (such as the springs) are constrained to move along only one mechanical axis, and therefore cannot generate Coriolis force. Because the angular gain factor is set by mass ratios, it is very stable and does not drift significantly over time.




Figure 2.1: Lissajous FM gyroscope block diagram. The red outline shows the readout circuits that will be analyzed in the rest of this work.

One major difficulty with the QFM operating mode is that any drift of the center frequency is directly added to the measured rate output. Because the resonant frequency is typically thousands of times larger than even the full-scale angular rate, even very small frequency shifts result in extremely large rate offsets. While offset drift may be significantly reduced by employing two counter-rotating proof masses (with opposite rate sensitivity), the achievable temperature tracking accuracy is not sufficient to make this approach competitive with AM gyroscopes [13].

2.3 Lissajous FM operation

One way to avoid the issue of QFM frequency drift is to periodically reverse the direction of the orbit. This may be done by periodically switching the phase shift between the X and Y oscillations from $+90^\circ$ to -90° , and vice versa. However, the lock time of the tuning loop would severely restrict the gyroscope bandwidth.

Periodic mode reversal may also be accomplished simply by allowing the two gyroscope axes to run at slightly different frequencies; the proof mass now follows a Lissajous curve instead of a circle [7]. This is mathematically equivalent to using a ramp as the setpoint for the QFM tuning loop, where the slope of the ramp is equal to the difference of the axis frequencies. As a result, the rate is modulated by a sinusoid with a frequency equal to the difference between the oscillation frequencies of the two axes. In the case of the ideal gyroscope transducer with no cross-coupling between the two axes, the axis frequencies are given by the following expression, where ω_{ox} and ω_{oy} are the natural axis frequencies of the X and Y axes, v_{xa} and v_{ya} are the velocity amplitudes of each axis, α_z is the angular gain factor of the transducer, Ω_z is the input angular rate, and

$$\Delta\phi_{xy}(t) = (\omega_{ox} - \omega_{oy})t.$$

$$\dot{\phi}_x(t) = \omega_{ox} - \frac{v_{ya}}{v_{xa}}\alpha_z\Omega_z(t) \sin \Delta\phi_{xy}(t) \quad (2.1)$$

$$\dot{\phi}_y(t) = \omega_{oy} - \frac{v_{xa}}{v_{ya}}\alpha_z\Omega_z(t) \sin \Delta\phi_{xy}(t) \quad (2.2)$$

Modulating the rate sensitivity causes the rate signal to move from DC to the frequency of modulation; this is identical to the chopper stabilization technique commonly used in precision amplifiers. As a result, the gyroscope is made insensitive to slow variations in the natural frequency of the transducer, as might result from temperature fluctuations and other drift sources. Furthermore, the scale factor of the gyroscope depends only on the angular gain factor α_z (a stable, dimensionless parameter set by the transducer geometry) and the velocity amplitude ratios v_{xa}/v_{ya} and v_{ya}/v_{xa} . If the demodulated X and Y channel outputs are summed, the sensitivity to velocity amplitude mismatch is greatly reduced due to the reciprocal summation of the two amplitude ratios. Suppose $v_{ya} = (1 + \varepsilon)v_{xa}$. The scale factor error is then given by:

$$\frac{\frac{1}{1+\varepsilon} + \frac{1+\varepsilon}{1}}{2} - 1 = \frac{\varepsilon^2}{2(1 + \varepsilon)} \approx \frac{\varepsilon^2}{2} \quad (2.3)$$

The above result illustrates a significant advantage of the LFM gyroscope: the scale factor accuracy is decoupled from the gain accuracy of the circuits. For example, a relatively large mismatch of 0.5% between the X and Y axis velocities results in a scale factor error of only 12 ppm.

LFM operation presents a trade-off between split frequency, bandwidth, and power. Because the rate signal is amplitude-modulated onto a sinusoidal carrier at the split frequency, the gyroscope bandwidth is always less than the split frequency. At the same time, increasing the split frequency requires decreasing the amount of white phase noise for a given resolution spec, increasing power. A somewhat similar trade-off is seen in AM gyroscopes: as the split frequency is increased, usable bandwidth and power also increase. However, unlike the AM gyroscope, LFM operation has excellent scale factor stability even with relatively low split frequencies, allowing power consumption to be significantly reduced if the required bandwidth is relatively small. A disadvantage of LFM operation is the possibility of aliasing in the presence of signal components that exceed the split frequency. For example, a sinusoidal rate input at exactly the split frequency will appear as a DC rate offset at the output of the gyroscope.

2.4 LFM gyroscope error sources

In order to start the design of the readout circuits, it is necessary to first understand the error sources that limit the ultimate performance of the LFM gyroscope system.

Transducer errors

Numerous errors arise from imperfections in the micromechanical gyroscope transducer. These errors arise from manufacturing imperfections and can also be due to the geometry of the structure. Gyroscope errors can be categorized into errors affecting a single axis, and errors that introduce parasitic coupling between the two axes.

Errors that affect a single axis can affect the effective mass, stiffness, and quality factor. Errors in the effective mass and stiffness affect the resonant frequency of the axis and the gyroscope split. LFM operation depends on these errors being reasonably well-controlled in order to ensure a predictable distribution of split frequencies, but they do not otherwise affect the performance of the gyroscope. Variations in quality factor affect the scale factor of the gyroscope in open-loop AM modes. In the LFM mode, the oscillator circuit cancels the axis losses, and variations in the quality factor do not contribute to scale factor errors.

Most errors in the gyroscope arise from parasitic coupling between the two axes. The parasitic coupling errors are generally classified into cross-stiffness and cross-damping terms. These arise from different mechanical sources. Because they have different phase shifts, they contribute different errors to the gyroscope. When these errors are included, the LFM equations assume the following form, where Ω_c terms represent cross-damping and Ω_k terms represent cross-stiffness errors [7].

$$\dot{\phi}_x(t) = \omega_{ox} - \frac{v_{ya}}{v_{xa}}(\alpha_z \Omega_z(t) - \Omega_{c,yx}) \sin \Delta\phi_{xy}(t) + \Omega_{k,yx} \frac{v_{ya}}{v_{xa}} \cos \Delta\phi_{xy}(t) \quad (2.4)$$

$$\dot{\phi}_y(t) = \omega_{oy} - \frac{v_{xa}}{v_{ya}}(\alpha_z \Omega_z(t) + \Omega_{c,xy}) \sin \Delta\phi_{xy}(t) + \Omega_{k,xy} \frac{v_{xa}}{v_{ya}} \cos \Delta\phi_{xy}(t) \quad (2.5)$$

Anisostiffness errors

Anisostiffness, commonly known as quadrature, is an important source of bias instability in gyroscopes. In the most general case, this error source includes any force that is applied to an axis of a gyroscope in proportion to the displacement of the other axis. In the linear model of the gyroscope, anisostiffness refers to the off-diagonal spring terms coupling the two nominally orthogonal axes. Linear errors may be modeled as

misalignment of the spring axes relative to the electrodes [7]. In the linear case, the cross-stiffness terms are symmetric.

This error may arise from non-linear sources. One such source is related to spring non-idealities. In the ideal case, gyroscope springs allow motion along only one axis, and do not cause displacement in the perpendicular direction. However, all MEMS flexures require at least some displacement in the orthogonal direction in order to function. While the use of folded flexures and symmetric layout minimizes these effects, any remaining asymmetry (e.g. due to fabrication mismatches) can nevertheless result in motion along the orthogonal direction. This is an entirely nonlinear and non-reciprocal effect, and symmetry of the stiffness matrix is therefore not assured.

The effect of quadrature error in the AM gyroscope is significant motion on the sense axis in phase with the displacement of the drive axis. This motion appears in quadrature with the Coriolis signal, and is nominally rejected by the demodulator. Likewise, in the LFM gyroscope, quadrature error results in a tone at the split frequency that is in quadrature with the rate signal. A large quadrature error uses up dynamic range in the analog front-end, potentially creating linearity issues and/or increasing the noise floor. While quadrature error may be rejected by synchronous demodulation, any phase errors in the demodulator reference can cause it to add to the gyroscope's zero rate output. If the amount of quadrature error or the phase error drifts over time, this can contribute a large bias drift component.

As an example, suppose the gyroscope has $1000^\circ/\text{s}$ of quadrature error, and the phase detector is accurate to within 0.1° ; both of these values are typical for consumer-grade gyroscopes. The zero-rate output is then $1.75^\circ/\text{s}$. If the phase detector error drifts by 1%, the zero-rate output will shift by $63^\circ/\text{hr}$. Clearly, a large quadrature error makes it very difficult to achieve good bias stability.

Anisodamping errors

Anisodamping is an error source that includes any force applied to a gyro axis in proportion to the velocity of the other axis. In the conventional operating mode, it is indistinguishable from applied angular rate. Like anisostiffness, it may be modeled linearly as a misalignment of the damper axes with the electrodes, in which case the cross-damping terms are symmetric ($\Omega_{c,xy} = \Omega_{c,yx}$) and will be cancelled when the FM-demodulated X and Y signals are added together. However, this error source may also arise from nonlinear effects (such as squeeze and slide film damping of an asymmetric structure), in which case it is not necessarily reciprocal and will not be completely cancelled. Max-

imizing the quality factor of the structure minimizes all damping terms and is the best defense against this error.

Oscillator errors

The oscillator loops contribute several errors to the LFM gyroscope, including noise, amplitude mismatch, and amplitude ripple.

Oscillator noise

An excellent analysis of Pierce oscillator noise can be found in [14]. The phase noise of a linear Pierce oscillator is given by:

$$S_{\phi_n^2}(\Delta\omega) = \frac{(1 + \gamma)kT\omega}{2QE_m\Delta\omega^2}$$

In the above equation, γ is a constant that represents the excess noise contributed by the sustaining circuit (if $\gamma = 0$, only the transducer thermal noise is included). ω represents the oscillation angular frequency, and $\Delta\omega$ represents the frequency offset. Q is the quality factor of the transducer, and E_m is the mechanical energy present in the resonator.

The noise spectrum predicted by the above expression has a $1/f^2$ shape around the oscillation frequency. The spectrum of the frequency noise is therefore white.

In the presence of transducer or circuit nonlinearity, it is possible for the flicker noise of the oscillator transistors to be upconverted to the oscillation frequency, changing the phase noise shape to a $1/f^3$ shape near the carrier. After FM demodulation, this error will appear as $1/f$ frequency noise. In practice, the oscillator should be designed to be sufficiently linear to keep the flicker noise corner far below the LFM split frequency. This is not difficult in the case of high-Q resonators, since only a very small drive signal is needed to sustain oscillation. It is important to note that the corner of the upconverted flicker noise is independent of the flicker noise corner in the baseband, and is determined entirely by the average value of the impulse sensitivity function. In a perfectly linear oscillator, the ISF has no DC component and $1/f$ upconversion does not occur [14,15].

Amplitude mismatch

As mentioned previously, amplitude mismatch between the X and Y axes can cause a scale factor error. Because of reciprocal summation, the matching requirements are




Figure 2.2: Layout of quad-mass gyroscope transducer.

quite modest: 0.5% matching is sufficient to ensure scale factor stability on the order of 12 ppm. While much better matching is generally possible for monolithic circuit components, the error budget must also include the parasitic capacitances at the sense terminals of the transducer. If these parasitics cannot be accurately matched, a feedback amplifier can be used to neutralize them.

Amplitude ripple

In LFM operation, the gyro alternates between QFM (90° phase shift) and whole angle (0° phase shift) operating modes. When the proof mass moves in a straight line, applied angular rate will attempt to transfer energy between the axes, resulting in sinusoidal amplitude variations as the phase shift between the axes varies. This effect is mostly suppressed by the amplitude regulator in the oscillator loop, though some residual amplitude variation will remain. Even in the absence of angular rate, quadrature and cross-damping errors will cause amplitude ripple to be present. The FM demodulator must reject this error source to avoid degrading scale factor and bias stability.

2.5 Test setup

Transducer

A quad-mass architecture [16] was chosen to implement the LFM gyroscope transducer. The layout is shown in fig. 2.2. The quad-mass layout is fully symmetric, thus maximiz-




Figure 2.3: Desired mode shapes of gyroscope transducer.




Figure 2.4: Scanning electron microscope photograph of fabricated structure.




Figure 2.5: Oscillator differential half-circuit.




Figure 2.6: Measured frequency noise spectral density for single oscillator channel.

ing quality factor and vibration rejection for both axes. Levers and coupling springs are used to move unwanted modes to higher frequencies. Differential drive and sense electrodes select the desired mode and help reject vibration and other common-mode disturbances. Folded springs are used to connect the proof mass to the frames; the springs are designed to minimize force transfer to the other mode, thus reducing quadrature. The transducer is designed to prevent the transfer of package stress to the springs by using a single anchor point for each frame. This reduces a significant source of drift and helps improve vibration rejection. The transducer was fabricated in the epi-seal HDXI process. An SEM micrograph is shown in fig. 2.4.

Sustaining loop

The sustaining loop was implemented as a Pierce oscillator, a low-power and high-performance circuit topology commonly used with crystals [14]. The circuit was implemented using op-amps and discrete components on a printed circuit board. Figure 2.5 shows the schematic of a differential half-circuit of one oscillator channel. The sense port of the MEMS transducer is connected to an active integrator. A very large resistor is used for DC biasing. By maintaining the sense terminal at a virtual ground, the active integrator ensures high linearity and maintains a more accurate phase shift and high CMRR by neutralizing parasitic capacitance at the sense node. The latter advantage is particularly important in the case of a PCB implementation, since the parasitic capacitances are both large and poorly controlled, resulting in significant mismatch between the two differential electrodes.

The output of the first integrator stage is fed to a second active integrator. It is then attenuated by a variable attenuator, inverted, and fed to the gyroscope drive port. An amplitude control loop adjusts the loop gain to maintain a constant envelope at the output of the first integrator. This output is also fed to the frequency demodulator. Figure 2.6 shows the measured noise spectral density for a single oscillator channel. Note that the frequency spectrum is generally white, implying a $1/f^2$ phase noise shape; this is the expected behavior. The vertical axis is scaled assuming $\alpha_z = 1$. The large tone at 100 Hz is the zero-rate output; the tone near 200 Hz is a harmonic of the split frequency, and is likely due to parasitic frequency modulation due to electrode nonlinearity. The tone at 60 Hz is a result of power line interference coupled into the sense node.

2.6 Sensor model

This scope of this work is limited to the design of readout circuits for the LFM gyroscope. The included circuits are delineated in figure 2.1; the transducer and the oscillator are excluded. In order to account for the errors introduced by the excluded components, it is necessary to create a model of the oscillator output.

The following expressions give the time-domain displacement of the X and Y channels in LFM operation. Each signal consists of a cosine with amplitude A and an AM component $x_{am}(t)$ or $y_{am}(t)$. The AM component represents the residual amplitude ripple due to finite amplitude control gain, as well as the AM component of amplitude noise. The argument of each cosine is comprised of a frequency modulation term (the integral), as well as a phase noise term $\phi_n(t)$. The arguments of the integrals represent the instantaneous angular frequencies, and are given by equations (2.4) and (2.5).

$$x(t) = [1 + x_{am}(t)]A_x \cos \left[\int_0^t \dot{\phi}_x(\tau) d\tau + \phi_{nx}(t) \right] \quad (2.6)$$

$$y(t) = [1 + y_{am}(t)]A_y \cos \left[\int_0^t \dot{\phi}_y(\tau) d\tau + \phi_{ny}(t) \right] \quad (2.7)$$

Angular rate is extracted by FM demodulating $x(t)$ and $y(t)$ to extract $\dot{\phi}_x$, $\dot{\phi}_y$ and the phase reference $\Delta\phi_{xy}$, and then synchronously demodulating the sine term from $\dot{\phi}_x$ and $\dot{\phi}_y$ to extract the angular rate Ω_z . The FM demodulator must be designed to reject the AM terms, since they do not contain usable information and can corrupt the FM measurements.

