DESIGN OF FULLY DIGITAL INDUCTORS FOR LOW BANDWIDTH FILTER APPLICATIONS

BY

BRAEDON SALZ

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2017

Urbana, Illinois

Adviser:

Associate Professor Pavan Kumar Hanumolu

ABSTRACT

We are currently living in the age of intelligent machines, where we are interested in acquiring data and making decisions all from some sort of embedded environment. Of particular value are personal health metrics, such as the analysis of heart rate, muscle action potentials, and brain waves. Collecting this data requires new advances in the circuitry behind much of classical filter design. In this thesis, we present a digital inductor based on time-domain signal processing. This approach uses the phase-domain theory that is well-known and understood in the fields of clocking and serial links and applies it to analog circuit design. By using a ring oscillator to integrate the input voltage and a switched transconductor to inject current into the input node, the proposed time-domain gyrator achieves inductive input impedance without using either large resistors or capacitors. Realizing the gyrator in this manner makes it significantly more amenable for technology scaling. Fabricated in 65 nm CMOS process, the inductor operates from a 0.7 V supply voltage and consumes 528 μW. Measurement results show inductance values in the range of 150 μH to 1.5 mH can be achieved.

You can't let your body control your destiny.

ACKNOWLEDGMENTS

I would like to begin by thanking my adviser, Pavan Hanumolu. Without his encouragement, none of this would have been possible. Anyone can make you work, but it takes a special kind of adviser to make you think you're having fun while doing so. I would also like to thank Bibhudatta Sahoo for keeping me motivated and helping out whenever he could during the slightly arduous tape-out process. It would not have been possible without you!

Thanks as well to everyone else in the research group: Guanghua Shu, Romesh Nandwana, Ahmed Elkholy, Woo-Seok Choi, Safwat Elmallah, Amr Khashaba, Dongwook Kim, Seong-Joong Kim, Karim Megawer, Timir Nandi, Da Wei, and Charlie Zhu. I am mostly just listing your names so when I am soon unemployed, I'll know who to call. Special thanks to Daniel Coombs, who would always make me find a better of way of doing things (for him). Additionally, thanks to Karen Lamb, who allowed me to ignore her far more than any two people in a relationship should be comfortable with in order to create this work.

Lastly, thanks to my family for all their hard work that allowed me to have this great opportunity.

TABLE OF CONTENTS

CHAPTER 1 INTRODUCTION 1
1.1 Motivation 1
1.2 Outline 2

CHAPTER 2 BACKGROUND MATERIAL 3
2.1 Active Inductors 3
2.2 Time-Domain Signal Processing 7

CHAPTER 3 INDUCTOR DESIGN 13
3.1 System Modeling 13
3.2 Stability 19
3.3 Circuit Implementation 24

CHAPTER 4 MEASUREMENTS 31

CHAPTER 5 CONCLUSION 34
5.1 Future Work 34

REFERENCES 36

CHAPTER 1

INTRODUCTION

1.1 Motivation

Technology scaling has allowed for huge benefits in our modern world. Passive devices however, such as resistors, capacitors, and inductors, have not enjoyed the same Dennard scaling as transistors [1]. As transistor dimensions shrink more and more, the area efficiency of using active designs over passive designs becomes more and more appealing. Designs which take advantage of this like the switched-capacitor resistor or capacitance multiplier [2] are now commonplace in both research and industry. While many active inductor circuits exist, none have yet taken advantage of the benefits of time-domain techniques.

This trend is also true in analog circuit designs, where the demand for lower supply voltage and power has driven interest in replicating the circuitry digitally. Analog circuits such as filters [3], data converters [4], and amplifiers [5] have all been redesigned to take advantage of the benefits of time-domain control techniques. However, these all focus on the upper end of the usage spectrum - high power and high performance. On the other end is an equally difficult challenge: creating single-Hertz filters with the same dynamic range and voltages that one would expect out of a typical system. Currently, the only solution space that exists requires large, bulky passive devices sitting outside the chip, as the area requirements simply cannot be met with an integration solution.

Furthermore, there exists a new need for these devices thanks to the era of "intelligent machines". While machine learning and deep learning systems are creating impressive new results every day, they need digital inputs. Some of the most promising applications of this are personal biomedical systems that monitor vital signals and make intelligent health decisions. However,

signals like electroencephalogram (EEG), electrocardiogram (ECG), or electromyogram (EMG) require ultra-low bandwidth systems with high precision. This data needs significant pre-processing in analog front ends to make it usable for these intelligent systems. The front end demands currently cannot be met in an integrated system due to the previously mentioned area limitations.

Much work has been done to meet these demands using traditional analog circuit design. $G_M - C$ filters remain the most popular approach [6, 7, 8], but suffer from large area. When implemented in a filter bank approach [9], they suffer a reduced voltage swing in order to maintain the necessary linearity requirements. In [10], they attempt to use a switched-capacitor design instead of a continuous-time one, but at the expense of power.

1.2 Outline

This thesis is organized as follows. Chapter 2 explains the background material, specifically previous work in active inductors and fundamental concepts in time-domain signal processing. Chapter 3 covers the simulation and design of a time-domain inductor, and measurement results are presented in Chapter 4. Finally, our conclusions are summarized in Chapter 5.

CHAPTER 2

BACKGROUND MATERIAL

In this chapter, we go over previous active inductor implementations and their limitations. Additionally, we cover basic time-domain processing techniques and their applications to circuit design.

2.1 Active Inductors

The idea of creating an active inductor is not new [11], especially for filter applications. Most implementations revolve around the gyrator circuit, shown in Figure 2.1. We will assume both G_M cells are ideal transconductors: zero input current, and output current linearly proportional to the input voltage. We see that the forward transconductor will have an output voltage $V_X = V_{in}G_{MF}/sC_G$, which will produce an output current $I_B = G_{MB}V_X$. We can find the impedance looking into this circuit as:

$$Z_G = \frac{V_{in}}{I_B} = \frac{sC_G}{G_{MF}G_{MB}} \tag{2.1}$$

We can see this circuit now behaves inductively, as the impedance scales linearly with s. Assuming both transconductors have the same gain, we can say the inductance $\hat{L}$ is given as:

$$\hat{L} = \frac{C_G}{{G_M}^2} \tag{2.2}$$

Next, we add some terms to model non-idealities of our circuit design, shown in Figure 2.2. The forward- and back-transconductance cells have their own output impedances R_{OF} and R_{OB}, respectively. Additionally, there is some parasitic capacitance at the input node of the circuit, C_O. We assume the input impedance for each transconductor is large enough to be negligible.

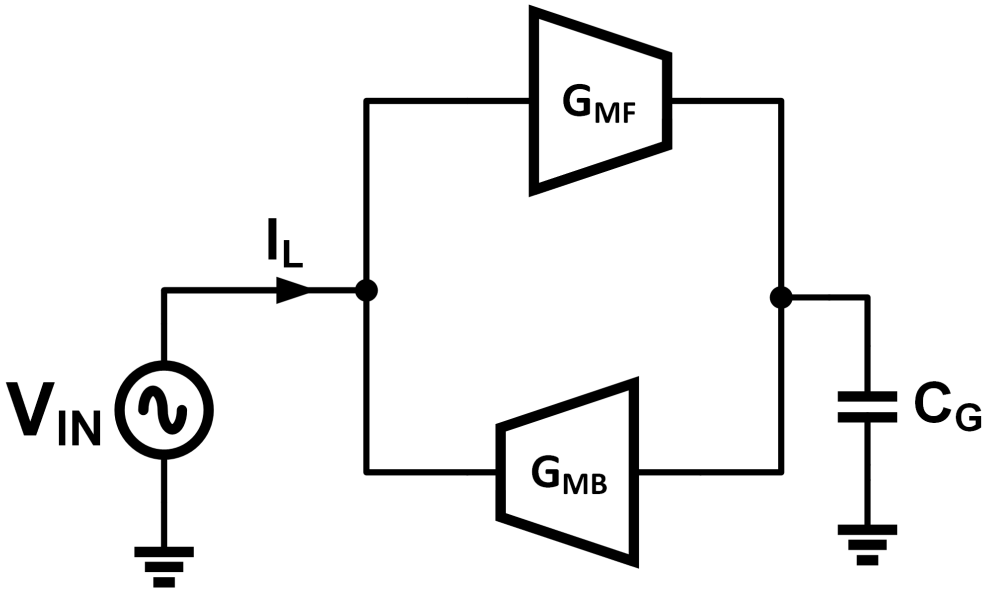

Figure 2.1: A common active inductor designed called the "gyrator" circuit.

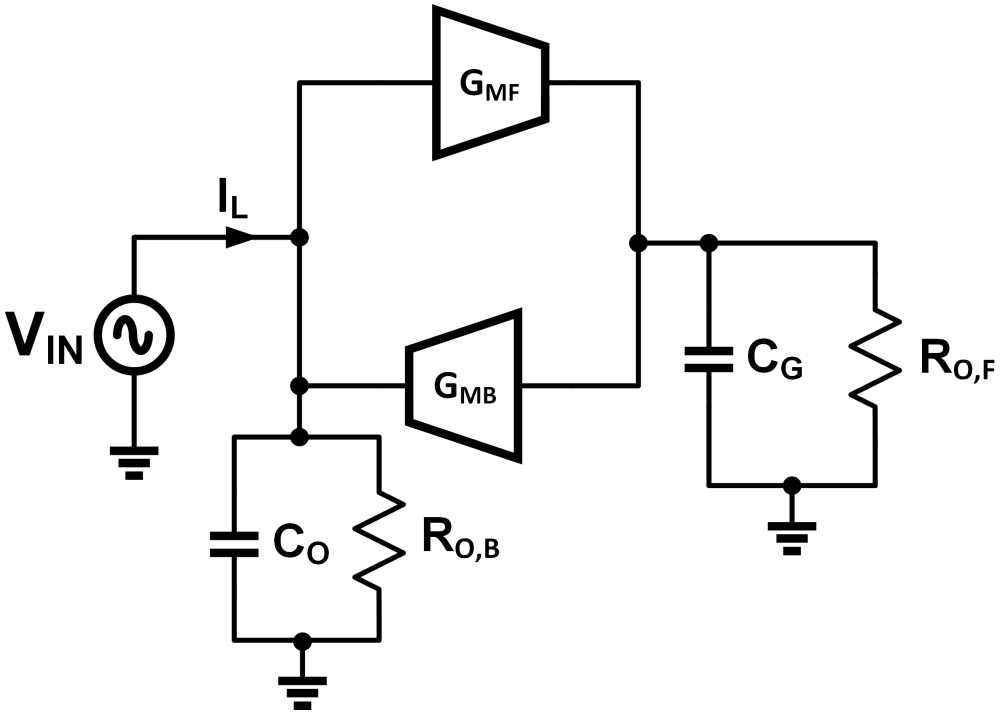

Figure 2.2: A gyrator circuit with practical limitations modeled.