As mentioned previously, the oscillator circuit noise and transducer noise consist of $1/f^2$ phase noise, which is equivalent to white FM noise. Note that only the tank current has this noise shape, whereas noise taken from the output of the oscillator will have an approximately white spectrum near the frequency of oscillation. For minimum noise, the FM demodulator signal chain should be fed directly from the transducer.

The FM components of the signal consist of the desired LFM tone, as well as cross-damping and quadrature components. For the purposes of error analysis, it will be assumed that the frequency split is 100 Hz, the quadrature offset is 150 deg/s, and the cross-damping offset is 1 deg/s. These values are close to the measured parameters of the test device.

Chapter 3

Frequency-to-Digital Conversion

3.1 Introduction

The frequency demodulator is a key part of any frequency-modulated gyroscope. This block largely determines the resolution, scale factor stability, bias stability, and power consumption of the sensor.

Frequency modulation was first used in early communication systems. Early high-power transmitters could not be keyed on and off, and frequency-shift keying was used to transmit telegraphic signals; the selectivity of the receiver converted the FSK signal to an amplitude-modulated one. Later work employed FM in an attempt to reduce the bandwidth required by an AM transmission by reducing the peak frequency deviation. Carson's analysis of FM [17] showed this to be impossible, and for over a decade FM was considered to be vastly inferior to AM in every respect. Armstrong [18] was the first to recognize the SNR advantages of wideband FM for high-fidelity broadcasting and to develop a practical FM system.

In the process of frequency modulation, a modulating signal $m(t)$ modulates the instantaneous frequency of a sinusoid:

$$x(t) = A \cos \left(\omega_0 t + \int_0^t m(\tau) d\tau \right)$$

Note that the amplitude of $m(t)$ has units of frequency. The maximum value of $|m(t)|$ is known as the frequency deviation. In wideband FM, the frequency deviation is larger than the bandwidth of $m(t)$. For example, broadcast FM uses 75 kHz of deviation to modulate an audio signal with a bandwidth of 15 kHz. However, FM gyroscopes are narrowband FM systems: a full-scale 2000 deg/s rate input corresponds to a deviation

of 5.6 Hz, while the bandwidth of the modulating signal is typically at least 50 Hz. In general, the bandwidth of $m(t)$ is much smaller than the center frequency ω_0 , and the maximum deviation is also assumed to be small compared to ω_0 .

In general, the amplitude A is not constant. For example, any white noise that is added to an FM signal by the processing electronics may be decomposed into an amplitude and a phase noise component. In FM gyroscopes, input rate and quadrature errors cause amplitude ripple, which may not be completely suppressed by the amplitude control loop. These effects can cause significant errors in the demodulator.

3.2 Performance requirements

The performance specifications for the FM demodulator are driven by the gyroscope performance objectives in Table 1.2. In the LFM mode, the amplitude-modulated rate signal is located in the band $f_{split} \pm BW$. In order to avoid aliasing, the split frequency must be somewhat larger than the highest-frequency input component; if the input has a bandwidth of 60 Hz, the split frequency should be 100 Hz or more.¹

The angle random walk is determined by the noise of the oscillator circuit, the comparator (if any), and the FM demodulator (fig. 2.1). If the 10 mdeg/s/ $\sqrt{\text{Hz}}$ budget is split equally between the front-end and the FM demodulator, the in-band noise density at the output of the FM demodulator must be 10mdeg/s/ $\sqrt{\text{Hz}}/(2 \cdot 360^\circ) = 13.9 \mu\text{Hz}/\sqrt{\text{Hz}}$; for a 60 Hz gyro bandwidth, the integrated in-band rms noise is 152 μHz . The factor of two in the denominator accounts for the noise allocation between the comparator and the oscillator, and for the two AM sidebands; each of these contribute a factor of $\sqrt{2}$.

The maximum angular rate is determined by the input range of the FM demodulator. To accommodate 2000 deg/s, the demodulator must allow an input deviation of ± 5.6 Hz, plus some additional range to allow for quadrature errors and center frequency drift.

Scale factor accuracy considerations require the FM readout scheme to be ratiometric to an reference clock frequency, either by directly using the reference clock for frequency measurement, or via calibration. As will be shown later, the LFM gyroscope is capable of achieving sub-10 ppm scale factor stability; the FM demodulator must not degrade this.

¹Theoretically, the minimum split frequency is equal to the bandwidth. In practice, some margin must be allowed for frequency drift, filter roll-off, and to prevent aliasing of out-of-band signals.

3.3 FM demodulation methods

Slope and quadrature detection

The slope detector is the oldest and the simplest approach for FM demodulation. The basic idea is to use a tuned circuit (or a digital filter) to convert frequency variation into amplitude variation. A high-Q resonant circuit is commonly used to maximize the change in amplitude for a given frequency shift.

The quadrature detector is a variation of this concept, and is the most common analog FM demodulator. The FM signal is fed through a resonant network tuned to the center frequency. This produces a nominally 90° phase shift, which varies with frequency. The resulting phase-shifted signal is mixed with the original, producing an output approximately proportional to the input frequency deviation. For a given Q, this topology is more linear than the simple slope detector.

These approaches have numerous drawbacks. While they can provide adequate performance for radio receivers using wideband FM, the achievable resolution, scale factor accuracy, and linearity are limited by analog imperfections. Because the achievable filter slope is limited by component quality factors, slope detectors are poorly suited for systems with small frequency deviation, such as gyroscopes. Finally, these circuits are not especially suitable for monolithic integration due to the required high-Q components.

Analog phase-locked loops

A phase-locked loop is another option for FM demodulation. An analog PLL uses a feedback loop to keep a voltage-controlled oscillator in a defined phase relationship (generally either 0° or 90°) with the input signal. The VCO control voltage is thus proportional to the frequency deviation from a center frequency. Adjusting the VCO offset and gain allows the PLL to have either high frequency resolution or a wide input frequency range. However, the offset stability, scale factor accuracy, and linearity are set by the VCO voltage-to-frequency characteristic, and are subject to analog errors. In an analog implementation, it is difficult to make these parameters good enough for a gyroscope readout system.

Another disadvantage of the PLL is its limited frequency response. A common “rule of thumb” states that the loop bandwidth of a PLL can be no more than 10% of the input frequency. This implies that for a 30 kHz gyroscope, the maximum bandwidth is on the order of 3 kHz. While this is not a problem in itself, this does imply that the loop has




Figure 3.1: Arctangent-based FM demodulator.

significant phase shift even at a typical LFM split frequency of 200 Hz, which changes with component drift, input frequency, and split frequency. This presents difficulties in recovering the phase of the LFM demodulation reference, and can significantly limit achievable quadrature rejection, resulting in poor bias stability.

DSP-based demodulators

Any of the preceding techniques can be implemented in the digital domain. DSP-based implementations of FM demodulators have significant advantages. They can benefit greatly from the precise timing reference available in most digital systems, eliminating many error sources that limit the performance of analog implementations. A DSP-based PLL is a particularly attractive implementation. The frequency and scale factor accuracy depend only on the accuracy of the reference clock source. Such an architecture is employed in laboratory instruments such as the Zurich Instruments HF2LI lock-in amplifier [19].

To analyze the noise performance of this structure, it is convenient to use the demodulator shown in Figure 3.1a as a model. This is a commonly used structure for DSP-based FM demodulators [20]; it is also quite similar to the front-end of more complicated structures, such as the Costas loop² [21]. The FM signal is first quadrature downconverted to either an IF frequency, or to DC. An arctangent function converts the I and Q components to a phase angle, and a derivative operation converts the phase angle to an angular frequency. White noise added to the signal splits equally between amplitude and phase components. Figure 3.1b shows this graphically: the signal is a

²The Costas loop was originally used to recover a phase reference from a suppressed-carrier AM signal by employing the AM rejection property of this structure.

vector of length A , and the added phase noise consists a vector representing a zero-mean random variable with standard deviation $\sigma_n/\sqrt{2}$. Unlike oscillator phase noise, this added noise is not accumulated and disturbs the phase in a zero-mean fashion.

For a signal with amplitude A and white noise density σ_n , the phase disturbance due to noise is (assuming $\sigma_n \ll A$):

$$\sigma_\theta = \tan^{-1} \left(\frac{\sigma_n}{\sqrt{2}A} \right) \approx \frac{\sigma_n}{\sqrt{2}A} \quad (3.1)$$

After the derivative operation, the frequency noise density is given by

$$\sigma_f(f) = \frac{\sigma_\theta}{2\pi} 2\pi f = \frac{\sigma_n f}{\sqrt{2}A} \quad (3.2)$$

Because the FM demodulator effectively differentiates its input, white phase noise at the input of the FM demodulator generates f^2 noise at the output. To determine the required data converter SNR, the total frequency noise power in the signal bandwidth B must first be found:

$$P_n(\sigma_n) = \int_{f_c-B}^{f_c+B} \sigma_f(f)^2 df = \left[\frac{B^3}{3} + Bf_c^2 \right] \frac{\sigma_n^2}{A^2} \quad (3.3)$$

By equating this expression with the desired total frequency noise specification, the required σ_n can be determined:

$$P_n(\sigma_{n,req}) = \sigma_{f,req}^2 B \quad (3.4)$$

$$\sigma_{n,req}^2 = \frac{\sigma_{f,req}^2 A^2}{B^2/3 + f_c^2} \quad (3.5)$$

The required SNR is then given by the following expression. Note that the bandwidth of interest is twice the gyro bandwidth, since both AM sidebands must be considered.

$$\text{SNR} = 10 \log \frac{2A^2}{\sigma_{n,req}^2 2B} = 10 \log \frac{B^2/3 + f_c^2}{\sigma_{f,req}^2 B} \quad (3.6)$$

To achieve a 5 mdeg/s/ $\sqrt{\text{Hz}}$ (13.9 $\mu\text{Hz}/\sqrt{\text{Hz}}$) ARW spec with a gyroscope bandwidth of 60 Hz and a 100 Hz frequency split requires an in-band SNR of 120 dB. It is very difficult to meet such a specification with a sub-1mW power budget. This approach is thus unattractive for low-power, high-resolution gyroscope systems.




Figure 3.2: FM demodulation by period measurement. The topmost line is the modulating signal; the middle line is the frequency modulated sinusoid. The bottom line shows the signal after the limiter.

Period measurement techniques

FM can also be demodulated by directly measuring the period of the incoming signal. In most cases, the signal is first converted to a square wave with a continuous-time comparator (limiter); limiting the signal yields well-defined, fast edges that can be processed by digital circuits. The time interval between edges is then measured; this measurement is typically quantized to some time increment. This quantization noise is first-order shaped.

Consider a frequency-modulated signal $x(t)$ with a modulating signal $m(t)$:

$$x(t) = A \cos \left(2\pi f_0 t + 2\pi \int_0^t m(t) dt \right) \quad (3.7)$$

When $x(t)$ is processed by a comparator, the resulting square wave signal can be equivalently represented as a sequence of time points $\tau_0, \tau_1 \dots \tau_n$, where each such point represents a rising edge. This process is shown graphically in figure 3.2. Assuming both the comparator and the signal are free of noise and offsets, each such point τ_i can be found from the relation:

$$2\pi f_0 \tau_i + 2\pi \int_0^{\tau_i} m(t) dt = 2\pi i \quad (3.8)$$

Note that the left side is simply the argument (i.e. phase) of the cosine in (3.7). The comparator produces a rising edge whenever the phase is equal to multiple of 2π .

Assuming that $|m(t)| \ll f_0$, that the bandwidth of $m(t)$ is much smaller than f_0 , and

that $m(t)$ has no DC component³, equation (3.8) can be approximately solved for τ_i to yield the following expression:

$$\tau_i \approx \frac{i}{f_0} + \sum_{k=0}^i \frac{m(k/f_0)}{f_0^2} \quad (3.9)$$

This expression assumes the maximum frequency deviation $|m(t)|$ is small enough that the frequency is approximately constant and the effects of nonuniform sampling can be ignored. In gyroscopes, this is a very good approximation, since a full-scale 2500 deg/s input is only a 230 ppm frequency shift with a 30 kHz transducer.

In order to measure the period, consecutive time points are differenced. In practice, each measured τ_i is corrupted by added edge jitter $e[i]$, which comprises quantization noise and various sources of random noise. The resulting expression for the period is thus:

$$\tau_i - \tau_{i-1} = \frac{1}{f_0} \left[1 + \frac{m(i/f_0)}{f_0} \right] + e[i] - e[i-1] \quad (3.10)$$

To recover the original $m(t)$, this expression is multiplied by f_0^2 to transform the period measurement back to a frequency:

$$f_0^2 \cdot (\tau_i - \tau_{i-1}) = f_0 + m(i/f_0) + f_0^2 \cdot (e[i] - e[i-1]) \quad (3.11)$$

The edge jitter undergoes first-order noise shaping (differentiation). This may be understood intuitively by considering that frequency is the derivative of phase; edge jitter is essentially phase noise, so it is differentiated by an FM demodulator.

Frequency counter

The simplest period measurement architecture is a time-to-digital converter based on a counter. A digital counter fed with a reference clock continuously counts up; the input edge samples the value of the counter, and the previous such sample is subtracted. The resulting value represents the period of the input signal quantized to the period of the reference clock.

Despite its simplicity, this topology has several interesting features. There are no analog errors: the accuracy is set solely by the accuracy of the reference clock. Because of the noise shaping characteristic, a counter is a simple, accurate, and highly precise

³A DC component of $m(t)$ is simply a static frequency shift, and can be lumped into f_0 without loss of generality.

instrument for frequency measurement when a low sample rate is sufficient. However, achieving high resolution requires extremely high clock speeds.

The FM noise density due to quantization jitter with quantization step Δ can be found using (3.11):

$$\sigma_f^2(f) = f_0^4 \frac{\Delta^2}{12(f_0/2)} \left| 1 - e^{-j2\pi f/f_0} \right|^2 = \frac{2}{3} f_0^3 \Delta^2 \sin^2(\pi f/f_0) \quad (3.12)$$

Achieving 5 mdeg/s/ $\sqrt{\text{Hz}}$ in a 60 Hz bandwidth with a 100 Hz split requires $\Delta \leq 221$ ps, necessitating a clock frequency of at least 4.53 GHz. Operating at such a high clock frequency requires a large amount of power, likely well over 7 mW in a 0.18 μm process [22].

Time to digital converters

A time-to-digital converter (TDC) can have higher resolution than a counter for a given reference clock frequency. TDCs are commonly implemented with a delay-locked loop (DLL) architecture. A voltage-controlled delay line is locked to the period of a reference clock. The multiple clock phases are then used to measure the time between two events. The most basic TDC is a thermometer-like design: a number of flip-flops are used to measure the position of a “start” pulse within the delay line. This limits resolution to the delay of a single delay element. Another common TDC architecture is the vernier TDC, which uses two delay lines with a different delay step to achieve a resolution equal to the delay difference. The trade-off is a larger number of stages for a given delay range. Numerous other architectures can also be used, such as gated ring oscillators [23,24].

In order to measure the period of a 30 kHz gyro, a TDC must be combined with a counter to extend its measurement range. This presents a direct tradeoff between counter clock frequency and the required number of elements in the TDC. For example, to meet the previously calculated 221 ps resolution requirement with a 10 MHz counter frequency and a thermometer TDC would require 450 delay elements and associated decoding logic. This would consume significant area. A higher reference clock frequency would reduce the number of delay elements, at the expense of higher counter power.

Another problem is linearity. While the DLL architecture ensures that the total delay of the delay line is always constant, mismatch between individual elements results in nonlinearity. This necessitates complex calibration schemes and generally limits the achievable performance. As a result of these limitations, TDCs tend to have fairly high power consumption – generally on the order of 5 to 50 mW.

Sigma-delta frequency to digital converters

The sigma-delta frequency-to-digital converter (FDC) takes advantage of the inherent oversampling of the modulating signal present in most FM schemes. For a typical MEMS gyroscope with a resonant frequency of 30 kHz, the sensor bandwidth is generally between 20 Hz and 2 kHz, corresponding to an oversampling ratio between 15 and 1500. Oversampling is commonly used in sigma-delta ADCs to achieve very high resolution with a coarse (often 1-bit) converter. The $\Sigma\Delta$ FDC realizes the same basic idea in an FM demodulator, by employing noise shaping to improve the resolution of a counter with a relatively low reference clock frequency.

This architecture has numerous advantages over other FM demodulator architectures. The feedback path is entirely digital; therefore, the scale factor accuracy and offset stability of the demodulator depend only on the accuracy of the reference clock. The signal transfer function is simply a delay; unlike a conventional PLL, the signal is not subject to bandwidth limitation and phase distortion. Analog imperfections in the charge pump can affect noise shaping performance, but can not introduce frequency offset or DC gain errors. Noise added by the analog integrator undergoes second-order shaping, and its impact is negligible.

Because of its significant advantages, this is the architecture chosen for this work, and it is fully analyzed in the remainder of this chapter.

3.4 Comparator design

Many FM demodulation methods rely on comparators to convert a sinusoidal FM signal to a square wave. Such comparators are not clocked, and are sometimes referred to as limiters. Limiting an FM signal removes the amplitude modulation components and allows FM signals to be processed by digital logic gates. In a modern CMOS process, digital gates have delays on the order of picoseconds, consume very little power, and have very low phase noise when supplied with fast edges and a clean power supply rail. Even a minimum-size inverter has a negligible noise contribution, provided that the input rise time is sufficiently fast.

The front-end comparator performs the function of an LNA in an RF front-end: it provides sufficient gain to render the noise of any subsequent stages negligible. Without a comparator, significant power must be expended at every stage in the analog chain to maintain a sufficiently high signal-to-noise ratio. For example, the Zurich Instruments HF2LI requires an analog front-end with over 100 dB of dynamic range and a 14-bit, 210

MS/s analog-to-digital converter in order to achieve frequency resolution corresponding to about $0.005\text{mdeg/s}/\sqrt{\text{Hz}}$ [19].

A major advantage of comparator-based systems is that the gain is not restricted by the supply voltage. In a linear system, the gain of the front-end amplifiers must be kept low enough to keep the signal swing within the supply voltage rails. This restriction is not present in a comparator-based system. For example, a comparator producing a 30 ps rise time with a 30 kHz signal on a 1.8 V supply generates the same slope as a sinusoidal signal with an amplitude of 318 kV. For a given noise voltage, such a signal is far less susceptible to jitter from subsequent stages than a lower-amplitude one, allowing the use of relatively noisy, low-power, minimum-size logic gates for further processing.

Another advantage of comparators is that they can use dynamic operation to greatly reduce power for a given noise specification. A linear class A amplifier must use a constant bias current to achieve a certain level of noise. A well-designed comparator operates similarly to a class C amplifier, and needs current only during the time it is transitioning. Since comparators generally have a fast transition time, dynamic operation can reduce power by an order of magnitude compared to a linear amplifier designed for the same noise specification.

Unfortunately, comparators also have some disadvantages. A fundamental problem is noise folding. This is easily seen intuitively: a very fast, high-gain comparator only looks at the signal during a very brief time window; as a result, it is sensitive to even very brief noise impulses, which can cause it to transition either before or after the true zero crossing. Another way to understand noise folding is to view the operation of the comparator as a sampling process. The comparator produces sharp edges at the zero crossings of the input. If the rise time is very fast, this waveform can be losslessly represented by a discrete-time sequence of timestamps sampled at a constant phase interval π , if both the rising and the falling edges are considered. If the effects of nonuniform sampling are neglected, the sampling rate is $2f_0$, where f_0 is the oscillation frequency of the gyroscope. Thus, any noise components above $f_s/2 = f_0$ will be folded down into the signal band.

Noise folding is less pronounced in communications systems, where the channel SNR is low and tuned circuits reject out-of-band noise. However, it can be a significant limitation in high-resolution FM systems, and proper measures must be taken to prevent it.




Figure 3.3: Simplified equivalent comparator model.

Comparator model

Analyzing the noise of a comparator is a difficult non-linear problem. The circuit may be analyzed numerically, using a periodic steady-state (PSS) solver, such as SpectreRF. While this is indispensable for final design verification, accurate PSS analysis of a comparator requires an extremely large number of harmonics to be computed. For example, the signals inside a comparator with a 1 ns rise time will contain frequencies up to approximately 400 MHz. With a 30 kHz input signal, over 13,000 harmonics must be computed to obtain an accurate result, making the simulation very slow. Therefore, it is highly desirable to obtain a simple analytical model that can guide the design process in the correct direction and provide intuition.

The idealized model of a comparator is shown in Fig. 3.3. The comparator is modeled as a g_m - C integrator with ideal diodes for voltage limiting. For simplicity, a bipolar supply is defined; the comparator switches at the input zero crossing and has a positive gain. The input in the vicinity of the zero crossing is assumed to be a straight line, which is an accurate approximation for sinusoidal inputs.

It is important to note that the actual comparator circuit is implemented differently. Nearly all comparator topologies operate dynamically: the g_m device is only active during the time the output is transitioning, and does not consume static current. Because the output slope in a well-designed comparator is generally much higher than the input slope, and the transistor is in saturation during the first half of the transition, the g_m can be assumed to be roughly constant during the switching event.

The switching waveforms for this comparator model are shown in Fig. 3.4. The switching behavior starts when the input crosses zero and the comparator begins integrating the input, with the output initially starting at $-V_{dd}/2$. The switching process is considered complete when the output crosses zero. It is useful to define a propagation delay t_p as the time between the input and output zero crossings. The input and output




Figure 3.4: Switching waveforms for comparator model.

slopes at the point of the zero crossing will be denoted by α_i and α_o . The comparator gain-bandwidth product is $\omega_0 = g_m/C$.

With an input $x(t) = \alpha_i t$, the output is:

$$y(t) = \frac{V_{dd}}{2} - \frac{1}{C} \int_0^t g_m x(t) dt = -\frac{V_{dd}}{2} + \omega_0 \alpha_i \frac{t^2}{2}. \quad (3.13)$$

The propagation delay t_p can be found from the relation $y(t) = 0$:

$$t_p = \sqrt{\frac{V_{dd}}{\omega_0 \alpha_i}}. \quad (3.14)$$

By differentiating $y(t)$ at $t = t_p$, we obtain $\alpha_o = \sqrt{\omega_0 \alpha_i V_{dd}}$. By solving this expression for ω_0 , we obtain:

$$\omega_0 = \frac{\alpha_o^2}{\alpha_i V_{dd}}. \quad (3.15)$$

The preceding equation allows the required gain-bandwidth product to be found for a given output slope. In further analysis, the comparator is specified using only α_i , α_o , and V_{dd} , making the analysis easier to relate to simulated or measured behavior.

Noise analysis

The comparator operates as a periodic integrator. Because of the clipping diodes, the output saturates after the edge transition.⁴ Once the output is clipped, its value is independent of the preceding integrator state. Thus, the integrator state is effectively reset after each comparison. From [25], if a source of current noise with spectral density $S_i(f)$

⁴In the actual circuit, clipping occurs when the output approaches the supply rails. Diodes are used to simulate this behavior in an otherwise linear model.




Figure 3.5: Calculated and simulated comparator edge jitter vs. C with different levels of input noise. Dashed lines show gyroscope ARW specs at 100 Hz split. Black circles show PSS simulation results.

is repeatedly integrated on a capacitor C for time t_p , the power spectral density of the resulting voltage is given by:

$$S_C(f) = \frac{1}{C^2} \left[\frac{\sin(\pi f t_p)}{\pi f} \right]^2 S_i(f) \quad (3.16)$$

If $S_i(f)$ is white, the rms voltage noise is:

$$\langle v_n^2 \rangle = \int_0^\infty S_C(f) df = \frac{S_i t_p}{2C^2} \quad (3.17)$$

The jitter is then $\langle t_n^2 \rangle = \langle v_n^2 \rangle / \alpha_0^2 = S_i t_p / 2\alpha_0^2 C^2$.

There are two noise sources in this system. The first is the noise of the g_m block, which for a single transistor is $S_i = 4kT\gamma g_m$. In this case, the jitter is given by:

$$\langle t_{n,int}^2 \rangle = \frac{4kT\gamma g_m t_p}{2C^2 \alpha_0^2} = \frac{2kT\gamma}{C \alpha_i \alpha_0} \quad (3.18)$$




Figure 3.6: Calculated comparator rise time vs. edge jitter due to different levels of input noise. Increasing C corresponds to moving left along the X axis, reducing the comparator bandwidth. This results in decreased edge jitter, but slower rise time.

This equation predicts that noise is decreased when g_m is increased; this corresponds to a larger $\alpha_o C$ product. Noise also decreases when α_i is increased, which corresponds to a larger input signal.

The second noise source is the constant white noise at the input. One such noise source can be a preamp before the comparator. In this case, the jitter is given by:

$$\langle t_{n,ext}^2 \rangle = \frac{S_v g_m^2 t_p}{2C^2 \alpha_o^2} = \frac{S_v}{\alpha_i^2} \frac{\alpha_o}{2V_{dd}} \quad (3.19)$$

Note that S_v/α_i^2 is simply the input jitter PSD, in units of s^2/Hz . The second term can therefore be considered the effective bandwidth over which the comparator folds input noise. If $\alpha_o = V_{dd}/t_p$ is substituted, we find that $BW_{eff} = 1/2t_p$. This implies that noise folding is a serious problem for any comparator: a comparator with a 100 ns rise

time has a noise bandwidth of 5 MHz, corresponding to a 22 dB noise penalty with a 30 kHz input. Furthermore, this noise penalty depends only on the time window over which the comparator integrates the input signal. The noise penalty disappears if the comparator observes the entire half-cycle of the input (in which case no clipping occurs, and the circuit is equivalent to a linear g_m -C integrator). Figure 3.6 illustrates the trade-off between comparator speed and jitter: a smaller noise bandwidth results in a slower comparator, but reduces the effects of noise folding.

It is important to recognize the trade-off involving the choice of α_o . For convenience, we can define $\alpha_o = G\alpha_i$, where G is the slope gain. The total edge jitter is then:

$$\langle t_n^2 \rangle = \frac{2kT\gamma}{CG\alpha_i^2} + \frac{S_v G}{2\alpha_i V_{dd}} \quad (3.20)$$

The total jitter is minimized for $G = G_{opt}$, where:

$$G_{opt} = \sqrt{\frac{4kT\gamma V_{dd}}{S_v C \alpha_i}} \quad (3.21)$$

Substituting (3.21) into (3.20), we obtain:

$$\langle t_{n,opt}^2 \rangle = \sqrt{\frac{4kT\gamma S_v}{C V_{dd} \alpha_i^3}} \quad (3.22)$$

By substituting G_{opt} into (3.15), we obtain the optimal gain bandwidth product:

$$\omega_{0,opt} = \frac{4kT\gamma}{S_v C} \quad (3.23)$$

Note that the optimum corresponds to selecting a value of g_m that contributes the same amount of input-referred noise as S_v .

Clearly, it is essential to minimize wideband noise at the comparator input. This can be accomplished by band-limiting or filtering the oscillator output, or by connecting the comparator directly to the MEMS transducer. Nevertheless, some amount of broadband noise will always be present (e.g. from the gate resistance), and the appropriate gain must be chosen.

The preceding analysis found the total edge jitter. This noise is white, and the density may be found simply by dividing the rms value by the bandwidth $\sqrt{f_0}$ (assuming double edge sampling). However, to conclude the analysis, we must derive expressions for the FM noise at the output of the comparator. Using (3.11) we can write:

$$S_{FM}(f) = f_0^4 S_\tau(f) \left| 1 - e^{-j2\pi f/2f_0} \right|^2 = 4f_0^4 S_\tau(f) \sin^2(\pi f/2f_0) \quad (3.24)$$

S_τ is the jitter spectral density (in units of s^2/Hz).




Figure 3.7: Example comparator design.

Transistor-level analysis

Despite its simplicity, the previously described model can be used to accurately predict the noise of a transistor-level comparator circuit. In order to calculate the noise, it is necessary to find α_i , α_o , and C , which is most readily done with a transient simulation. α_i and α_o can be found by differentiating the input and output waveforms at the transition point. C can be accurately calculated by measuring α_o with and without a known load capacitance connected to the output, using the following equation:

$$C = \frac{C_{load}}{(\alpha_o / \alpha_{o,load})^2 - 1} \quad (3.25)$$

The noise may then be calculated using equations (3.20) and (3.24).

The comparator shown in figure 3.7 was analyzed using this procedure. From the transient simulation with and without a 1 pF load capacitor on each output, the following values were obtained:

$$\begin{aligned} f_o &= 33 \text{ kHz} \\ V_{dd} &= 1.8 \text{ V} \\ \alpha_i &= 1.70 \times 10^5 \text{ V/s} \\ \alpha_o &= 5.08 \times 10^6 \text{ V/s} \\ \alpha_{o,load} &= 3.35 \times 10^6 \text{ V/s} \\ C &= 530 \text{ fF} \end{aligned}$$

Equation (3.20) predicts an rms jitter density of $740 \text{ fs}/\sqrt{\text{Hz}}$ for these parameters; PSS simulation reports a jitter density of $912 \text{ fs}/\sqrt{\text{Hz}}$. Given the approximations made in




Figure 3.8: AM to PM conversion in the presence of offset.

the above analysis, this is a reasonable degree of agreement. As a further check, the input amplitude was reduced such that $\alpha_i = 6.89 \times 10^4$ V/s and $\alpha_o = 2.58 \times 10^6$ V/s. Equation (3.20) predicts a jitter density of 1.63 ps/ $\sqrt{\text{Hz}}$, while PSS simulation reports a jitter density of 2.04 ps/ $\sqrt{\text{Hz}}$. These results are likewise in reasonably close agreement.

Other errors

AM to PM conversion

Amplitude modulation can result in phase modulation when the comparator offset is non-zero. This process is graphically shown in Figure 3.8. The comparator transitions when the signal crosses the offset voltage V_{os} (instead of zero). Note that because V_{os} is generally small, the sine wave is approximately linear in the region of interest and has slope $\alpha = 2\pi f_0 A$, where A is the amplitude and f_0 is the input frequency. Thus, the time delay between the true zero crossing and the comparator transition is a function of amplitude.

Because the process of frequency demodulation differences consecutive time samples, a static delay does not affect the measured frequency. However, a time-varying delay will appear as FM.