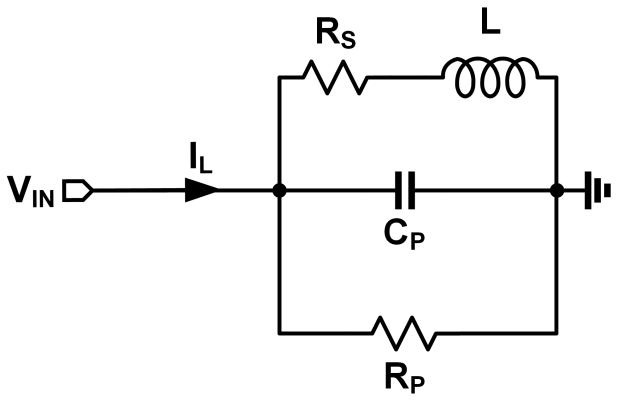

Figure 2.3: RLC representation of the gyrator circuit.

Lastly, we assume the parasitic impedance at node V_X is negligible compared to C_G, or already accounted for. Full derivations are given in [12], but we can reduce the gyrator model to an RLC model, shown in Figure 2.3. It is clear that as the R_O terms go to infinity and C_P goes to zero, this becomes our ideal model from before. The component values in terms of the gyrator configuration are given as:

$$\begin{aligned} R_P &= R_{O,B} \\ C_P &= C_O \\ R_S &= \frac{1}{R_{O,F} G_{MF} G_{MB}} \\ L &= \frac{C_G}{G_{MF} G_{MB}} \end{aligned} \tag{2.3}$$

The total impedance looking into the gyrator can be calculated as:

$$Z_G = \left(\frac{R_S}{C_P L}\right) \frac{s\dfrac{L}{R_S} + 1}{s^2 + s\left(\dfrac{1}{R_P C_P} + \dfrac{R_S}{L}\right) + \dfrac{R_P + R_S}{R_P C_P L}} \tag{2.4}$$

If we plot the magnitude of the impedance over frequency, we get the approximate Bode plot shown in Figure 2.4. The low-frequency magnitude is a constant resistance until the zero at f_z; then we see inductive behavior until the double pole at f_p, after which it appears capacitive. While it may be obvious that we want to maximize the range of inductive behavior, it is not necessarily clear how to do so.

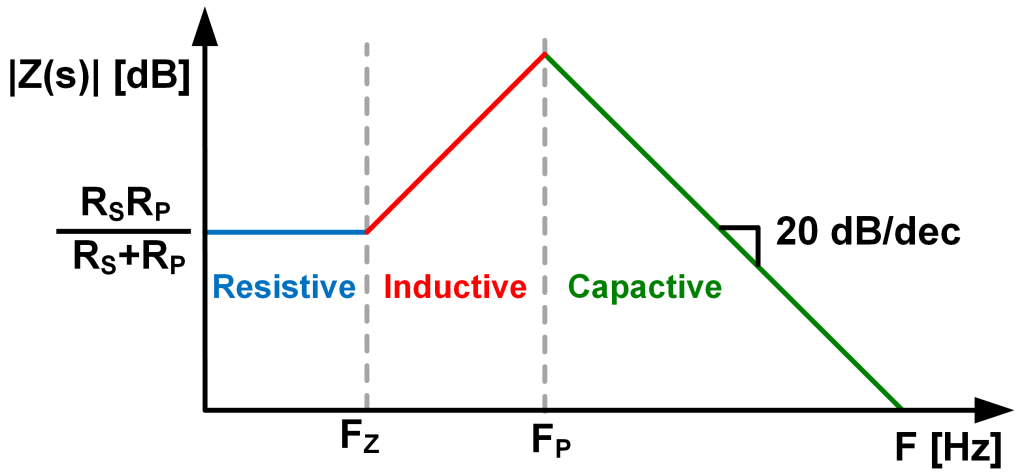

Figure 2.4: Bode plot of the gyrator's impedance Z_G.

Instead of plotting the magnitude of the impedance, we judge our active inductor's performance by plotting the Q, or quality factor. The Q of a system has many definitions, but can be broadly described as a measurement of loss. In bandpass filters, Q is defined as the center frequency of a filter divided by its bandwidth. In filter components, Q is defined as [12]:

$$Q = 2\pi \times \frac{\text{Net magnetic energy stored}}{\text{Energy dissipated in one cycle}} = \frac{\Im m[Z(jw)]}{\Re e[Z(jw)]} \tag{2.5}$$

Rather than analyze one large equation, it is simpler to decompose our Q into $Q = Q_1 \cdot Q_2 \cdot Q_3$, as given in Eq. 2.6. Q_1 is the "classical" Q of a real inductor, as every inductor has some DC resistance R_S. Q_2 represents the limitation in impedance due to the output impedance of the transconductance cell. Lastly, Q_3 is the limitation due to the natural frequency of the gyrator set by the parasitic capacitance.

$$Q_1 = \frac{\omega L}{R_S} \tag{2.6}$$

$$Q_2 = \frac{R_P}{R_P + R_S\left[1 + \left(\frac{\omega L}{R_S}\right)^2\right]} \tag{2.7}$$

$$Q_3 = 1 - \frac{{R_S}^2 C_P}{L} - \omega^2 L C_P \tag{2.8}$$

We can plot these with some common values ($R_P = 1$ kΩ, $R_S = 5$ Ω,

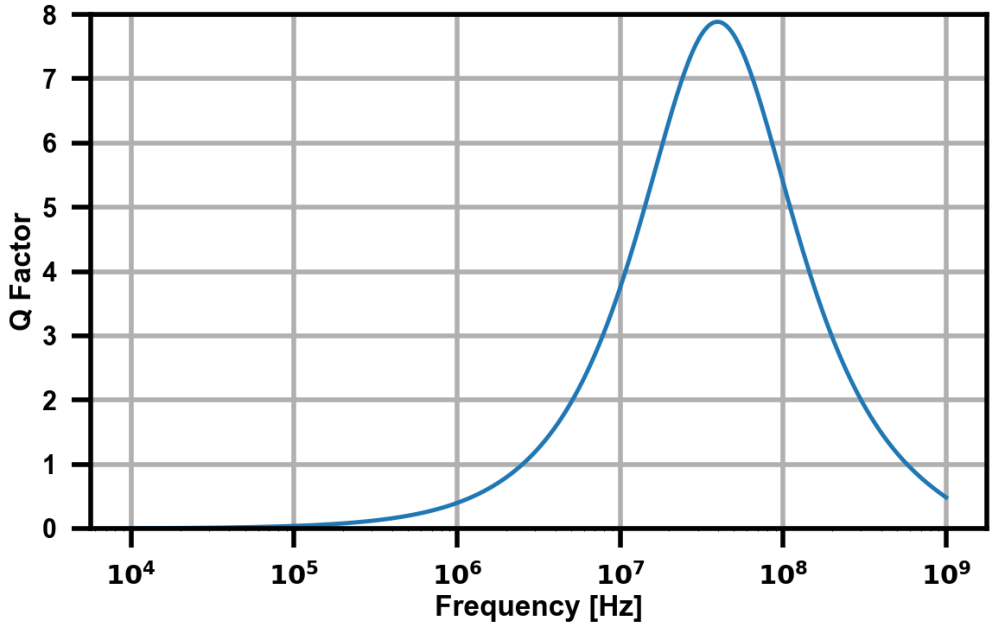

Figure 2.5: Active inductor Q over frequency.

$C_P = 150$ fF, and $L = 1.6$ μH). The plot shown in Figure 2.5 has a very similar shape to the Bode plot shown in Figure 2.4, starting flat, rising around peak inductance point, then falling off again at high frequencies. Therefore it is indeed a useful visualization to compare various active inductor designs.

While the chart is useful visually, the numbers are currently presented without context. Large passive inductors have Q's in the range of 60-80 while operating at the intended frequencies. These are designed for use on printed circuit boards (PCBs). Next, spiral on-chip inductors have Q's in the range of 25-50, but take up significant die area [13]. Lastly, gyrators and other active inductor designs usually only reach a Q of 10-12 at maximum.

2.2 Time-Domain Signal Processing

All classical control systems need three essential blocks: differencers, gain, and integrators. In analog systems, these are implemented by various amplifier designs. In digital systems, these are implemented using Boolean algebra to perform the math. In time-domain signal processing, we need to create these blocks as well. Here, we go over the design of these classical building blocks.