For example, consider the following amplitude-modulated signal:

$$x(t) = (1 + k_{am} \cos(\omega_{am}t)) \cos(2\pi f_0 t) \quad (3.26)$$

If this signal is processed by a comparator with an offset V_{os} , the delay variation will be given by

$$\tau_D = \frac{V_{os}k_{am}}{\pi f_0 A} \cos(\omega_{am}t) \quad (3.27)$$

where A is the amplitude and f_0 is the input frequency. The differentiating action of the FM demodulator applies the transfer function $|H(j\omega)| = 2 \sin(\omega/2f_0) \approx \omega/f_0$. At the output of the FM modulator, the AM component is therefore

$$y_{am}(t) = \frac{2V_{os}k_{am}f_{am}}{A} \sin(\omega_{am}t) \quad (3.28)$$

AM to FM conversion can be a significant problem in the LFM gyroscope. If the amplitude regulator loop does not have sufficient gain, applied angular rate (or quadrature) can create amplitude ripple at the split frequency. For example, suppose the comparator has an offset $V_{os} = 5$ mV, the amplitude control loop allows 1% amplitude ripple ($k_{am} = 0.01$), the split frequency is $f_{am} = 100$ Hz, and the oscillation amplitude is $A = 0.5$ V. Due to the offset, the AM ripple becomes an FM signal with a deviation amplitude of 0.02 Hz (or 7.2 deg/s) at the split frequency, thus adding a significant offset to the rate output and degrading performance.

Offset drift to FM conversion

From Figure 3.8, it is clear that V_{os} variation can also change the delay between the true zero crossing and the comparator edge. This delay is given by:

$$\tau_D = \frac{V_{os}}{\alpha} = \frac{V_{os}}{2\pi f_0 A} \quad (3.29)$$

Just like in the AM case, the FM demodulator differentiates this delay. If V_{os} has a power spectral density $S_{os}(f)$, the output of the FM demodulator will have a spectrum given by:

$$|S_{FM}(f)| = |S_{os}(f)| \left(2 \sin \frac{\pi f}{f_0} \right)^2 \quad (3.30)$$

Double edge sampling

Sampling both rising and falling edges eliminates both of these error sources. Consider a sine wave processed by a comparator with no offset: the output sequence is $\tau_{r1}, \tau_{f2}, \tau_{r3}, \tau_{f4}, \dots$, corresponding to a sequence of rising and falling edges. Because the rising and falling zero crossings have opposite slopes, an offset or amplitude change

delays the rising and falling edges in opposite directions. For example, a positive offset will delay rising edges and advance falling edges with respect to their nominal positions. If τ_{Dn} is the delay of the n -th edge, the actual edge sequence will be $\tau_{r1} + \tau_{D1}, \tau_{f2} - \tau_{D2}, \tau_{r3} + \tau_{D3}, \dots$

When the FM demodulator differences consecutive edge times, the following sequence is produced:

$$\begin{aligned}\Delta\tau_1 &= (\tau_{f2} - \tau_{D2}) - (\tau_{r1} + \tau_{D1}) = (\tau_{f2} - \tau_{r1}) - (\tau_{D1} + \tau_{D2}) \\ \Delta\tau_2 &= (\tau_{r3} + \tau_{D3}) - (\tau_{f2} - \tau_{D2}) = (\tau_{r3} - \tau_{f2}) + (\tau_{D3} + \tau_{D2}) \\ \Delta\tau_3 &= (\tau_{f4} - \tau_{D4}) - (\tau_{r3} + \tau_{D3}) = (\tau_{f4} - \tau_{r3}) - (\tau_{D4} + \tau_{D3}) \\ &\dots\end{aligned}$$

Now, the added time errors are not differenced, but are instead added together; the resulting error is then modulated by a $+1, -1, +1, \dots$ sequence, modulating the error to the Nyquist frequency $f_s/2 = f_0$.

Sampling both edges effectively chops the errors due to offset drift and amplitude modulation, and moves them out of band. This comes at the cost of greatly increasing the dynamic range requirement for the FM demodulator, since it must now be capable of handling $\pm 2\tau_D$ in addition to the FM signal, even for static offsets.

An implicit assumption made in this analysis is that the offset is the same for the rising and the falling edges. This is not the case for most pseudo-differential comparator topologies, where different transistors are used to detect rising and falling edges. Therefore, the pseudo-differential comparator must be designed for sufficiently low flicker noise either via appropriate sizing or by employing an auto-zero scheme. However, double edge sampling is still useful for rejecting errors arising from other offsets and AM.

Circuit design

The theory developed in the previous section can be readily applied to the practical design of high-performance, low-power comparator circuits. The main trade-off that must be considered in the design process is the trade-off between comparator speed, comparator size, and noise folding. The goal of the comparator signal chain is to produce very fast edges at the output even with a small, low-amplitude input signal while adding minimum noise and consuming the minimum possible energy.

A comparator acts as an extremely high-gain amplifier. For example, a sinusoidal signal with an amplitude of 100 mV and a frequency of 30 kHz has a slope of about 19

kV/s at the zero crossing. The rise time of a minimum-size inverter is on the order of 50 ps; with a 1.8V supply, the slope is 36 GV/s. This corresponds to a gain of about 126 dB. Multiple stages are necessary to achieve this.

Preamp

In most systems, the comparator must be preceded by a linear preamp. This block is necessary for a number of reasons. Because the switching point of the cross-coupled comparator is generally tied to the threshold of the input devices, it is necessary to amplify the signal and control the common-mode voltage in order to center the signal around the switching point. Furthermore, the signal must be sufficiently large to ensure rapid regeneration to avoid large shoot-through currents, and to overcome any offset and hysteresis of the comparator. The preamp can greatly reduce kickback from the comparator, preventing the injection of disturbances into the oscillator loop.

Perhaps the most important reason to use a preamp is to provide common-mode rejection for the comparator. This may appear to be unnecessary, since the cross-coupled comparator has a differential input. However, the cross-coupled comparator essentially operates as two independent inverters: on any given comparison cycle, only one side is performing the comparison, and the cross-coupling is only needed for efficient regeneration. Thus, despite having a differential input, it does not provide common-mode rejection. A linear preamp avoids this problem.

The design requirements for the preamp are somewhat peculiar. The preamp output noise spectrum must not extend past the input frequency in order to avoid noise folding in the subsequent stages. This implies an open-loop integrator topology. At the same time, it is not necessary or desirable to have a high DC gain. Ideally, the integrator pole should be located slightly below the input frequency to minimize noise folding and avoid excessive low-frequency gain.

The output amplitude of the amplifier must be accurately controlled in order to maximize swing while avoiding clipping. The latter requirement may seem unnecessary, since the subsequent comparator stage does not operate near the peaks of the signal. However, allowing the integrator to clip would lead to its state being periodically reset, which would result in noise folding. This may be easily avoided by employing an automatic gain control circuit similar to that used in Pierce oscillators [14]; such a circuit adjusts the current through the differential pair to control the amplifier gain and ensure a constant output amplitude over a range of input amplitudes. This has the additional benefit of avoiding the use of feedback in the amplifier itself, thus minimizing output

noise and reducing the required amplifier power.

First comparator stage

The first stage of the comparator must operate on the full-swing linear signal from the preamp. Together with the preamp, this is the most critical section of the comparator signal chain because it contributes much of the noise. Unlike the preamp flicker noise, flicker noise from this stage is not removed by double-edge sampling. Thus, very low- V^* devices must be used, limiting the maximum slope gain of the comparator stage.

The design procedure must consider the trade-off between the gain of the stage, the amount of noise the stage contributes, and the power efficiency of the stage. Very large stages have a small gain and consume significant power, but add less noise than smaller, higher-gain stages.

The power consumption for a properly designed dynamic comparator stage is primarily comprised of the energy required to recharge the capacitors during each input clock cycle. This value is quite low. For example, if the total comparator capacitance is 100 pF, the power required for 30 kHz operation at 1.8 V is about 10 μ W.

If a cross-coupled transistor pair is used for regeneration (fig. 3.7), additional parasitic power will be necessary due to shoot-through current from the regeneration devices. This incurs an unfortunate trade-off: making these devices weaker reduces crowbar current in the comparator stage, but also increases the time required to reach a valid output, which creates crowbar current in the following stage. Possible alternatives to cross-coupled regeneration are described later.

Following comparator stages

The pseudo-differential output waveforms of a properly designed comparator stage are highly asymmetric. For an NMOS comparator, these waveforms will have a fast and low-noise falling edge, and a slow and noisy rising edge. Only the falling edge should be used by the subsequent stage, implying that the next stage should use PMOS input devices (which will, in turn, generate a clean rising edge). Thus, the stage polarities must alternate.

Because each subsequent stage operates with a higher input and output slope, it contributes less noise and must be sized to be smaller and faster than the preceding one. The comparator thus tapers down towards the back-end as the edges become faster.




Figure 3.9: Comparator implementation.

Edge detector and SR latch

The final comparator produces two pseudo-differential outputs, each of which has only one usable low-noise edge. The low-noise edge may be either rising or falling, depending on the polarity of the last comparator stage. These must be combined into one signal where both edges have equally low noise. This may be done with a combination of two edge detectors and an SR latch (fig. 3.9b). Upon receipt of an edge of the appropriate polarity, an edge detector generates either a “set” or “reset” pulse, toggling the output of the SR latch. The edge detectors are necessary because the SR latch would otherwise use the “dirty” regeneration edge to generate the output transitions.

Regeneration circuits

Proper design of the regeneration circuits is critical for minimizing crowbar current and thus controlling power consumption. As described previously, the conventional cross-coupled regeneration circuit has significant parasitic power consumption, particularly when the preceding comparator stage is relatively slow.

Ideally, the regeneration logic should wait until one of the input devices is completely turned off before regenerating it, thus avoiding crowbar current. This may appear to be recursive, since it is necessary to build a comparator to derive the digital signals required by this logic. However, this signal does not need to have low jitter, since it is not used to produce the final output. Thus, a minimum-size skewed inverter can be used to gauge when the input voltage has decreased sufficiently to start regeneration; this topology is shown in figure 3.9a. The inverter transition point needs to be set to below the threshold voltage of the input devices to avoid shoot-through current.

Later stages that operate with relatively fast input edges do not need any special

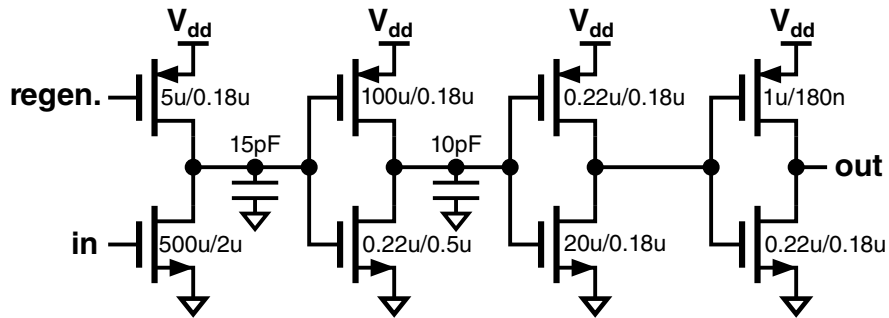


Figure 3.10: Example comparator circuit schematic.




Figure 3.11: Example comparator switching waveforms.

regeneration logic; it is sufficient to simply make the regeneration devices small enough to avoid significant crowbar current, making the structure identical to a skewed inverter.




Figure 3.12: Simulated comparator ARW vs. split frequency at the output of each stage.

Example circuit

To illustrate the operation of the cascaded dynamic comparator, an example circuit was designed and simulated. The schematic is shown in figure 3.10. The comparator consists of four stages, and is designed to meet a $5\text{mdeg/s}/\sqrt{\text{Hz}}$ ARW spec at a frequency split of 300Hz. The regeneration circuit is not shown.

The switching waveforms from a transient simulation are shown in figure 3.11. Note that the first stage switches very slowly, and each subsequent stage switches much faster.

The noise (from Pnoise simulation) is shown in fig. 3.12; it is expressed as gyroscope rate noise, in units of $\text{mdeg/s}/\sqrt{\text{Hz}}$. Note that the first stage dominates the noise, and subsequent stages add very little noise. Flicker noise is the dominant noise source, and the first two stages are sized quite large to minimize it. An autozero technique could potentially be employed to reduce the area and input capacitance.

This circuit is almost completely dynamic, and the power consumed by the half-circuit shown is $1.43\mu\text{A}$ at 1.8V. Two such half-circuits are needed for double-edge sampling.

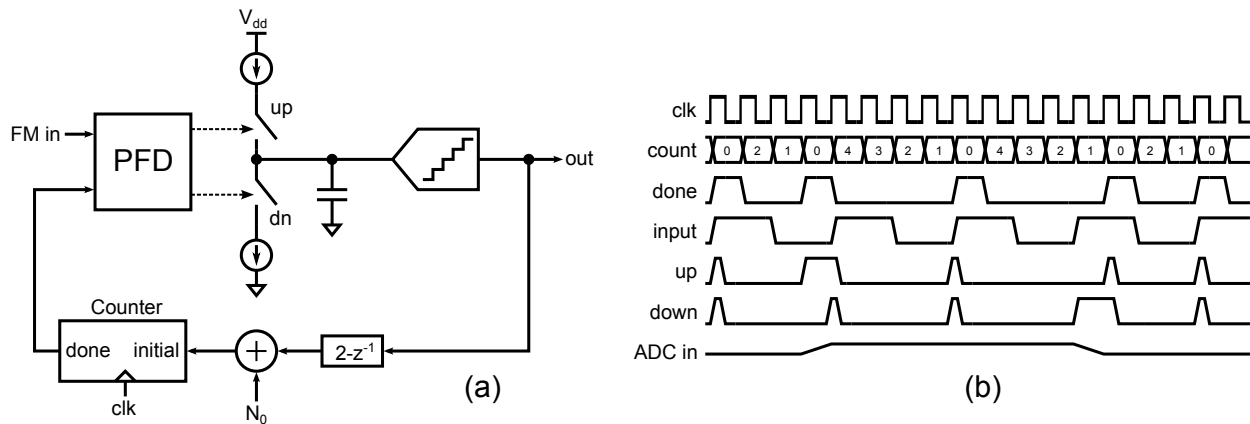


Figure 3.13: $\Sigma\Delta$ FDC: (a) conceptual block diagram; (b) timing diagram.

3.5 $\Sigma\Delta$ FDC Architecture

Fig. 3.13(a) shows the conceptual block diagram of a second-order $\Sigma\Delta$ FDC. The input is a frequency-modulated square wave with center frequency f_0 . A digital counter generates time intervals quantized to the period of the reference clock T_{ref} . The front-end consists of a phase and frequency detector (PFD) and a charge pump, which measure the time error between the input edges and the quantized feedback edges, and output a charge proportional to this error. This charge is then accumulated by an integrator. An ADC quantizes the output of the integrator; the digital code is then processed by a digital compensator block and fed back to the counter with a center frequency offset (N_0). The digital output of the FDC represents the message signal in units of T_{ref} with the addition of shaped quantization noise and with an offset $-N_0$ (representing the subtracted center frequency).

Fig. 3.13(b) shows the timing diagram for the conceptual $\Sigma\Delta$ FDC. The counter is decremented during every clock cycle. The “done” signal is produced when the counter reaches zero. The phase and frequency detector compares this signal to the input and switches the current sources appropriately, slewing the output of the integrator. The quantization error due to the counter is integrated by the capacitor at the output of the charge pump.

Fig. 3.15 shows the linearized equivalent model of the $\Sigma\Delta$ FDC. As described in the preceding chapter, the FM modulator acts as a discrete-time integrator in the time domain, producing a sequence of time points that form a monotonically increasing ramp. In a similar way, the counter also integrates its input. For example, if an initial “done” edge occurs at $t_0 = 0$ and the value N_1 is loaded into the counter, the next edge will

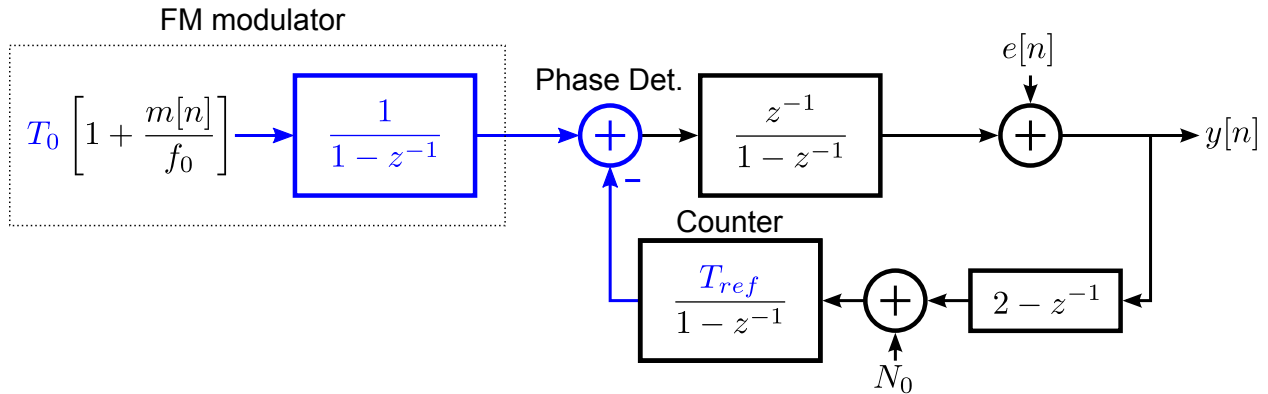


Figure 3.14: Linearized equivalent model of $\Sigma\Delta$ FDC. Time signals are highlighted in blue.

occur at $t_1 = T_{ref}N_1$; if N_2 is then loaded into the counter, the following edge will occur at $t_2 = T_{ref}(N_1 + N_2)$; and so on.

In order to analyze this structure, we must first remove the N_0 addition from the feedback loop. This may be done simply by replacing this addition with an offset $-T_{ref}N_0$ added to the FM modulator input. The addition of N_0 in the feedback path removes most of the center frequency offset from the FM signal, allowing the phase detector to process only the deviation of the FM signal from its center frequency.

The forward transfer function (between the phase detector and the output) is $A = (1/T_{ref})z^{-1}/(1 - z^{-1})$. The charge pump must have a gain $1/T_{ref}$ to undo the corresponding scaling by the counter. The feedback transfer function is $\beta = T_{ref}(2 - z^{-1})/(1 - z^{-1})$. The closed-loop transfer function is then $(1 - z^{-1})z^{-1}/T_{ref}$. The FM modulator integrates the signal, pre-multiplying it by $1/(1 - z^{-1})$; the signal transfer function (from the input of the FM integrator) is thus simply z^{-1}/T_{ref} . By a similar process, it can be shown that the noise transfer function (from $e[n]$ to $y[n]$) is $(1 - z^{-1})^2$. This structure thus performs second-order noise shaping.

Fig. 3.15 shows how the $\Sigma\Delta$ FDC structure is synthesized from a conventional second-order modulator [26]. The first integrator is replaced with a pair of integrators before the subtraction point, which represent the counter and the FM modulator. The subtraction point is replaced with a phase detector. Next, the second feedback path is removed and replaced with a $2 - z^{-1}$ compensator; this avoids the necessity of building a voltage-mode DAC and alleviates concerns due to gain mismatches between the two feedback paths. The penalty is increased dynamic range at the phase detector input.




Figure 3.15: Synthesis of $\Sigma\Delta$ FDC structure from second-order $\Sigma\Delta$ modulator. Time signals are highlighted in blue.

3.6 Double edge sampling

Sampling both the rising and the falling edges of the input signal doubles the oversampling ratio, significantly reducing quantization noise. Furthermore, sampling both edges effectively chops the errors due to oscillator and comparator offset drift, flicker noise, and AM sensitivity, and moves them out of band.

To implement this feature in the $\Sigma\Delta$ FDC, it is necessary to trigger the charge pump comparison on both rising and falling edges; this can be performed by inverting the signal after every comparison. Since most comparators have a pseudo-differential output, the inversion operation can be performed by a butterfly switch, which also minimizes errors due to any delay asymmetry between rising and falling edges at the output of the comparator. In the case of a single-ended input, an exclusive-OR gate can be used instead of the butterfly switch.

Figure 3.17 compares the AM rejection of the double-edge FDC to the single-edge

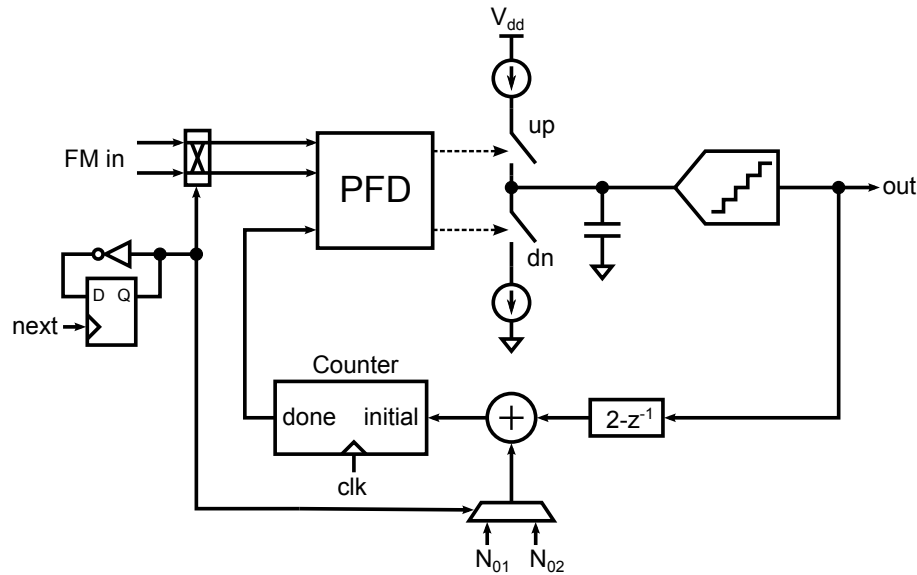


Figure 3.16: $\Sigma\Delta$ FDC configuration for double edge sampling.

FDC. The 100 Hz tone appears near $f_s/2$ for the double-edge FDC, but aliases directly into the signal band in the single-edge version. The double-edge FDC also has significantly lower quantization noise due to the higher oversampling ratio.

Double edge sampling requires increased dynamic range to handle offsets and amplitude modulation, since these phenomena are converted to amplitude modulation at the Nyquist frequency $f_s/2 = f_0$. Even a small duty cycle error produces a signal much larger than any expected input due to rate: for example, with a 30 kHz gyro, a 1% duty cycle error will result in a $f_s/2$ component with a 150 Hz amplitude.

With the $\Sigma\Delta$ FDC, this problem can be avoided by using feedforward cancellation. Since offset drift and amplitude modulation are generally slow phenomena (relative to the sampling rate) and since there is no bandwidth restriction in the $\Sigma\Delta$ FDC loop, the N_0 input can be used to effectively subtract the expected contribution of these components from the $\Sigma\Delta$ FDC input.

Such a scheme is shown in figure 3.16 where two separate N_0 values are used for the positive and negative half-cycles. These values can be static (if offset drift is negligible), but a more practical approach is to dynamically adjust them during operation to minimize the signal at the output. Because the N_0 value is effectively subtracted from the input to the FM modulator, and the signal transfer function is simply a delay, no artifacts occur when the N_0 value is changed. Such a scheme has the additional advantages of expanding the dynamic range of the PLL for slowly varying signals, and automatically




Figure 3.17: AM rejection via double edge sampling.

tracking any shifts in the resonant frequency of the transducer.

3.7 High-level design

The primary design choices that must be made at the beginning of the design process are the loop order and the reference clock frequency. These two parameters determine the loop's quantization noise for a given oversampling ratio, and specify the design parameters for the rest of the loop circuits (such as the charge pump). The choice of loop order and reference frequency is driven by a combination of multiple co-dependent factors, such as the ratio between analog and digital power, the process used, the clock frequencies available (and the power required to generate them), and others. These trade-offs will be examined in the following sections.

Quantization noise

Noise in the $\Sigma\Delta$ FDC consists of quantization noise and circuit noise. Unlike a voltage-type $\Sigma\Delta$ ADC, the frequency quantization step size of the $\Sigma\Delta$ FDC is determined by the

clock frequency of the feedback counter and is equal to $\Delta = 2f_0^2 / f_{ref}$; the factor of two occurs because the double edge loop measures the period of a half-cycle. The quantization noise is shaped by the loop; in the case of the second-order loop, the NTF is $(1 - z^{-1})^2$, and the noise spectrum is given by:

$$N(f) = 4\sqrt{\frac{\Delta^2}{12f_0}} \sin^2\left(\frac{\pi f}{2f_0}\right) \quad (3.31)$$

In the more general case of an n -th order sigma-delta loop having $NTF(z) = (1 - z^{-1})^n$, the quantization noise spectrum is given by

$$N(f) = 2^{1+n/2} \sqrt{\frac{\Delta^2}{12f_0}} \sin^n\left(\frac{\pi f}{2f_0}\right) \quad (3.32)$$

In most $\Sigma\Delta$ converters, it is desirable to make the quantization noise floor significantly smaller than the circuit noise floor. However, in the $\Sigma\Delta$ FDC, reducing the quantization noise floor requires increasing the clock frequency, which increases power consumption in the counter and in the clock generator and clock distribution path. It is therefore more power-efficient to set the noise contributions from quantization and circuit noise to be approximately equal. In this case, it is important to verify that the modulator does not have strong tonal behavior.

Circuit noise

The circuit noise is composed of jitter from the counter and PFD, as well as thermal and flicker noise from the charge pump, integrator, and ADC. Generally, the digital jitter is a negligible source of noise, and the charge pump noise dominates. The charge pump noise undergoes first-order noise shaping, while the output of the n -th integrator undergoes $(n + 1)$ -order shaping. Integrator voltage noise generally does not significantly contribute to the overall noise floor in a second-order loop. However, the first integrator stage may be a significant limiting factor in high-order loops, particularly if the flicker noise is significant.

The ADC adds both quantization noise and voltage noise. Generally, the voltage noise is much smaller than the quantization noise in a coarse ADC; however, this may be an issue in a high-order loop with a relatively high-resolution ADC.




Figure 3.18: ARW due to quantization noise as a function of reference clock frequency for 2nd order loop, $f_{split} = 200$ Hz, $BW = 100$ Hz.

Loop order

The counter in the feedback path of the $\Sigma\Delta$ FDC behaves like a nearly-ideal DAC with a large number of bits, allowing stable high-order multi-bit loops to be implemented without encountering DAC linearity issues. For example, Figure 3.19 shows a fourth-order structure.

While increased loop order is a seemingly attractive option for reducing the required reference clock frequency and thus apparently reducing power, reducing the clock frequency imposes much more stringent requirements on the phase detector, charge pump, and integrators, which must handle a much larger dynamic range while maintaining the same timing resolution. With a reduced clock frequency, the loop relies increasingly on the charge pump and integrators to convert time to voltage. Since the charge pump swing is generally limited by the supply voltage, and since the input range is now larger, this change requires a reduction in charge pump and integrator noise, requiring a quadratic increase in power for these components for a linear decrease in clock

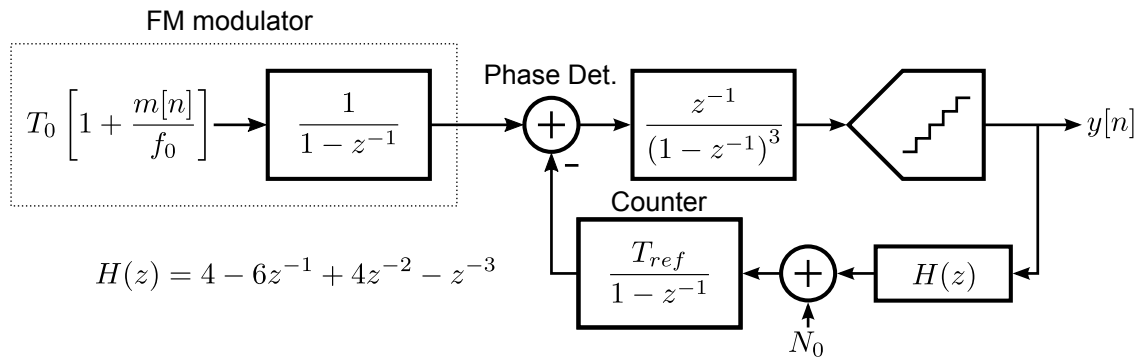


Figure 3.19: Example fourth-order $\Sigma\Delta$ FDC structure with combined feedback path.

frequency. This is a very unfavorable trade-off, since both digital power and clock PLL power generally scale linearly with clock frequency.

Furthermore, if the distributed feedback paths are merged into a single path through the counter, the first integrator sees larger and larger swing with increasing loop order due to the increasing gain of the compensator block. Figure 3.20 shows a simulated histogram of the phase detector error with a small input. While the second-order loop requires a charge pump range of ± 3 LSBs, a fourth-order loop increases this requirement to ± 56 LSBs. The actual increase in the required dynamic range is even larger, since a high-order loop would use a reduced counter frequency and thus would have coarser LSB steps. If the distributed feedback paths are not combined, a linear multi-bit DAC would need to be provided, and performance would be degraded by any mismatch between the counter and DAC gains.

The appropriate choice of loop order is thus dictated by the relative power consumption of the analog and digital blocks. If power is dominated by the clock and counter logic, increasing loop order will reduce clock frequency and digital power. However, the linear reduction in digital power comes at the expense of an increase in dynamic range, necessitating a quadratic increase in analog power to maintain the same noise floor. A reasonable design approach is to increase loop order only if the noise requirements cannot be met without increasing clock frequency past the limits of the process.

3.8 Circuit implementation

A second-order loop was implemented in a $0.18 \mu\text{m}$ CMOS process. The design goal was to achieve approximately $5 \text{ mdeg/s}/\sqrt{\text{Hz}}$ in a 100 Hz bandwidth at a 200 Hz split frequency. The noise budget was split approximately equally between quantization noise

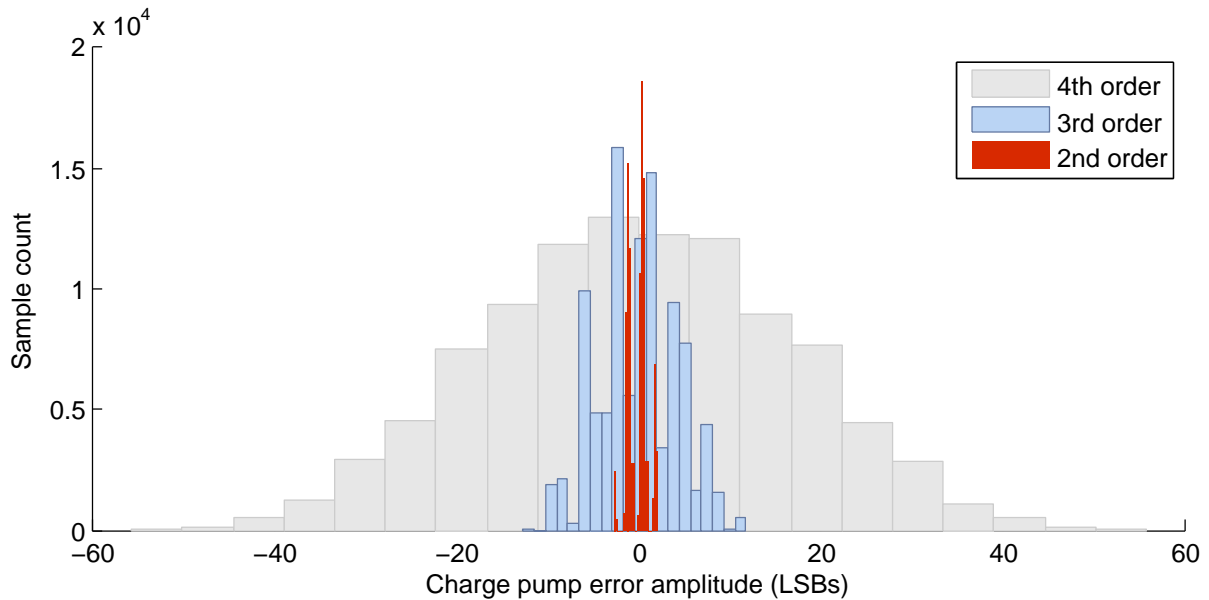


Figure 3.20: Simulated phase detector error histogram for 2nd, 3rd, and 4th order loops with combined feedback.

and circuit noise. To achieve this, $f_{ref} = 256\text{MHz}$ was chosen. As will be shown, the second-order loop presents a good trade-off between analog and digital power consumption.