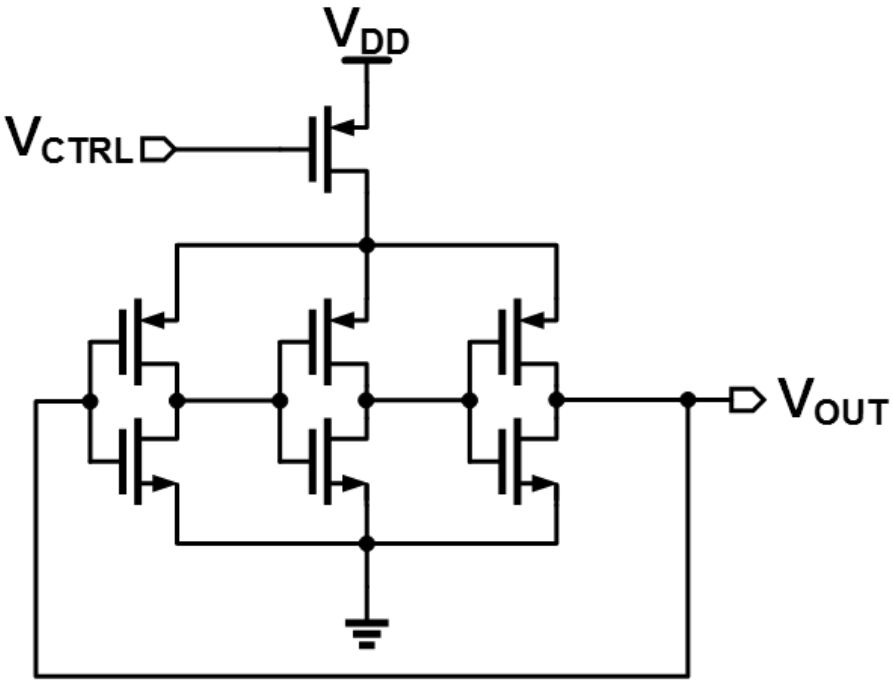

Figure 2.6: A three-stage inverter chain used as a voltage-controlled ring oscillator.

2.2.1 Voltage Controlled Oscillator

We will first consider the case of using a voltage-controlled oscillator (VCO) as an integrator. We can "black-box" the VCO as a system that takes a voltage input and produces a frequency output with perfect linearity. We can define the frequency-domain gain of this block as:

$$K_{VCO} = \frac{F_{OUT}}{V_{IN}} \tag{2.9}$$

Now instead of measuring frequency, one can instead measure the phase of the output. Phase is the integral of frequency, or in the Laplace domain: $\Phi_{OUT} = F_{OUT}/s$. Now the transfer function from input voltage to output phase is given as:

$$H_{VCO}(s) = \frac{\Phi_{OUT}}{V_{IN}} = \frac{F_{OUT}}{sV_{IN}} = \frac{K_{VCO}}{s} \tag{2.10}$$

This means the oscillator is an ideal voltage-to-phase integrator. We note there is no idea of limited DC gain, as the VCO is a true integrator with infinite DC gain. The bandwidth of the integrator is set by the gain K_{VCO}.

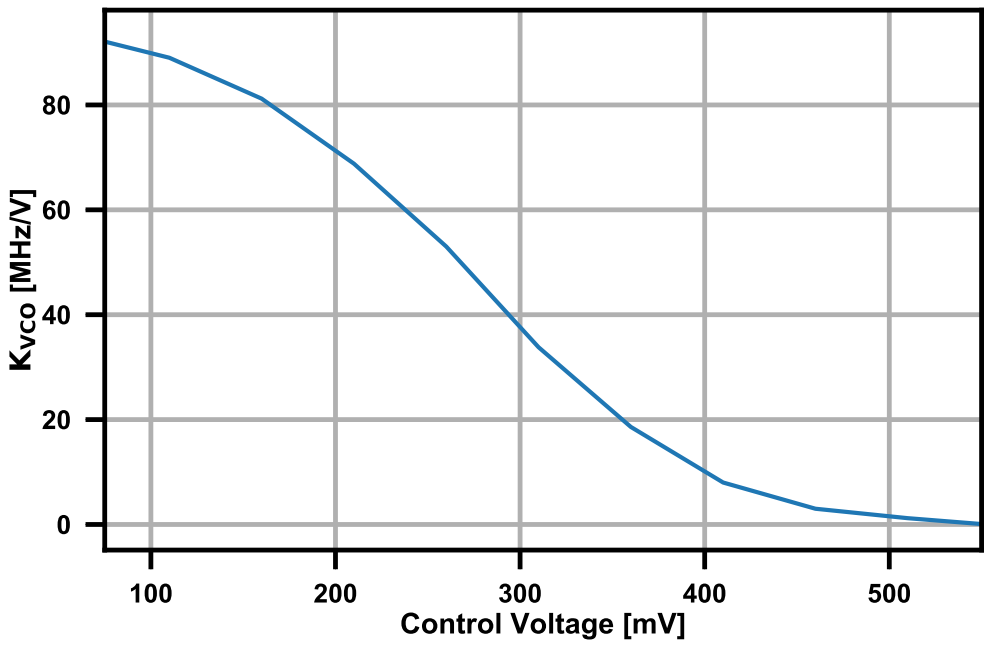

Figure 2.7: The nonlinearity of K_{VCO} in an example ring oscillator.

One common implementation of a VCO is through a ring oscillator. This design consists of an odd number of ring inverters tied in feedback, which will produce a square wave output. The frequency of that output, F_{OUT}, will be inversely proportional to the propagation delay of each stage:

$$F_{OUT} = \frac{1}{2\sum_{i} t_{p,i}} \tag{2.11}$$

A current-limiting block is used to alter the delay in each stage, ultimately changing the total output frequency. In Figure 2.6, this block is a simple PMOS device. As V_{CTRL} changes, the PMOS will allow more or less current into the oscillator. Higher currents correspond to higher output frequencies, and vice versa. The output of Figure 2.6 is labeled as V_{OUT}, as the inverter outputs a voltage-domain signal. However, it is important to note that the integration occurs in the phase domain, not the voltage domain. This is expanded upon later in Section 2.2.2.

Another important criterion for any gain block is the linearity. We simulated a simple ring oscillator and measured the output frequency, then used a numerical derivative to find K_{VCO}. We can use known formulas [14] to

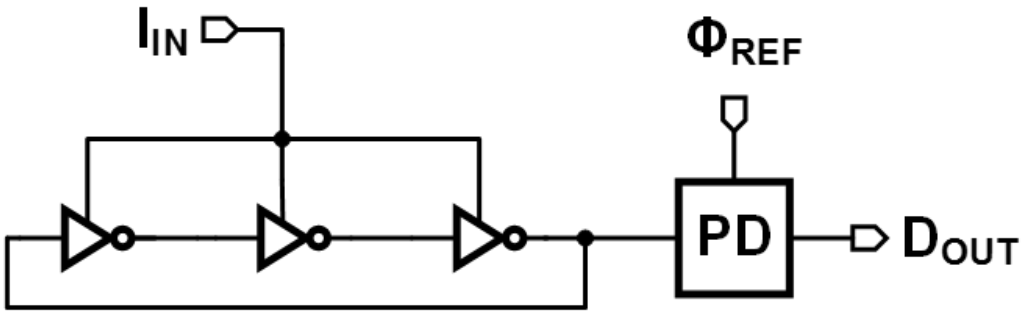

Figure 2.8: Black box interpretation of a phase detector.

approximate the distortion. Namely, we fit our simulation to a third-order polynomial of the form $\sum \alpha_i x^i$. Then we assume a sinusoidal input signal, and are able to calculate the second- and third-order distortion terms as:

$$\mathrm{HD}_2 = \frac{\alpha_2/2}{\alpha_1 + 3\alpha_3/4} \tag{2.12}$$

$$\mathrm{HD}_3 = \frac{\alpha_3/4}{\alpha_1 + 3\alpha_3/4} \tag{2.13}$$

For the example curve in Figure 2.7, we have approximately $\mathrm{HD}_2 = -13$ dB and $\mathrm{HD}_3 = -57$ dB. We tend to ignore the very high second harmonic, as we assume the final design will be used in a differential system, which safely negates the majority of even harmonics.

2.2.2 Phase Detector

As mentioned before, a VCO has ideal integration in the phase domain, but unfortunately, classical circuits measure changes in the voltage and current domains. The next critical block in our design is a system that can measure a phase change, and turn it into a usable voltage signal. We begin by analyzing the black box model in Figure 2.8.

We have our black box phase detector (PD), which measures a phase difference and produces some digital output D_{OUT}. We have some current-controlled oscillator on the left, which produces a Φ_{VCO}, and a reference phase Φ_{REF}. Much like how voltages are always defined as a potential relative to ground, phases must also be defined relative to some reference. There-

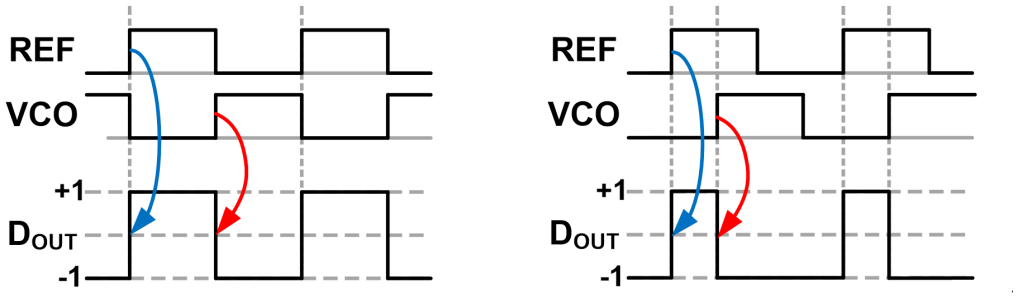

Figure 2.9: Example PD waveforms.

fore the measurement of interest will be the difference between the two phase inputs. The PD block will perform a very simple function: on the rising edge of Φ_{VCO}, D_{OUT} will output a high level, and on the rising edge of Φ_{REF}, D_{OUT} will output a low level.

In order to see how this works, consider the two cases shown in Figure 2.9. On the left-hand side, Φ_{VCO} and Φ_{REF} are delayed by π radians. This will produce an output signal with a duty cycle of 50%. On the right-hand side, the delay between the inputs is halved to $\pi/2$ radians. The corresponding output duty cycle has also shrunk linearly to only 25%. Since the output duty cycle changes linearly with the input phase difference, we can say the transfer function is a linear gain K_{PD} V/rad. In order to change the output from a PWM-voltage to a constant voltage, one can use a low-pass filter with bandwidth well below the VCO's output frequency.