PFD, charge pump, and active integrator

The analog front-end of the $\Sigma\Delta\text{FDC}$ consists of the charge pump and integrator. The charge pump constitutes a time interval to charge converter and the integrator accumulates charge to produce a voltage proportional to the integrated time error. These blocks are critical to the performance of the converter; their noise and nonlinearity determine the ultimate resolution of the converter.

While the charge pump structure is similar to that used in conventional PLLs, it has different requirements. In a PLL, the charge pump error output is nearly zero when the loop is in a locked condition, so linearity over a wide input range is not required; non-linearity can affect loop stability, but is not critical otherwise. In the $\Sigma\Delta\text{FDC}$, the error signal stays relatively large (± 2 LSBs peak-to-peak in the second-order loop) even in steady state operation, which makes linearity a significant issue. Another issue relates to dynamic behavior. In conventional PLLs, good matching between the up and down current sources is critical to minimize reference spurs due to periodic error impulses.

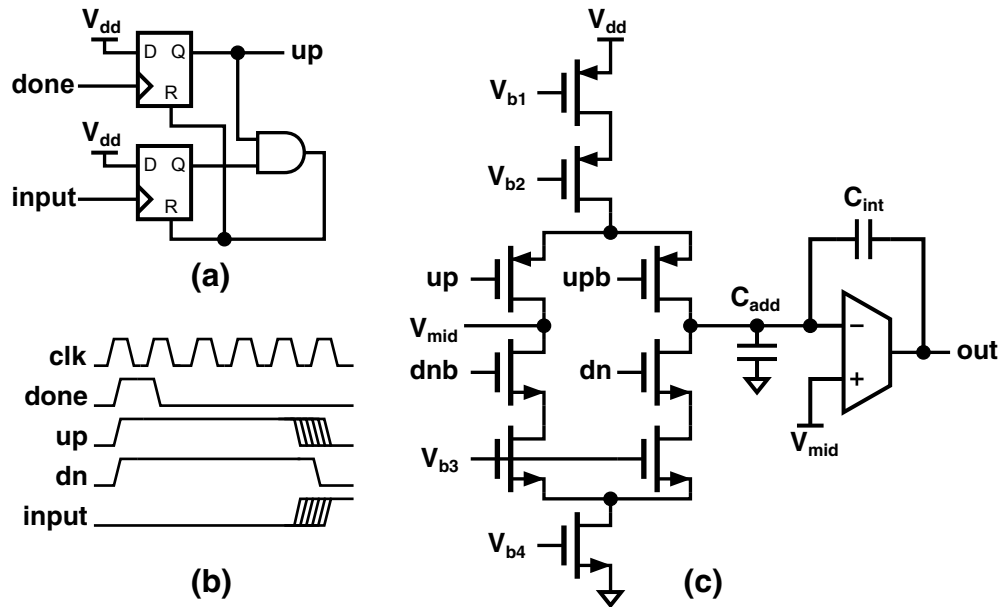


Figure 3.21: PFD and charge pump circuits: (a) PFD schematic, (b) timing diagram, (c) charge pump and integrator schematic.

The $\Sigma\Delta$ FDC is a discrete-time system, so dynamic errors do not degrade performance, as long as the integrator output is correct when it is sampled by the ADC.

The design shown in figure 3.21 takes into account the specific requirements. The PFD operates in a somewhat unconventional way, similar to that described in [27]. When the “done” signal from the counter is generated, both the “up” and the “down” switches are enabled. The “down” current source is always switched on for a fixed duration. The “up” switch is disabled by the input edge. This process is shown graphically in figure 3.21b. This operating mode has the advantage of having the same gain for both positive and negative phase errors, even if the “up” and “down” current sources are mismatched.

The charge pump (fig. 3.21c) uses cascoded current sources for linearity. An active integrator restores the output to the reference voltage before each comparison. This reduces nonlinearity due to voltage-dependent charge injection and charge redistribution from parasitic capacitances. Because the active integrator cannot be made sufficiently fast to maintain a virtual ground at its input during charge pump slewing, an additional capacitor C_{add} is used to reduce voltage swing at the output of the charge pump to maintain linearity. The active integrator transfers charge to the integrating capacitor C_{int} after each comparison operation, amplifying and buffering the signal for the ADC. The integrator is never explicitly reset and does not have local DC feedback; this function is




Figure 3.22: Counter topologies

performed by the outer sigma-delta feedback loop.

The flicker noise of the charge pump dominates the noise performance at low frequencies. In the LFM case, it is sufficient to set the flicker noise corner to the minimum expected split frequency, since no benefit is obtained from lower noise near DC. This was accomplished simply by making the charge pump current source transistors sufficiently large. Note that in this case, the biasing devices cannot be scaled down; this was addressed by power-gating the bias network when the charge pump is not in use. In this case, the charge pump duty cycle is on the order of 15%, reducing the static power by a factor of 6.

In future work, if it is necessary to reduce FDC noise at low frequencies (as would be the case for the QFM operating mode), a chopping or auto-zero arrangement could be employed to reduce the charge pump flicker noise. A possible implementation of charge pump chopping consists of a differential integrator and charge pump, and butterfly switches to periodically swap the "up" and "down" current sources (for example, this could be done on every rising input edge). This would move the error away from the baseband and up to a higher frequency. This approach has the disadvantage that the common-mode voltage of the integrator is difficult to control: the FDC loop provides differential-mode DC feedback, but the common-mode voltage must be controlled by another means. Another possible scheme could calibrate each current source to match a stable reference immediately before each comparison. This has the advantage of not requiring common-mode control for the differential integrator.

Dual Stage Counter

To accommodate gyros with a range of resonant frequencies between 1 kHz and 50 kHz, an 18-bit counter is required. Such a counter could be implemented as a synchronous counter, as shown in fig. 3.22a. However, this topology has very high power consumption: all 18 flip-flops are clocked at the full 250MHz rate, and numerous additional transitions occur in the adder on every cycle as the input signals propagate.




Figure 3.23: Slow counter schematic. The slow counter generates a “done” pulse when the count reaches zero. It also generates clock enable pulses used by the control state machine to time long-duration events, such as charge pump and amplifier settling.

An asynchronous counter, also known as a ripple counter, is another interesting option. This topology is shown in figure 3.22b. This counter is essentially a cascade of divide-by-two stages; an edge on the last flip-flop indicates when the count has reached zero. Because each subsequent flip-flop runs at half of the clock rate of the preceding one, the power consumption asymptotically approaches that of two flip-flops as the number of stages is increased. Because there is no combinational logic, this is a very power-efficient topology.

Unfortunately, it is not possible to directly use an 18-stage asynchronous counter at the 250 MHz input frequency. Because each stage uses the output of the previous stage as a clock, the propagation delay of the flip-flops accumulates. For the flip-flops used in this work, the nominal CLK to Q delay is about 280ps. For an 18-stage counter, the output edge will be delayed by about 5ns relative to the input edge; this skew exceeds one clock cycle. Furthermore, the exact delay is dependent on process, supply voltage, and temperature, and cannot be reliably compensated for. The need to initialize the counter to a particular starting value presents additional timing complications.

The solution chosen in this work is a dual-stage counter. The 18-bit counter is partitioned into a 5-bit “fast counter,” which runs at the full 256 MHz, and a 13-bit “slow counter” that runs at 8 MHz. The power-efficient slow counter is used to generate long delays with a coarse quantization step; the fast counter is enabled at the end of each cycle to generate the final high-resolution delay. The 5-bit counter has a propagation delay of 1.4 ns, which provides sufficient timing margin for the output flip-flop.

Figure 3.23 shows the schematic of the slow counter. The slow counter is clocked continuously. The slow counter also generates a `clk16` signal that acts as a $8\text{MHz}/16 = 500\text{kHz}$ timing signal for the control state machine, where it is used to generate delays




Figure 3.24: Fast counter schematic. The counter generates a “set” pulse that turns on the charge pump “up” current source, and a “down” signal that controls the “down” current source. The “down” signal has a duration of 5 clock periods. Clock gating is used to minimize power. The gating signal also controls the asynchronous set/reset inputs.

for charge pump and amplifier settling. Including layout parasitics, the slow counter consumes a current of $1.91\mu\text{A}$, comprising nearly all of the digital power consumed by the FDC circuit.

Figure 3.24 shows the schematic of the fast counter. The fast counter employs clock gating to minimize power consumption. The clock is enabled only while the counter is running, and the maximum duty cycle is about 0.1%; this makes the power consumption negligible. The clock is enabled by the rising edge of the start signal, and is disabled by the falling edge of the down signal. The fast counter generates two outputs. set is a single-cycle pulse produced when the count reaches zero and used to enable the charge pump “up” current source, which is subsequently turned off by the input edge). down enables the “down” current source for a duration of 5 clock periods; it starts simultaneously with set. The initial value is asynchronously loaded into the counter using the set/reset inputs of the flip-flops. The clock gating logic enables the set/reset signals whenever the clock is turned off, and disables them when the start signal edge occurs.

The slow counter does not stop while the fast counter is running; doing so would make it impossible to run the slow counter from a lower-frequency clock. This leads to an error, since the time interval during which the fast counter is running is double-counted by the slow counter. This error must be removed by keeping track of the accumulated overlap time and advancing the fast counter accordingly. When the advance amount




Figure 3.25: Dual-stage counter block diagram. The arithmetic circuit corrects the overlap error. In the actual system, the fast counter is triggered by the control state machine, so the connection is shown as a dotted line.

exceeds the range of the fast counter, the slow counter is advanced one cycle. A block diagram of this implementation is shown in figure 3.25. The input 18-bit frequency control word is split into the 13 MSBs that go to the slow counter, and the LSBs that are used to control the fast counter. An integrator accumulates the overlap error time, which is added to the fast counter control word. The carry-out bit is added to the control word for the slow counter.

The operation of the compensation circuit is illustrated in figure 3.26. For simplicity, the fast counter is assumed to run from a clock with twice the frequency of the slow counter, and the time interval is set to five fast clock cycles. The fast counter starts when the slow counter reaches zero, and the output pulse is produced when the fast counter reaches zero. The overlap error is shown with red highlighting. The additional cycles added by the compensation loop are highlighted in green. In the first cycle, one additional cycle is added to the fast counter. Because the range of the fast counter would be exceeded in the second cycle, the slow counter is advanced instead.

The bit allocation between the coarse and fine counters needs to be optimized to minimize total power consumption. The slow counter has the equivalent of 4 flip-flops running at the slow clock rate, and it is always on. Therefore, its average power is approximately $4E_{FF}f_{ref}/2^{N_f}$, where E_{FF} is the energy consumed by a flip-flop per clock cycle, and N_f is the number of bits in the fast counter. The fast counter, on average, is enabled for a gating time equal to $T_g = (2^{N_f-1} + 5)/f_{ref}$ (including the “down” pulse) and has the equivalent of 5 flip-flops (the combinational logic can be approximated as another flip-flop). Its average power is thus $5E_{FF}f_{ref}T_gf_{in} = 5E_{FF}(2^{N_f-1} + 5)f_{in}$. The total power consumption is thus equal to:

$$P_{ctr} = 4E_{FF}f_{ref}/2^{N_f} + 5E_{FF}(2^{N_f-1} + 5)f_{in}$$




Figure 3.26: Dual-stage counter overlap compensation example. The counter is set to generate output pulses at 5 fast clock cycle intervals (equal to two slow plus one fast cycle). The red highlighting shows the overlap error. The green highlighting shows extra overlap compensation cycles, which cancel this error.

Control Signal				Next Trigger	State Description
cp_en	fc_start	adc_sh	load		
0	0	0	0	ctr_done	Wait for slow counter
1	0	0	0	clk16	Enable charge pump; wait for settling
1	1	1	0	—	Start fast counter and enable ADC S/H
1	0	1	0	—	Wait for charge pump
0	0	1	0	clk16	Wait for integrator settling
0	0	0	0	adc_done	Wait for SAR conversion to finish
0	0	0	1	—	Generate load signal

Table 3.1: Controller state table. The state machine steps through the states in the order shown. It produces control signals that enable the charge pump, start the fast counter, enable the ADC sample and hold, and generate a load signal.

For the flip-flops used in this work, $E_{FF} = 67fJ$. For $f_{ref} = 256MHz$ and $f_{in} = 2 \times 30kHz$ (for double-edge sampling), the power consumption as a function of fast counter bits is plotted in figure 3.27. The minimum power is achieved for $N_f = 6$, and is equal to $2.56\mu W$. In the actual implementation, $N_f = 5$ was chosen to provide additional timing margin, since the optimum is quite shallow.

Controller

The FDC is controlled by a state machine, which generates all of the required control signals. These include: an enable signal which gates power to the charge pump; the fast counter “start” signal; the ADC “start” signal; the “load” signal that reloads the




Figure 3.27: Total counter dynamic power as a function of the number of fast counter bits. The minimum occurs for $N_f = 6$.




Figure 3.28: SAR ADC block diagram.

slow counter and updates the z^{-1} registers and the edge polarity selector signal. The state table for the controller is shown in table 3.1. The state machine steps through the steps sequentially. To minimize power consumption, clock gating is used to control when the state is advanced. The clock is gated by one of the input trigger signals using combinational logic. The current state is used to select the appropriate input trigger signal and generate the correct outputs.

The operating sequence is as follows. Initially, the controller waits for the slow counter to generate a “done” signal. Upon receipt of this signal, the controller enables the charge pump bias and waits $\approx 2\mu\text{s}$ (which ensures full settling). It then enables the ADC “start” signal and starts the fast counter; the charge pump is disabled one clock cycle later. After waiting another $2\mu\text{s}$ for the active integrator to settle, the ADC “start” signal is released, which initiates the SAR conversion. Upon receipt of the “done” signal from the ADC, the controller generates a “load” signal, which reloads the counters with new values, updates the lead compensator and overlap compensation registers, and inverts the edge selection signal. The controller then returns to its initial state.

The state machine uses the `clk16` output of the slow counter to generate slow delays used to time analog settling. By reusing the signals already generated by the counter, no additional power or area is used by the state machine to generate long delays. The analog circuits can thus take advantage of the relatively long interval between comparison operations, greatly reducing analog power.

The relatively long interval between when the slow counter generates a “done” signal and when it is reloaded may appear to cause a problem. However, this is not the case, because this delay is generated by digital logic and is always constant. Therefore, it is sufficient to simply adjust the programmed N_0 offset to account for this interval.

ADC

The number of bits in the ADC determines the input range of the FDC. Note that a minimum of 3 levels are required to allow the center frequency offset control to cover a continuous range of input frequencies. The input range must allow for the maximum expected input rate signal together with the maximum possible quadrature tone. It must also have some margin to allow for frequency drift and duty cycle mismatch during operation. The second problem is particularly significant: if the input to the comparator is small, even a small change in the DC level will cause a large mismatch between the durations of the positive and negative half-cycles. The input range of the FDC must be large

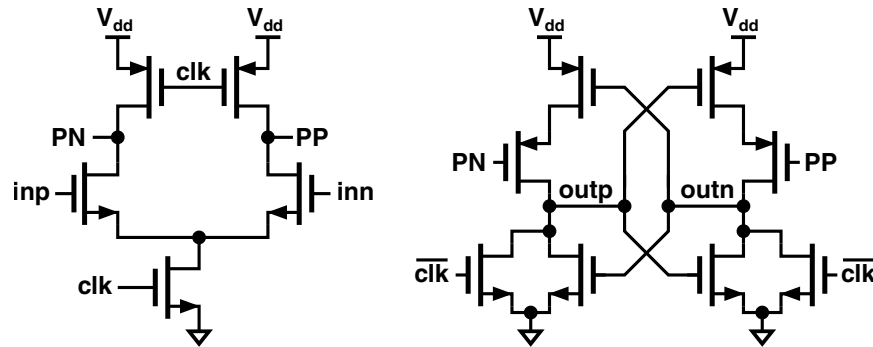


Figure 3.29: ADC comparator schematic.

enough to allow the duty cycle feedback loop to correct the drift without encountering nonlinearity.

The ADC block diagram is shown in figure 3.28. The ADC uses a conventional single-ended bottom-plate sampling charge-redistribution SAR topology with binary-weighted switches [28]. The ADC zero threshold is offset by 1 LSB by connecting the last capacitor to V_{dd} instead of the reference voltage V_{mid} . This ensures the active integrator has sufficient swing to drive the ADC over its full range.

The ADC uses the fully dynamic two-stage clocked comparator shown in fig. 3.29, which avoids the shoot-through current of the conventional “StrongARM” latch design, thus minimizing power [29]. The two-stage architecture also reduces kickback, which is an important consideration for a single-ended design. Power is further minimized by gating the clock feeding the SAR logic.

3.9 Experimental characterization

The FDC was characterized with a function generator input. The measured power spectral density is shown in fig. 3.30. The FDC exhibits the expected 40dB per decade quantization noise slope above 800 Hz. Below that frequency, the noise is dominated by the charge pump and active integrator noise. The 20dB per decade slope corresponds to white phase noise, which could arise either from signal source jitter or from white current noise at the input of the active integrator.

Figure 3.31 shows the power spectrum of the FDC with applied frequency modulation. The modulating signal has a frequency of 100 Hz and the deviation amplitude is 10 Hz, corresponding to a full-scale $\pm 2500^\circ/s$ input with $\alpha_z = 0.7$. The SFDR is about




Figure 3.30: Measured FDC noise density with 30kHz function generator input.

63 dB, and is constant over a wide range of deviation amplitude, suggesting that the linearity may be limited by the function generator used for this test.

Figure 3.32 compares the noise floor with and without applied full-scale modulation. No significant change is observed, proving the FDC has sufficient dynamic range to process full-scale inputs without any degradation in the noise floor. The dynamic range at $\Delta f = 100\text{Hz}$ is 91 dB assuming a gyroscope full-scale amplitude of $\pm 2500^\circ/\text{s}$. This is significantly higher than virtually all low-power AM gyroscopes; the gyroscope thus does not require range switching and can process inputs over the full input range without any degradation in noise performance.

The simulated power consumption of a single FDC is $10.7\mu\text{A}$ at 1.8V, of which $6.1\mu\text{A}$ is consumed by the analog blocks. Another $6.5\mu\text{A}$ is consumed by the output pin driver (assuming a 20pF load capacitance). The measured total power consumption of two FDCs and a clock multiplier PLL is $64\mu\text{A}$ during operation, which agrees with the simulation results.




Figure 3.31: Measured FDC power spectrum with 100 Hz sinusoidal frequency modulation having 10 Hz deviation amplitude.




Figure 3.32: Noise floor with and without applied full-scale modulation.

Chapter 4

LFM Signal Processing

4.1 Background

The LFM demodulator must extract angular rate from the delta-sigma modulated data streams produced by the frequency-to-digital converters. The demodulated FM signal contains angular rate information amplitude-modulated to the split frequency Δf . In addition, this signal contains a large amount of shaped quantization noise, frequency drift components at low frequencies, a large quadrature tone at Δf , and other spurious components. The demodulator must recover the rate information while rejecting the errors to the maximum extent possible.

The LFM demodulator operates as follows. First, the X and Y input signals are filtered and decimated to remove quantization noise. They are then integrated to produce two non-uniformly sampled phase ramps. These phase ramps are then resampled to a uniform sample rate, after which they are subtracted to extract the AM reference phase. The phase ramps are differentiated to produce a pair of resampled frequency signals. The AM reference is then used to demodulate the X and Y rate components from the resampled signals, which are then filtered and added to produce the final rate output. The remainder of this chapter analyzes each of these steps in detail.

4.2 Filtering and resampling

The output of the frequency-to-digital converters is the measured period of the incoming FM signals. Because the period can only be accurately measured at the zero crossings of the signal, the signal is sampled at a uniform phase interval equal to π (for double-

Algorithm 4.1: Integration and resampling procedure

Data: decimated FDC samples (decimation ratio D)
Result: resampled phase ramp

$T_s \leftarrow (D/2)/60\text{kHz}$; // Resampled period
 $T_{clk} \leftarrow 4\text{ns}$; // FDC LSB size
 $N_{os} \leftarrow \text{offset}$; // FDC counter offset
 $T_{curr} \leftarrow 0$; // Current sample time

foreach filtered FDC sample N_i **do**

$T_i \leftarrow D \times T_{clk} \times (N_i + N_{os})$; // Convert LSBs to period

while $T_{curr} \leq T_i$ **do**

$\phi_i \leftarrow T_{curr}/T_i$; // Interpolate next phase sample

$T_{curr} \leftarrow T_{curr} + T_s$; // Next sample time

end

$T_{curr} \leftarrow T_{curr} - T_i$; // Reset origin

end

edge sampling). Because the period is not constant, the signal is non-uniformly sampled in time.¹ If the sequence of measured periods is discrete-time integrated (summed), the result is simply a sequence of timestamps sampled at a uniform phase step π . For further processing, it is necessary to convert this signal to a sequence of phase samples sampled at a uniform time interval. This may be accomplished by interpolation.

Before interpolation can be performed, the oversampled X and Y signals are filtered and decimated to reduce the sampling rate and remove most of the quantization noise. A cascaded integrator-comb (CIC) decimator is an excellent choice for this task [30]. The CIC decimator is a computationally efficient implementation of a cascaded moving-average filter; it requires only adders and memory elements. It is an FIR filter, so it has linear phase and does not introduce phase distortion; this is an extremely important consideration for this application, since accurate phase alignment between the X and Y channels is critical for effectively rejecting quadrature.

At least a 3-stage CIC filter must be used to effectively attenuate second-order quantization noise [30]. The choice of the decimation ratio is a trade-off between the effective attenuation of quantization noise and attenuation of the signal. In LFM operation, the angular rate is represented by the amplitude of the tone at the split frequency. Thus, any passband droop contributes a scale factor error. For a constant split frequency, this

¹This is a fundamental limitation of frequency modulation, and is not specific to any FM demodulator architecture.




Figure 4.1: Split frequency sensitivity as a function of decimation ratio and split frequency for a 3-stage CIC decimator.

is simply a static gain error which can be easily trimmed out. However, the impact of split frequency drift on scale factor stability must also be considered. Figure 4.1 plots the scale factor change per Hz of split frequency as a function of decimation ratio and split frequency.

Figure 4.2 shows the signals before and after a $24\times$, 3-stage CIC decimator. The test signal was generated with a behavioral model of the FDC; full-scale sinusoidal modulation² and white FM noise corresponding to $6 \text{ mdeg/s}/\sqrt{\text{Hz}}$ was added to the source FM signal. The noise floor is unchanged by the decimator, indicating sufficient alias rejection. The second harmonic tone is caused by non-uniform sampling, and will be significantly attenuated after resampling.

The split frequency is not greatly affected by temperature or bias variations, because both gyroscope axis frequencies change by practically the same amount. Thus, the expected frequency drift during operation is quite small. In some cases, the split frequency may drift due to asymmetric package stress or in the presence of a bias voltage offset between the X and Y axes. Since the split frequency is accurately measured during de-

²10 Hz deviation amplitude, 100 Hz frequency.




Figure 4.2: Power spectral density before and after $24\times$, 3-stage CIC decimator with full-scale modulating signal. White FM noise corresponding to $6 \text{ mdeg/s}/\sqrt{\text{Hz}}$ is present.




Figure 4.3: Power spectral density after resampling.

modulation, it is also possible to add a correction factor to the scale factor to account for the filter attenuation, or to employ feedback control to lock the split frequency to a fixed reference value.

After decimation, the resulting signals are integrated to produce corresponding phase ramps. As mentioned earlier, these ramps are sequences of timestamps sampled at a uniform phase interval (and thus a non-uniform time interval). Two-point linear interpolation is used to resample each sequence to a sequence of phase samples sampled at a uniform time interval. To avoid aliasing of interpolation discontinuities, the resampled rate should be at least twice the decimated rate. In order to keep the integration result bounded, a modulo operator is applied to the phase ramps; appropriate steps must be taken to prevent errors near the wrapping points. Algorithm 4.1 implements this procedure.

For verification, this algorithm was used to process the data from fig. 4.2. Figure 4.3 shows the power spectral density after resampling. The noise floor is unchanged from fig. 4.2, and the second harmonic is attenuated by about 24 dB, increasing the SFDR to 96 dB.

4.3 Demodulation reference extraction

In the LFM operating mode, the angular rate is amplitude-modulated onto a suppressed carrier at the split frequency Δf . Furthermore, the undesired quadrature component is modulated onto a quadrature-shifted carrier at the same frequency. In order to extract angular rate while rejecting quadrature, it is necessary to recover the carrier and use it for synchronous AM demodulation.

From the LFM governing equations [7], it is easily seen that the carrier phase is simply the difference between the X and Y channel phase ramps. Therefore, the carrier is recovered simply by differencing the two phase ramps obtained by the resampling operation, and applying a sine function to the resulting phase ramp. Unfortunately, the frequency-to-digital converters do not provide information about the initial phase relationship between the X and Y channels, and so the resulting carrier has a random phase offset $\Delta\phi_0$ equal to the initial phase difference between the X and Y channels at the start of the integration process. Figure 4.4 illustrates this error graphically. The integration is started on the first falling edge of channel Y, when a phase shift $\Delta\phi_0$ is present between the X and Y channels, the the recovered phase ramp is offset by this value. Because the frequency-to-digital converters measure only the time between successive edges and




Figure 4.4: Reference phase extraction and initial error. The top traces show the time-domain FM waveforms. The middle plot shows the phase ramps as a function of time; FDC samples are indicated with dots. The bottom plot shows the true reference phase and the recovered version. They are identical except for an initial phase offset $\Delta\phi_0$. The zero-phase points of the reference ramp are generally not aligned with the edges of the FM waveform, introducing a quantization error if a phase detector is used to estimate $\Delta\phi_0$.

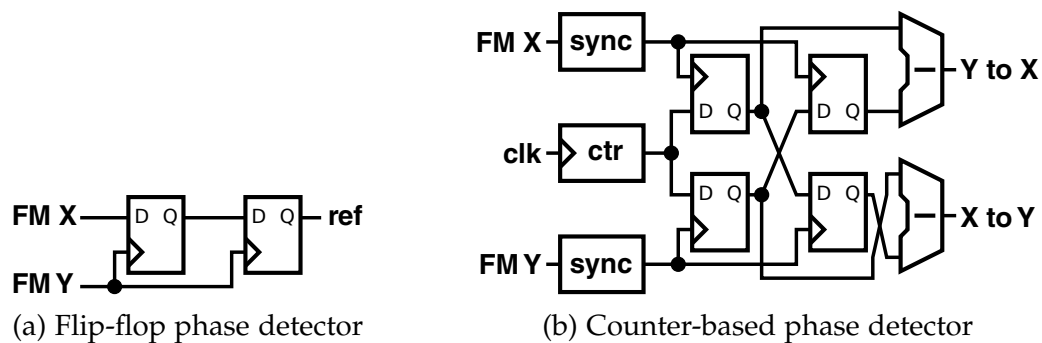


Figure 4.5: Reference phase extraction circuits.




Figure 4.6: Simulated phase estimate error as a function of time for the bang-bang phase detector phase recovery circuit for $\Delta f = 98\text{Hz}$ and $\Delta f = 100\text{Hz}$ (nominal frequency 30kHz).

do not produce an absolute timestamp, it is not possible to reconstruct $\Delta\phi_0$ from this measurement unless Δt_0 can be measured.

In order to achieve good bias stability and repeatability, this phase error must be known precisely. In general, the maximum allowable error is

$$\Delta\phi_{0,err} = \tan^{-1} \left(\frac{\Omega_{bias}}{\Omega_{quad}} \right).$$

For a bias stability of $5^\circ/\text{hr}$ and a quadrature error of $150^\circ/\text{s}$, this translates to a maximum error of $\pm 5 \times 10^{-4}$ degrees, corresponding to a time resolution of 15ns for $\Delta f = 100\text{Hz}$.

This error may be estimated and removed by several methods. The simplest such method takes advantage of the fact that the quadrature signal is independent of input rate. By applying rate impulses and adjusting the phase offset to minimize the sensitivity of the quadrature channel to input rate, it is possible to find the correct $\Delta\phi$. By measuring the relative sensitivity of the in-phase and quadrature channel, the phase error can be estimated. This is a useful lab calibration technique that can be used to verify the correct



(a) $\Delta f = 98\text{Hz}$ (b) $\Delta f = 101\text{Hz}$

Figure 4.7: Simulated phase estimate error as a function of time for the counter-based phase recovery circuit with a time resolution of 200ns for $\Delta f = 98\text{Hz}$ and $\Delta f = 100\text{Hz}$ (nominal frequency 30kHz).

operation of the phase initialization circuit, but it cannot be used in the field unless a suitable calibration input can be provided.

A bang-bang phase detector can be used to detect when the input edges are approximately coincident; the phase of the reference at this instant can be adjusted to be zero. This circuit is shown in figure 4.5a. With this approach, the quantization error is equal to at least one period of the center frequency. Figure 4.6 shows the simulated error over time for this topology for two different values of Δf . There is a significant offset, which is caused by the one-cycle delay of the synchronizing flip-flop. The error varies by about 1.2° , corresponding to a quantization step of one 30kHz cycle. Furthermore, the error has a static offset that is very sensitive to the value of the split frequency. For some split frequencies, the period of variation can be very slow, and will appear as simply a time-varying offset. Non-ideal behavior of the flip-flops, such as metastability, as well as variations in setup and hold time over temperature and supply voltage further degrade performance of this circuit.