Lastly, the voltage and phase comparison is not entirely apt. While both are continuous and measured as difference, phase signals are bounded and modular. We can visualize this effect by looking at the waveforms shown in Figure 2.10. We have one reference rising edge highlighted in blue, and two separate VCO signals. In particular, we highlight one rising edge on each of the two VCO signals. In the time domain, it is clear that the second VCO is delayed by one period, or 2π radians. To a phase detector however, these signals will appear exactly the same. This leads to the steady-state transfer function shown on the right in Figure 2.10, which is linear but modular around 2π intervals. This is fairly intuitive, as without the specific highlighting, both VCO signals would look exactly the same.

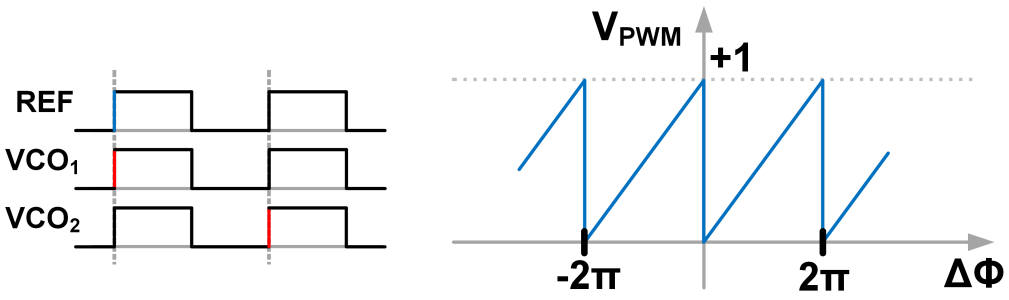

Figure 2.10: Two example waveforms showing the modular nonlinearity of a phase detector.

CHAPTER 3

INDUCTOR DESIGN

3.1 System Modeling

We begin by looking at the classical gyrator system model and implementing the integrator using the time-domain techniques previously discussed. Figure 3.1 shows the implementation where we directly replace the integrator with two VCOs and a PD. The second input to the PD is a VCO with input bias voltage V_B. For now, we will assume that V_B is a constant voltage, which will produce a constant frequency output. The effective inductance given by this circuit can be found as:

$$Z_L = \frac{V_{in}}{I_L} = \frac{s}{K_{VCO}K_{PD}G_M} \tag{3.1}$$

$$\hat{L} = \frac{1}{K_{VCO}K_{PD}G_M} \tag{3.2}$$

We first verified this functionality by creating and simulating Verilog testbenches. The VCOs were ideal models, with no distortion or saturation. We used a simple XOR gate as our phase detector and an ideal G_M cell, im-

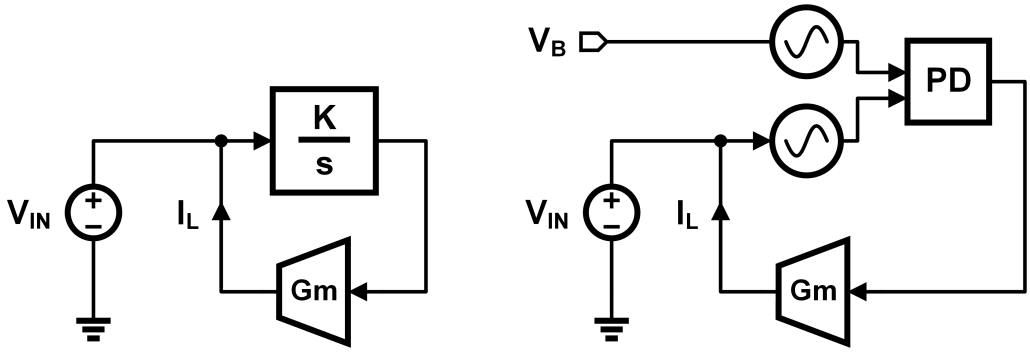

Figure 3.1: Left, the system level model of an inductor. Right, the time-domain implementation of that model.

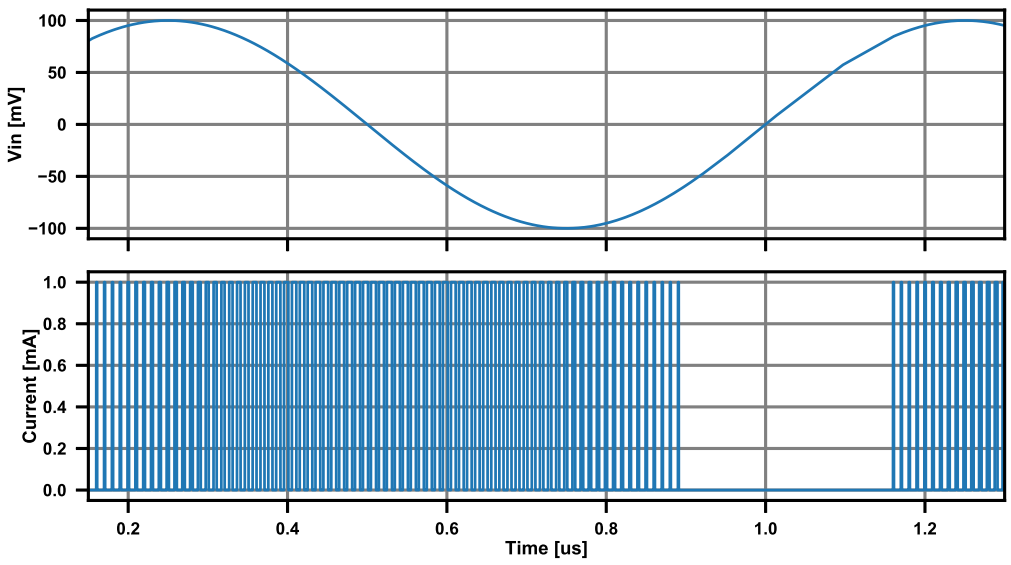

Figure 3.2: Input voltage and output current of the one-phase, XOR-PD system.

plemented by a voltage-controlled current source (VCCS). The input was an ideal sinusoidal voltage source, with an amplitude of 100 mV and a frequency of 1 MHz. The VCOs operated at 50 MHz with a gain $K_{VCO} = 10$ MHz/V. The VCCS had a gain of 1 mA/V.

The results are shown in Figure 3.2. With a sinusoidal input, an inductor should draw a sinusoidal current, with the same frequency and a phase offset of $\pi/2$. That output is indeed there, but it is instead shown through the modulation of the VCO. The phase detector will produce an output:

$$V_{PD} = \int V_{in}(\tau)\mathrm{d}\tau = \frac{V_{in}K_{VCO}K_{PD}}{s} \tag{3.3}$$

We can view this as a combination of pulse-frequency modulation (PFM) and pulse-width modulation (PWM). The difference between the modes is subtle but important to note. PWM is a process in which one modulates an input to a fixed frequency, with a variable pulse width relative to the input amplitude. PFM is a process in which one produces a constant duration pulse, but the frequency of the pulses changes relative to the input amplitude.

Performing analysis on both operating modes simultaneously is difficult, so we make a few assumptions. If the VCOs operating frequency, F_{VCO}, is significantly higher than the change in VCO frequency $K_{VCO} \cdot V_{IN}$, then

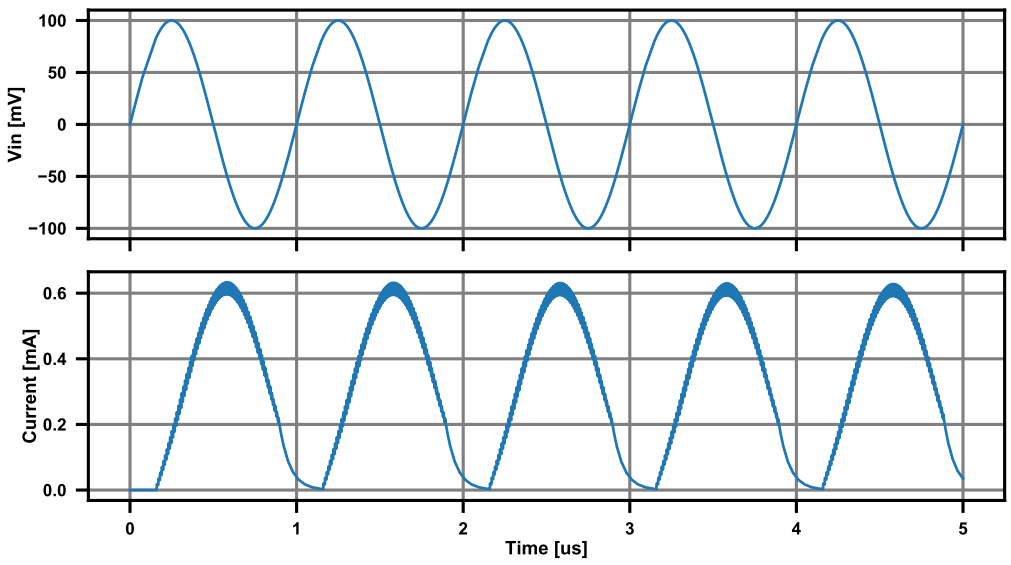

Figure 3.3: Input voltage and output current of the one-phase, XOR-PD system with an RC filter.

we assume that the modulation is solely PWM, as the output frequency is relatively constant with small variation. We can validate this assumption by passing our generated signal through a low-pass filter with a cut-off frequency in between the input signal's bandwidth and the F_{VCO}. This was implemented using a first-order RC filter with $f_c = 5$ MHz, and we get the results shown in Figure 3.3.

Of course, the irony of using an RC filter to solve this problem should not be lost on anyone. While it may work well for the simulation, scaling this solution to lower frequencies results in our circling back to the original problem. We do not actually need to create an incredibly low-bandwidth low-pass filter if we separate the modulation frequency from the input frequency by a significant margin. The further we separate the VCO tones from that transconductor's bandwidth, the more the dynamic range will improve.

One way to do this is to move from a one-level modulation to a multi-level output. A common way to implement this is to have M modulators in parallel, each with phase-offset update rates. The output signal is the sum of all parallel modulators. This effect is shown in Figure 3.4. Assuming each G_M cell is the same as before, the total system transconductance is now $M \cdot G_M$. This effect will increase the gain by a factor of M, so our new inductance can be found as:

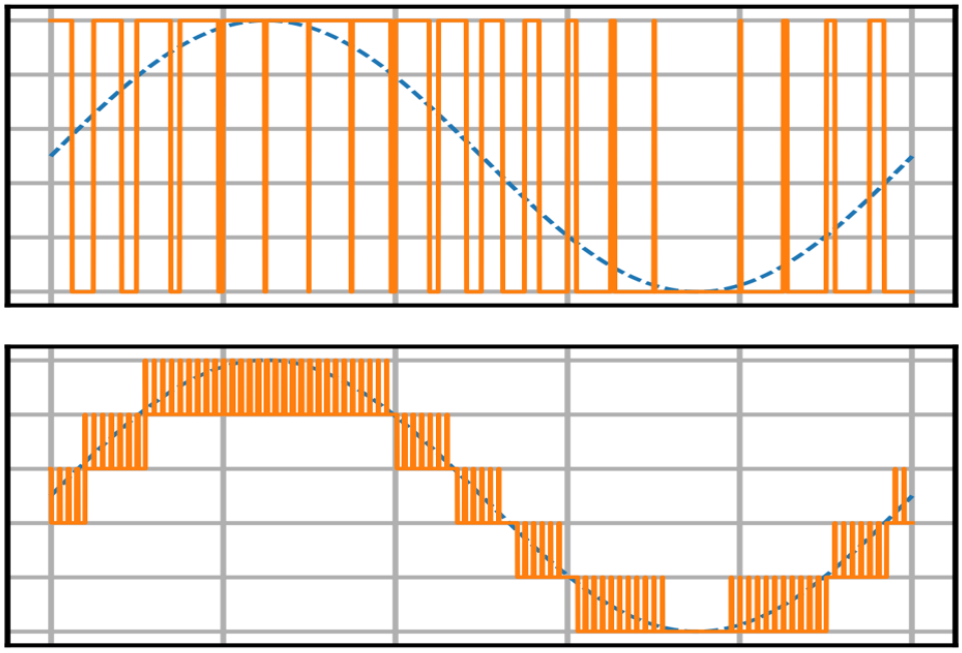

Figure 3.4: Above, the output of a one-bit PWM modulator. Below, the output of a four-level modulator.

$$\hat{L} = \frac{1}{K_{VCO} K_{PD} G_M M} \tag{3.4}$$

If the parallel modulators are evenly spaced, the effective modulation rate will also increase by M. The benefits of this are shown in Fig. 3.5. If we assume the transconductance cells have some bandwidth ω_{GM}, we want to maximize the suppression we get by increasing the modulation frequency ω_{VCO}. However, a faster update rate means more power draw for the rest of the system. Ring oscillator based VCOs naturally lend themselves to this multi-level update, as we can simply tap out each sub-phase after each inverter in the ring oscillator. An example of this is shown in Fig. 3.6. If proper care is taken during layout, we can assume that each of the phases are evenly spaced. If we take out 10 phases, for example, the modulation tones will move from ω_{VCO} to $10 \cdot \omega_{VCO}$. If we design $\omega_{VCO} > \omega_{GM}$ initially, then this will result in a 20 dB improvement in the modulation tone suppression. This result, however, only holds if the phase-updates are evenly spaced.

For the purpose of this work, merely comparing the Fourier spectra of the modulated signals is sufficient for estimating parameters such as distortion,

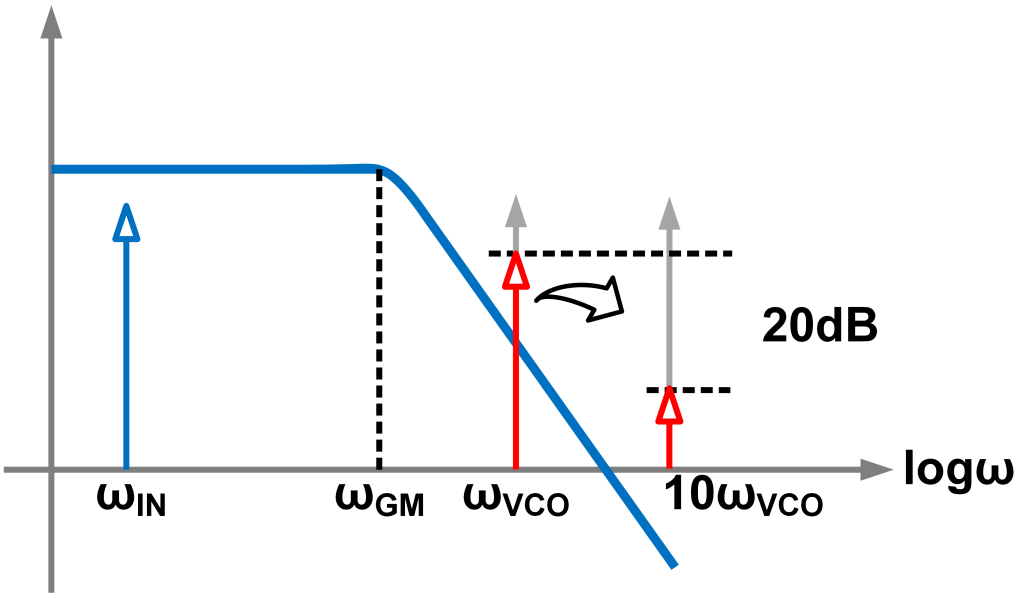

Figure 3.5: Suppression of modulation tones due to multi-level pulse-width modulation.

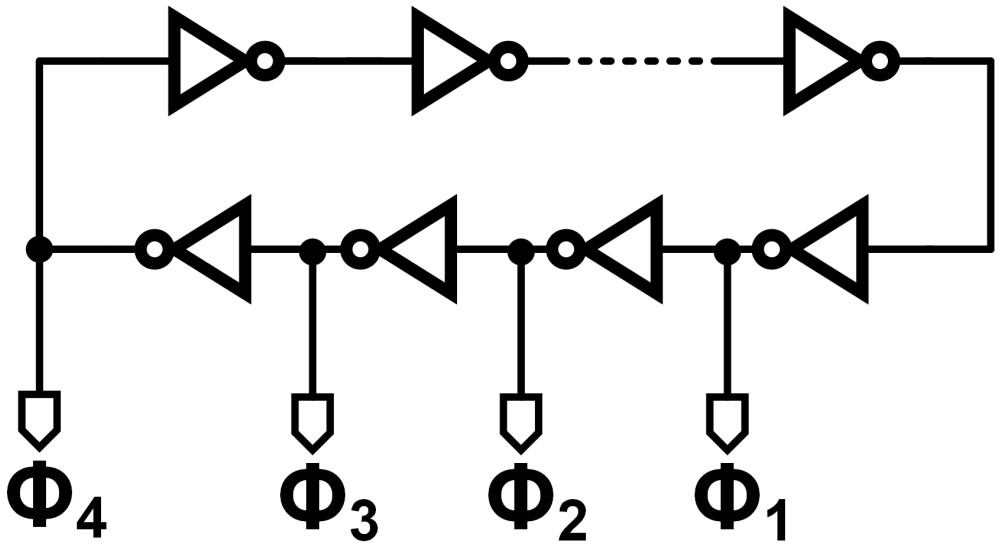

Figure 3.6: An example of using multiple phases in a ring oscillator.

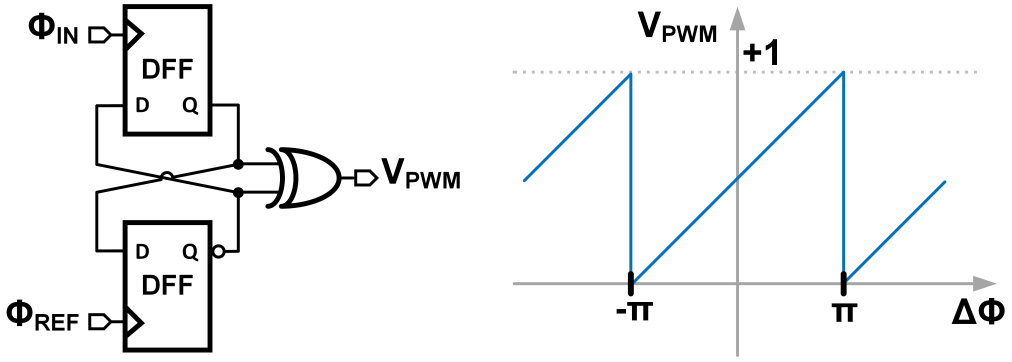

Figure 3.7: Two-state PD and its transfer function.

dynamic range, and SNR. The exact details involve decomposing the multi-level modulation signals into a double Fourier series expression, which is slightly beyond the scope of this work. Interested readers are encouraged to see [15, 16] for more information.

The next step is to increase the linear phase detection range. A larger phase range will correspond to a larger maximum input swing without distortion. We moved from the previous XOR-PD to the two-state PD, shown in Figure 3.7. While it doubles the maximum input phase difference from π to 2π, it has another benefit as well. When the two input signals are locked with zero phase error ($\Delta\Phi = 0$), the two-state PD has a normalized average output voltage of 0.5 V. The XOR PD, on the other hand, has a normalized average output voltage of 0 V in locked state.

Earlier, we assumed that the transconductor would act as a natural low-pass filter for the modulation tones. This is only useful if the rise and fall times of the modulator pulses are very small compared to the pulse widths. If the pulse widths are also filtered out, we have lost the integration information in the filtering. This will ultimately end up limiting our dynamic range. When adding a small input phase difference, the XOR PD will create very narrow pulses that will get filtered out. The two-state PD, however, will go from a pulse width with 50% duty cycle to a slightly duty cycle, and therefore the integration information is maintained.