A much better phase reference circuit comprises a counter or TDC that measures the time Δt_0 between X and Y axis edges, as shown in figure 4.4. This may be implemented simply with a counter running at a suitably high clock rate, two registers that record the counts at the X and Y input edges, and appropriate arithmetic logic; figure 4.5b shows such an implementation. Since the X and Y axis frequencies are known, $\Delta\phi_0$ can be easily found from Δt_0 . The error of a counter is first-order shaped, since the quantization error in a given cycle is carried forward to the next cycle. Thus, unlike the previous case, averaging is highly effective for reducing the quantization errors. The simulated performance of this phase recovery circuit for two different split frequencies is shown in figure 4.7. There is no steady-state error, and even with a quite modest counter clock frequency of 5 MHz, the circuit achieves a phase error of well under 1×10^{-6} degrees with some averaging. Unlike the bang-bang phase detector, drift in the split frequency does not introduce a time-varying error.

Algorithm 4.2 describes the procedure for generating the reference phase from the sequence of time differences produced by the counter-based phase detector. The input FDC samples are converted to timestamps T_i and used together with the frequency split Δf to construct the phase ramp, which has the correct frequency but an unknown phase offset. The phase detector samples Δt_i represent the time delta to the nearest preceding edge on the other channel; multiplying this value by $2\pi\Delta f$ yields an estimate of the relative phase between the two channels at the sampling instant. By subtracting the reconstructed phase value at this instant from the estimate and filtering the result, an accurate estimate of the phase offset ϕ_{os} is obtained. ϕ_{os} is then subtracted from the

reconstructed phase ramp to produce the final phase ramp, to which a sin operator is applied. An additional phase shift or time delay may be added to compensate for processing delays (e.g. from the decimation filter).

Algorithm 4.2: Reference phase generation procedure

Data: filtered FDC samples N_i ; phase det. samples Δt_i ; freq. split Δf

Result: reference phase ramp

$T_s \leftarrow 1/60\text{kHz}$; // Resampled period

$N_{os} \leftarrow \text{offset}$; // FDC counter offset

$T_{curr} \leftarrow 0$; // Current sample time

$\phi_{ref} \leftarrow 0$; // Reference phase integral

foreach pair(FDC sample N_i , time difference Δt_i) **do**

$T_i \leftarrow T_{clk} \times (N_i + N_{os})$; // Convert FDC value to period

 // Update phase offset using measured Δt_i

$\phi_{os} \leftarrow \text{filter}([2\pi\Delta f\Delta t_i - \phi_{ref} - 2\pi\Delta f(T_i - T_{curr})] \bmod 2\pi)$;

while $T_{curr} \leq T_i$ **do**

$\phi_{ref} \leftarrow (\phi_{ref} + 2\pi\Delta fT_s) \bmod 2\pi$; // Phase ramp

$\phi_{out,i} \leftarrow \phi_{ref} + \phi_{os}$; // Generate output phase value

$T_{curr} \leftarrow T_{curr} + T_s$; // Next sample time

end

$T_{curr} \leftarrow T_{curr} - T_i$; // Reset origin

end

4.4 Synchronous demodulation

The resampled phase ramps for the X and Y channels produced in the resampling step can be differentiated to produce a pair of uniformly sampled frequency signals. The rate signal is extracted by mixing each of these signals with the recovered carrier, and low-pass filtering the result. The X and Y signals are added together to produce the final rate output. Demodulating and summing both the X and Y signals increases SNR, cancels the linear component of cross-damping error, and greatly reduces sensitivity to amplitude variations because of reciprocal summation.

The final low-pass filter removes unwanted components, including any remaining quantization noise and slow frequency drift components, which become a tone at the split frequency after mixing. The specific design of this filter is determined by the ap-

plication requirements. The required computational resources are quite modest because this filter operates at a low sampling rate.

4.5 Quadrature extraction

A similar procedure can be followed to extract the quadrature component from the modulated LFM signal. While the quadrature carries no rate information, it is useful for a variety of diagnostic and calibration-related purposes. For example, the measured quadrature signal can be used to adjust the bias voltage to quadrature cancellation electrodes. Furthermore, because the amount of quadrature error is inversely correlated with the symmetry of the mechanical structure, such a measurement can be used during wafer test to detect defective or out-of-spec transducers. Finally, by measuring the change in the quadrature signal as a function of applied rate, it is possible to estimate and trim out any residual phase errors in the reference phase detector. Such errors may arise due to mismatched delays between the X and Y oscillator and comparator circuits.

To measure the quadrature component, it is necessary to generate both sine and cosine signals in the reference generation block. The cosine output is then connected to another synchronous demodulator, which produces the quadrature output.

4.6 Experimental results

The complete gyroscope system was assembled on top of a rate table. The circuits previously described were used to drive the transducer and to demodulate the FM signals. The resulting data streams were then fed to a personal computer, where LFM demodulation was performed in software.

The gyroscope bias and scale factor stability were characterized with a 50-hour test. The test consisted of repeated $+90$, 0 , and $-90^\circ/\text{s}$ angular rate measurements, which allowed bias and scale factor stability to be measured simultaneously. Testing was performed at room temperature with no active temperature regulation.

Figure 4.11 shows the measured scale factor error drift over time. Due to velocity amplitude mismatch between the X and Y channels caused by circuit imperfections and drift, individual X and Y axis signals exhibited a 1.9% initial scale factor error and additional drift of 0.45% during the test; this is a typical performance level for a conventional AM gyro with no calibration circuits. Combining the X and Y signals almost entirely eliminates this error source: fig. 4.11 shows no significant drift trend for the summed




Figure 4.8: Bias Allan deviation.




Figure 4.9: Scale factor Allan deviation.

output. This result vividly illustrates the benefits of reciprocal velocity ratio sensitivity provided by the LFM operating mode.

The Allan variance of the scale factor reaches a minimum of 6.7 ppm at $\tau = 15400$ s (fig. 4.9), which is near the specified performance limit of the rate table used for this test. This level of scale factor stability is one of the lowest reported to date, and is extremely difficult to match with an amplitude-modulated gyroscope without sophisticated and power-hungry calibration layers.

Bias stability reaches a minimum of $5.9^\circ/\text{hr}$ at $\tau = 3800$ s (fig. 4.8). This represents a $250\times$ improvement over the single-transducer circular orbit FM result in [13], and a $60\times$ improvement over the counter-rotating circular orbit FM implementation. The ARW in fig. 4.8 is inaccurate due to noise folding resulting from the test protocol: since the rate table was allowed to fully settle before a sample was taken, some aliasing was unavoidable.

Figure 4.10 shows the measured rate noise density. The rate noise density (i.e. angle random walk) is $0.014^\circ/\text{s}/\sqrt{\text{Hz}}$, and is limited by the phase noise of the discrete oscillator circuits, which are in turn limited by the PCB parasitics. The bandwidth of the gyroscope is 50Hz; the full-scale rate exceeds $\pm 1000^\circ/\text{s}$. The maximum tested full-scale rate was $360^\circ/\text{s}$ due to equipment limitations.

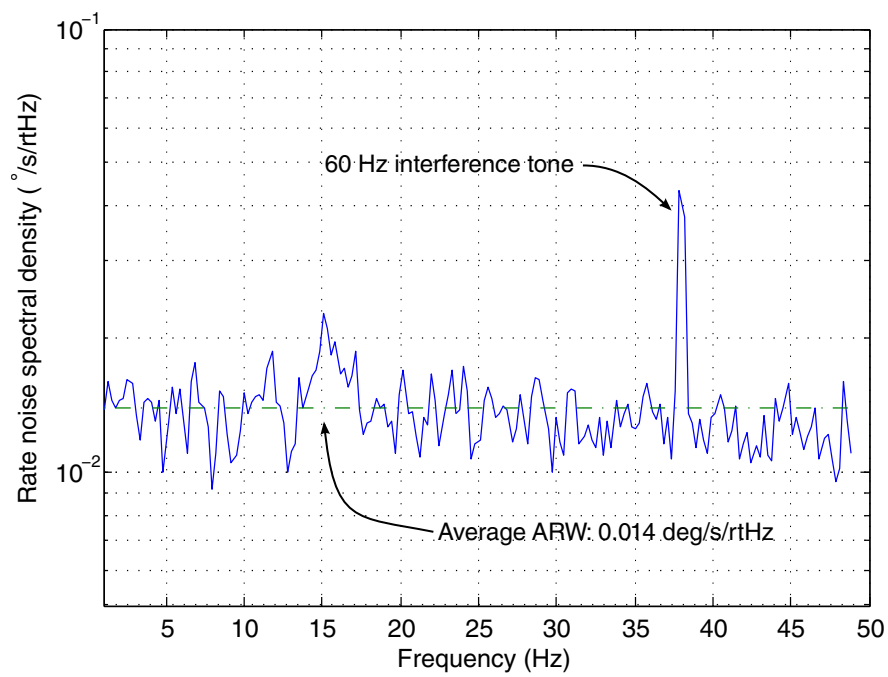


Figure 4.10: In-band rate noise power spectral density (angle random walk). The tone is caused by 60Hz power line interference coupling to the sustaining circuits. The rate signal is at about 100Hz; after demodulation the 60Hz tone appears near 40Hz.




Figure 4.11: Time-domain scale factor measurement. The individual X and Y axis measurements show significant drift during the test. Summing them almost entirely eliminates the error. A one-hour moving average was applied to the combined curve to reduce short-term ARW-related noise and show the long-term trend.

Chapter 5

Conclusion

This work has presented a full implementation of the Lissajous FM gyroscope. The experimental results clearly demonstrate the advantages of LFM operation in the areas of bias and scale factor stability, as well as dynamic range. The monolithic implementation of the sigma-delta frequency-to-digital converter demonstrates very high resolution in combination with exceptionally low power consumption. While the sustaining circuits, comparators, and DSP blocks have been implemented off-chip in this prototype, a full on-chip implementation of this architecture could enable a new generation of consumer-grade gyroscopes, with sub-100 μ W power consumption and the highest performance in their class.

Bibliography

- [1] *UBX-G7020 Product Summary*, u-blox AG, Sep. 2014. [Online]. Available: <http://www.u-blox.com/en/gps-chips/stand-alone-gps-chips/ubx-g7020-ctkka.html>
- [2] H. Liu, H. Darabi, P. Banerjee, and J. Liu, "Survey of wireless indoor positioning techniques and systems," *IEEE Transactions on Systems, Man, and Cybernetics—Part C: Applications and Reviews*, vol. 37, no. 6, pp. 1067–1080, Nov 2007.
- [3] *MPU-6000/MPU-6050 Product Specification*, InvenSense Inc., Aug. 2013.
- [4] *BMG160 Data sheet*, Bosch Sensortec GmbH, May 2014.
- [5] *A3G4250D 3-axis digital output gyroscope*, STMicroelectronics, Feb. 2012.
- [6] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*. Boston, MA: Artech House, 2008.
- [7] M. H. Kline, "Frequency modulated gyroscopes," Ph.D. dissertation, University of California, Berkeley, Dec. 2013.
- [8] "IEEE standard specification format guide and test procedure for single-axis interferometric fiber optic gyros," *IEEE Std 952-1997*, 1998.
- [9] W. R. Bussone, "Linear and angular head accelerations in daily life," Master's thesis, Virginia Polytechnic Institute and State University, Jul. 2005. [Online]. Available: <http://hdl.handle.net/10919/34615>
- [10] (2014, Nov.) Nexus 5 tech specs. Google Inc. Mountain View, CA. [Online]. Available: <https://support.google.com/nexus/answer/3467463?hl=en>
- [11] *CRG20 Digital Angular Rate Sensor*, Silicon Sensing Systems Ltd., Jan. 2010. [Online]. Available: <http://www.siliconsensing.com/information-centre/downloads-library/>

- [12] C. D. Ezekwe and B. E. Boser, "A mode-matching $\Sigma\Delta$ closed-loop vibratory gyroscope readout interface with a $0.004^\circ/\text{s}/\sqrt{\text{Hz}}$ noise floor over a 50 Hz band," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 12, pp. 3039–3048, Dec. 2008.
- [13] M. H. Kline, Y. Yeh, B. Eminoglu, H. Najjar, M. Daneman, D. A. Horsley, and B. E. Boser, "Quadrature FM gyroscope," in *2013 IEEE 26th International Conference on Micro Electro Mechanical Systems (MEMS)*, 2013, pp. 604–608.
- [14] E. A. Vittoz, *Low-Power Crystal and MEMS Oscillators*. Berlin: Springer, 2010.
- [15] A. Hajimiri and T. Lee, "A general theory of phase noise in electrical oscillators," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 2, pp. 179–194, Feb. 1998.
- [16] A. Trusov, I. Prikhodko, S. Zotov, A. Schofield, and A. Shkel, "Ultra-high Q silicon gyroscopes with interchangeable rate and whole angle modes of operation," in *2010 IEEE Sensors*, Nov. 2010, pp. 864–867.
- [17] J. R. Carson, "Notes on the theory of modulation," *Proceedings of the Institute of Radio Engineers*, vol. 10, no. 1, pp. 57–64, Feb. 1922.
- [18] E. H. Armstrong, "A method of reducing disturbances in radio signaling by a system of frequency modulation," *Proceedings of the Institute of Radio Engineers*, vol. 24, no. 5, pp. 689–740, May 1936.
- [19] *HF2 User Manual*, Zurich Instruments AG, Nov. 2014. [Online]. Available: <http://www.zhinst.com/manuals/hf2>
- [20] *Si4700-B15 Broadcast FM Radio Tuner for Portable Applications*, Silicon Laboratories Inc., Jul. 2006.
- [21] J. P. Costas, "Synchronous communications," *IRE Transactions on Communications Systems*, vol. 5, no. 1, pp. 99–105, March 1957.
- [22] Z. Deng and A. M. Niknejad, "The speed-power trade-off in the design of CMOS true-single-phase-clock dividers," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 11, pp. 2457–2465, Nov. 2010.
- [23] M. Straayer and M. Perrott, "A multi-path gated ring oscillator TDC with first-order noise shaping," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 4, pp. 1089–1098, Apr. 2009.

- [24] A. Elshazly, S. Rao, B. Young, and P. K. Hanumolu, "A 13b 315fsrms 2mW 500MS/s 1MHz bandwidth highly digital time-to-digital converter using switched ring oscillators," in *2012 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb. 2012, pp. 464–466.
- [25] A. Abidi, "Phase noise and jitter in CMOS ring oscillators," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 8, pp. 1803–1816, Aug. 2006.
- [26] I. Galton, "Analog-input digital phase-locked loops for precise frequency and phase demodulation," *IEEE Transactions on Circuits and Systems—Part II: Analog and Digital Signal Processing*, vol. 42, no. 10, pp. 621–630, Oct. 1995.
- [27] I. Galton, W. Huff, P. Carbone, and E. Siragusa, "A delta-sigma PLL for 14-b, 50 kSample/s frequency-to-digital conversion of a 10 MHz FM signal," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 12, pp. 2042–2053, Dec. 1998.
- [28] J. McCreary and P. Gray, "All-MOS charge redistribution analog-to-digital conversion techniques—part I," *IEEE Journal of Solid-State Circuits*, vol. 10, no. 6, pp. 371–379, Dec. 1975.
- [29] P. Harpe, Y. Zhang, G. Dolmans, K. Philips, and H. de Groot, "A 7-to-10b 0-to-4MS/s flexible SAR ADC with 6.5-to-16fJ/conversion-step," in *2012 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb. 2012, pp. 472–474.
- [30] S. R. Norsworthy, R. Schreier, and G. C. Temes, *Delta-Sigma Data Converters: Theory, Design, and Simulation*. Hoboken, NJ: Wiley–IEEE Press, 1996.