Lastly, we discuss the modeling of the transconductor. Unlike previous active inductors shown Section 2.1, linearity is not a concern with our design. Since each transconductor's input comes from a phase detector, the output should have only two values: $+g_m V_{DD}$ or $-g_m V_{DD}$. This binary system is perfectly linear, as any two points make a line. The only parameter that

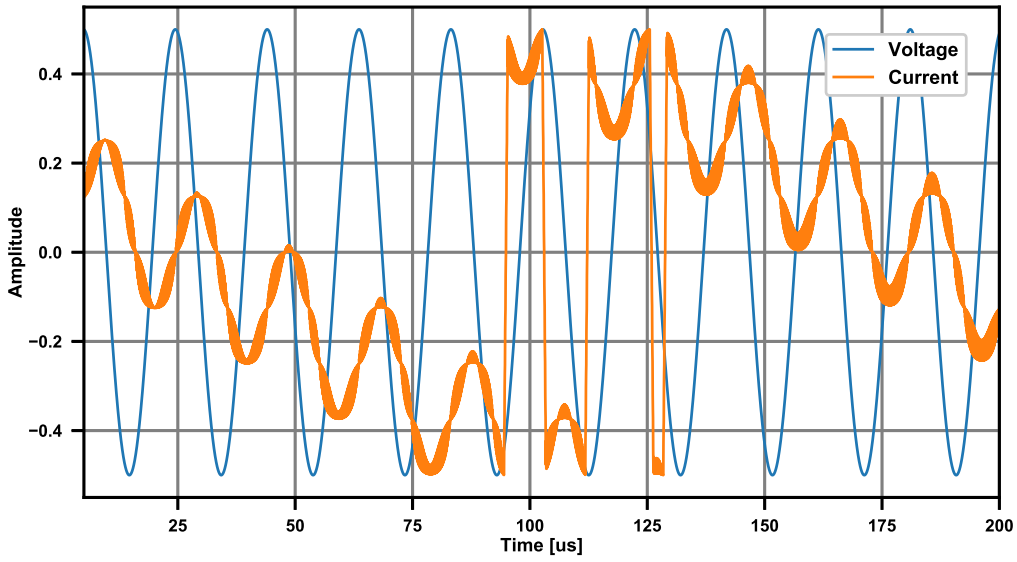

Figure 3.8: Normalized input voltage and current waveforms in the presence of an input offset.

matters is the bandwidth, which as discussed before, should be in between the input bandwidth and modulation frequency.

3.2 Stability

Despite the appearance of a closed loop, this circuit is still operating in open loop. The VCO input is voltage-driven, and the transconductor output is current-driving, so there is no feedback between the two. Due to the infinite gain of the VCO, any small input offset will be integrated to infinity, creating a saturated system. As an example, we simulate the system with a 0.5% frequency offset between the input and reference VCOs. The output of this simulation is shown in Figure 3.8. We can analyze the results by breaking the resulting current into two parts: the offset induced portion and the input sinusoid. The sinusoidal voltage is successfully integrated, producing the cosine we expect. The input offset is continuously integrated, producing a linear function with slope $K_{VCO}K_{PD}$.

Any phase detector has a finite linear region, at which point it wraps back around, creating a distortion effect. If we define the linear phase range as $\Delta\Phi_{\max}$, we can calculate the limit of the input swing. In the case of a two-

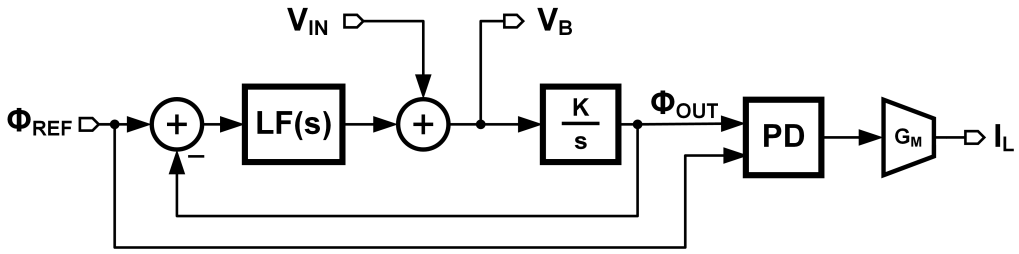

Figure 3.9: System diagram with added PLL feedback.

state PD, $\Delta\Phi_{\max} = 2\pi$. Consider the case of a sinusoidal signal with peak-to-peak magnitude V_{in} and frequency F_{in}. This will generate a frequency change of $K_{VCO}V_{in}$. We can then calculate the corresponding phase change:

$$\Delta\Phi = \frac{V_{in}K_{VCO}}{F_{in}} \leq \Delta\Phi_{\max} \tag{3.5}$$

This sets a lower bound on our input frequency as a function of the input voltage and VCO gain. Any input frequency below that will cause significant distortion in the output current spectrum. However, even DC frequency offset, as shown in 3.8, will cause this same problem. In order to solve it, we add a low-frequency phase-locked loop (PLL) to track the low-frequency input voltage, modeled in Figure 3.9. The goal of this loop is to lock the two reference phases to each other, with perfect low-frequency tracking.

We can model the system using superposition. $LF(s)$ represents the loop filter transfer function in the PLL:

$$\left.\frac{I_L}{\Phi_{\text{Ref}}}\right|_{V_{in}=0} = \frac{LF(s)K/s}{1+LF(s)K/s}G_M K_{PD} \tag{3.6}$$

$$\left.\frac{I_L}{V_{in}}\right|_{\Phi_{\text{Ref}}=0} = \frac{K/s}{1+LF(s)K/s}G_M K_{PD} = \frac{I_L}{\Phi_{\text{Ref}}}\cdot\frac{1}{LF(s)} \tag{3.7}$$

There are several ways to validate this model. First, when the PLL is disabled, the system should behave like an open loop integrator with infinite DC gain. Taking the limit as $LF(s) \to 0$, this still holds. Using the parameters: $LF(s) = 1.5$ mV/rad, $K_{VCO} = 50$ MHz/V, and $G_M K_{PD} = 16$ μA/rad, we get the results shown in Figure 3.10.

The DC gain of our system is $G_M K_{PD}/LF(s)$. The closed loop bandwidth is $LF(s)K_{VCO}$. Adding the PLL solves our DC tracking problem, illustrated

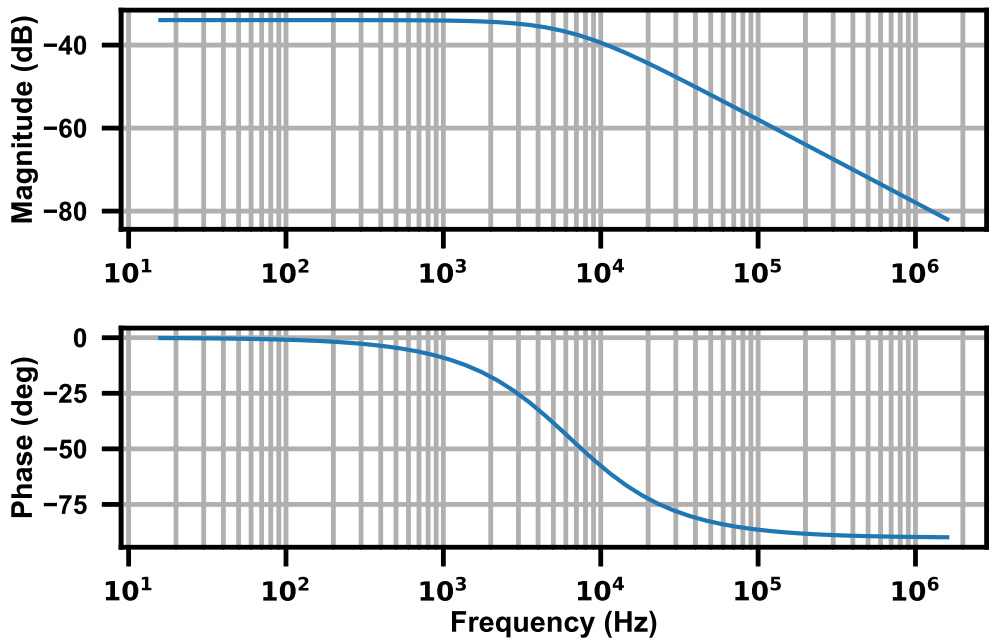

Figure 3.10: Bode plot of inductor with added PLL.

in Figure 3.11. Any input with a frequency inside the PLL's bandwidth will be tracked by the PLL, creating a fixed output current amplitude. We also no longer need to provide an external bias for the reference VCO, as the loop now generates that for us.

In terms of modeling, we need a way to represent the effect of the PLL in the overall inductor model. As mentioned before, all low frequency inputs will be tracked by the loop and produce a constant output. This output will have magnitude $I_L = V_{in}G_M K_{PD}/LF(s)$. This manifests itself as the

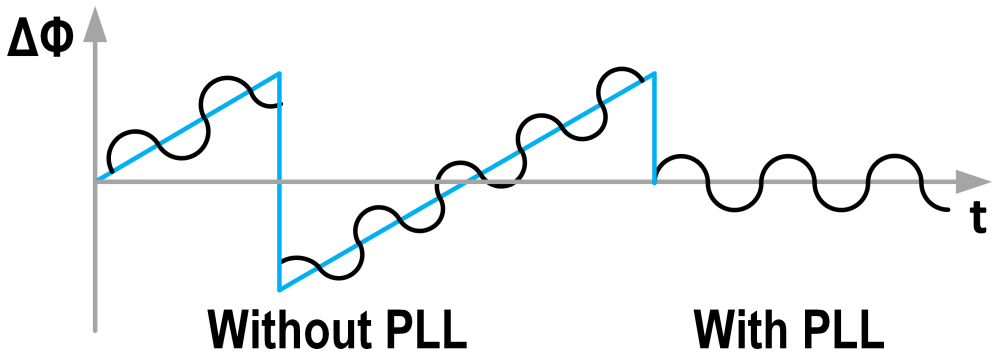

Figure 3.11: Effect of using a low-frequency tracking PLL in the inductor.

series resistance associated with active inductors. Beyond that frequency, we see the inductive behavior, with slope of -20 dB per decade indicating the single pole behavior. With this, we can find an analytical expression for the low-frequency Q of the inductor:

$$Q_1 = w\frac{\frac{1}{K_{VCO}K_{PD}G_M}}{\frac{LF(s)}{G_M K_{PD}}} = \frac{w}{LF(s)K_{VCO}} \tag{3.8}$$

In reality, the high-frequency Q will be limited by the output impedance of the transconductor, as mentioned in Section 2.1. We can rewrite Eq. 2.6 in terms of our model:

$$Q_2 = \frac{R_o}{R_o + \frac{LF(s)}{G_M K_{PD}}\left[1 + \left(\frac{w}{K_{VCO}LF(s)}\right)^2\right]} \tag{3.9}$$

Lastly, our high frequency Q_3 is:

$$Q_3 = 1 - (LF(s))^2 C_p K_{VCO} K_{PD} G_M M - \frac{w^2 C_p}{K_{VCO}K_{PD}G_M M} \tag{3.10}$$

We will estimate the parasitic output capacitance of the transconductor $C_p = 15$ fF, and the output impedance $R_o = 50$ kΩ. The results are shown in Figure 3.12. Similar to previous active inductors, our low-frequency Q is limited by the series resistance, and our high frequency Q is limited by the finite output impedance. Q_3 has very little impact on this design, and does not show up until extremely high frequencies. Unlike the previous active inductor models, the series resistance and the inductance are both functions of the same variable (K_{VCO}), so we cannot optimize this equation as easily for maximum Q.

Figure 3.13 shows the Q when sweeping the K_{VCO} values. The curve moves to the right with higher K_{VCO}, but otherwise remains the same. Additionally, the maximum value of Q does not change. This is similar to sweeping the forward G_M value in the gyrator. Similarly, we can design our inductor for some optimum frequency range to maximize the available Q.

Next, we sweep the PLL bandwidth as shown in Figure 3.14. Here we see a more striking change, as the peak Q changes in both magnitude and frequency with the changing bandwidth. As PLL bandwidth goes to zero, the R_s term goes to zero as well. No resistance means we have a "true"

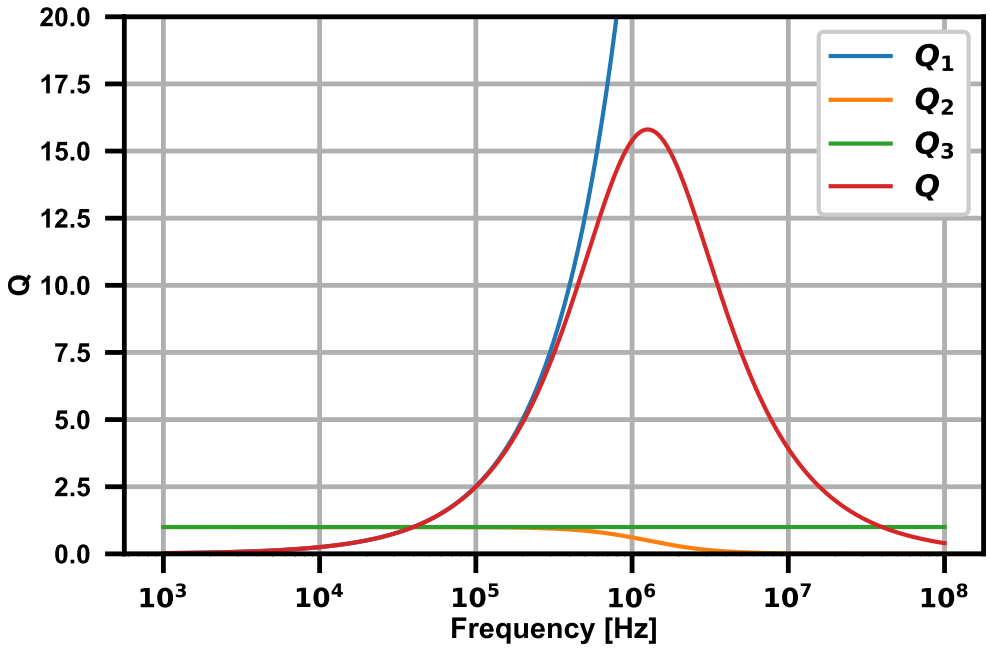

Figure 3.12: Q of the time-based inductor over frequency.

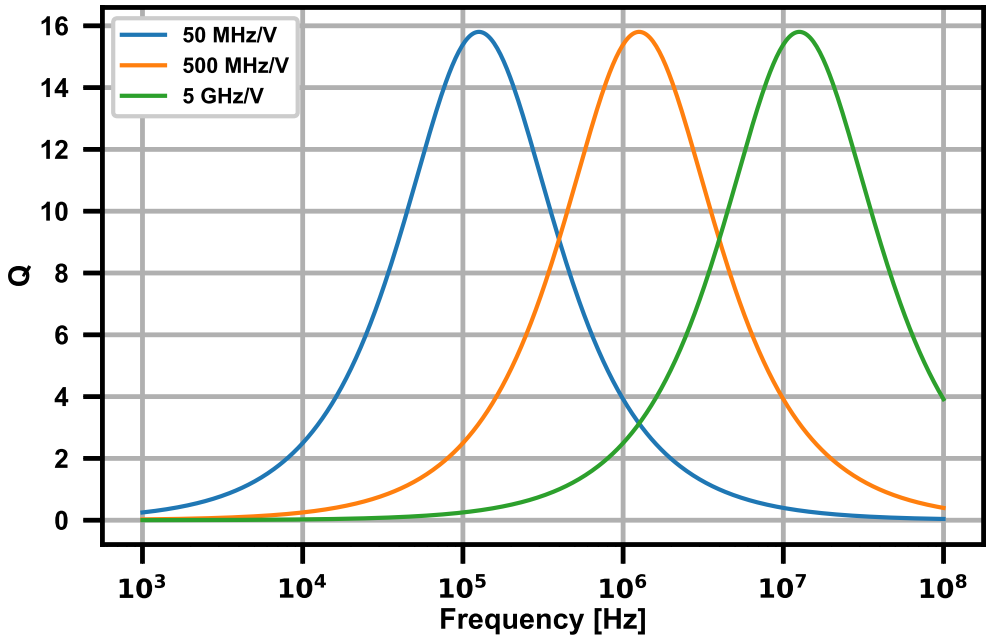

Figure 3.13: Simulated Q values across various K_{VCO} values.

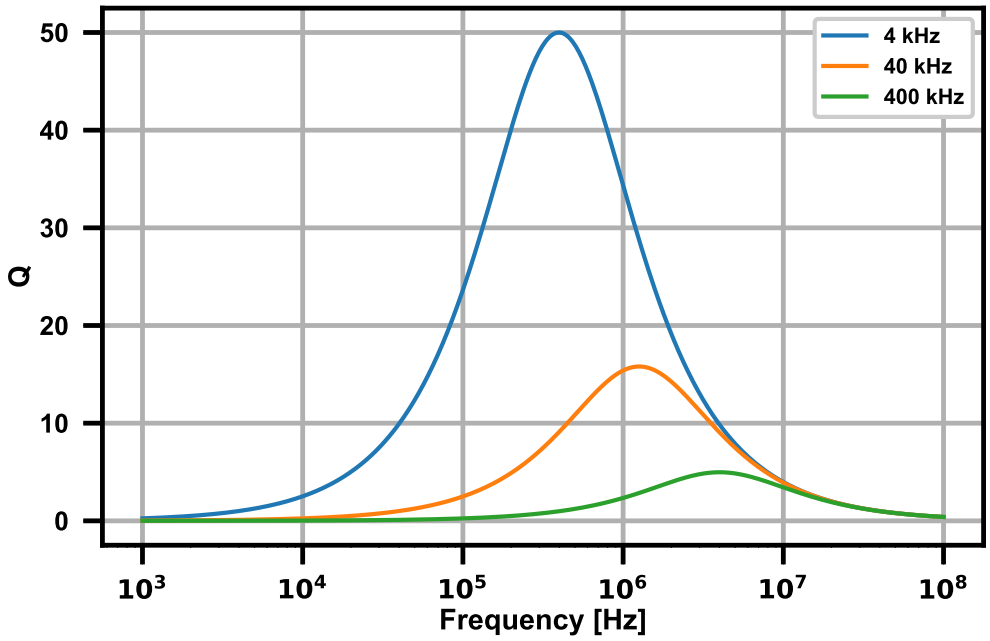

Figure 3.14: Simulated Q values across various PLL bandwidth values.

inductor, which should have infinite Q across all frequencies. Therefore, it is desirable to keep the PLL bandwidth as low as possible. Assuming the PLL bandwidth is finite and non-zero however, it is ideal to set K_{VCO} at the correct value to maximize Q in the desired operating range.

In summary, we have two conflicting goals. For stability, we require the PLL bandwidth to be sufficiently high in order to prevent any phase-induced distortion in the output current spectrum. In order to maximize Q, we require as low a bandwidth as possible, which will minimize the series resistance. As a design compromise, we aim to set the PLL bandwidth to be the lower bound of the input frequency range as calculated in Eq. 3.5.

3.3 Circuit Implementation

Here we discuss the detailed schematics of each of the primary building blocks of the time-based inductor. The full block diagram is shown in Figure 3.15. All circuits shown in this section were designed using the TSMC 65 nm process, with $V_{DD} = 0.7$ unless otherwise mentioned. Similarly, device dimensions are given as width by length (W/L).

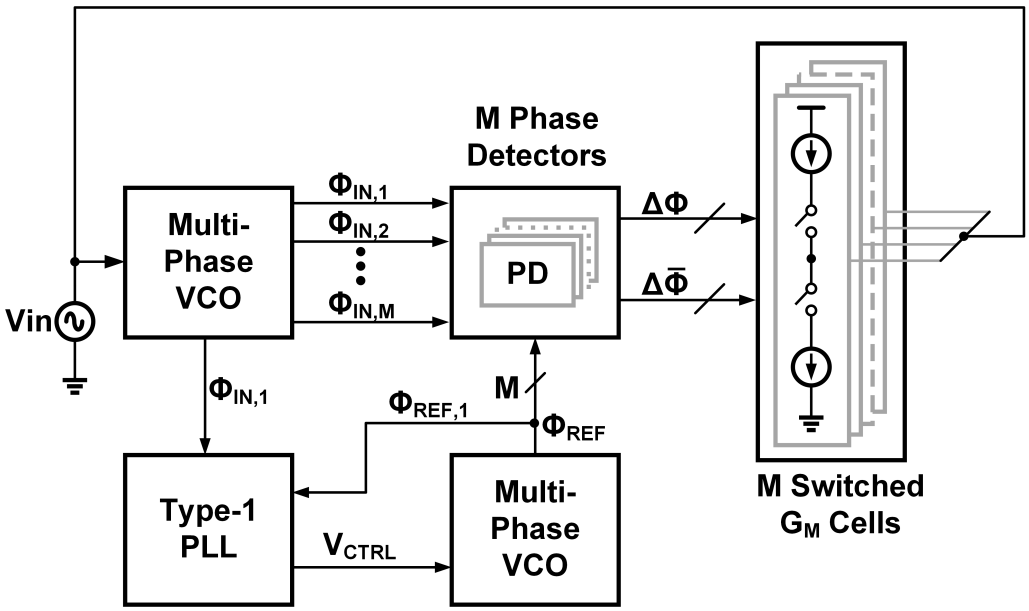

Figure 3.15: Inductor block diagram with added PLL.

3.3.1 VCO + PD

The VCO was designed as a 33-stage single-ended ring oscillator. A single-ended architecture was chosen to minimize power compared to a differential architecture. Such a high number of delays stages was used as we wanted to reduce the modulation distortion by at least -20 dB. A -20 dB change requires moving the modulation tone upwards by at least a decade, which in turn requires at least 10 phase outputs. A single-ended architecture requires an odd number of delay stages, meaning the minimum number of stages is 11. Furthermore, to reduce any layout mismatches between cells, we opted to make each delay stage three inverters, which results in the 33-stage design.

The VCO bias network was designed using two PMOS branches with three two-bit controls. One branch provided bias current for the nominal frequency operating point. The other branch provided bias current proportional to the input magnitude. Each control activated current mirrors to increase the strength of the respective branch by up to a factor of eight. The last control knob changed the input-proportional current by a factor of ten, to verify operation over a wide frequency range.

The schematic of the delay cell is shown in Figure 3.16. Each of the delay inverters was sized at 8 μm/250 nm. Since we use a PMOS bias, the output swing of the inverter will be restricted to V_{CTRL}. In order to restore the

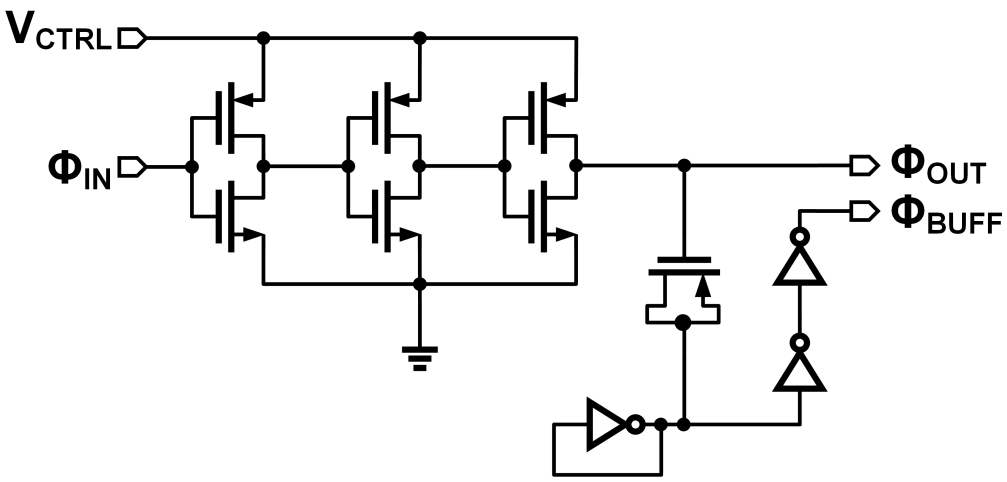

Figure 3.16: Delay cell schematic.

output to full swing, we AC-coupled the signal through a MOSCAP sized at 2 μm/3.6 μm. The output DC operating point is set by a self-tied inverter sized at 150 nm/4 μm. The self-tied inverter used high-V_{th} devices in order to reduce power. This was then passed to two buffers (sized at 200 nm/60 nm and 400 nm/60 nm) to bring the phase difference back up to full scale. This output is referred to as Φ_{Buff}.

VCO simulations are shown in Figure 3.17. We show the output frequency range and the K_{VCO} range for the low and high settings. In Figure 3.18 the phase noise is plotted for the VCO operating at 65 MHz and 175 MHz. The phase noise at a 1 MHz offset is -103 dBc/Hz and -115 dBc/Hz, respectively.

The phase detector was implemented with the same schematic shown in Figure 3.7. In order to aid with driving the large input of the switched-G_M cells, several buffers were added to the output. Additionally, an inverting output path was added to convert the single-ended XOR output of the two-state PD to a differential signal.

3.3.2 Switched-G_M Cell

The transconductor cell was implemented using a mirrored differential amplifier, shown in Figure 3.19. The current sources were controlled using two four-bit digitally programmable current mirrors, each of which was biased externally. The transconductor cell was designed to provide up to 50 μA at maximum. All of the 11 cells (one for each $\Delta\Phi_i$) had their outputs tied

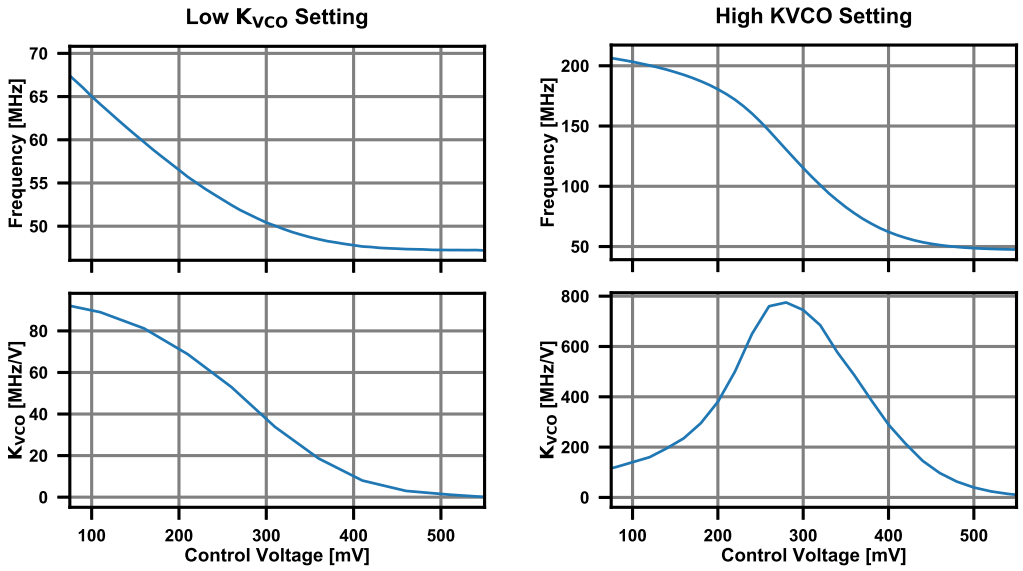

Figure 3.17: VCO output frequency and gain across two settings.

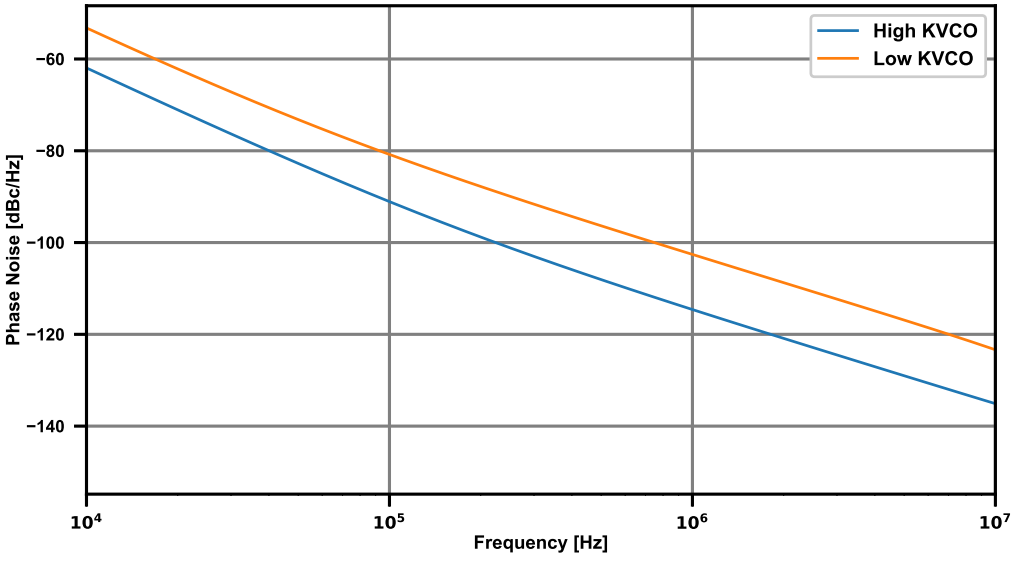

Figure 3.18: Phase noise simulations of the VCO.

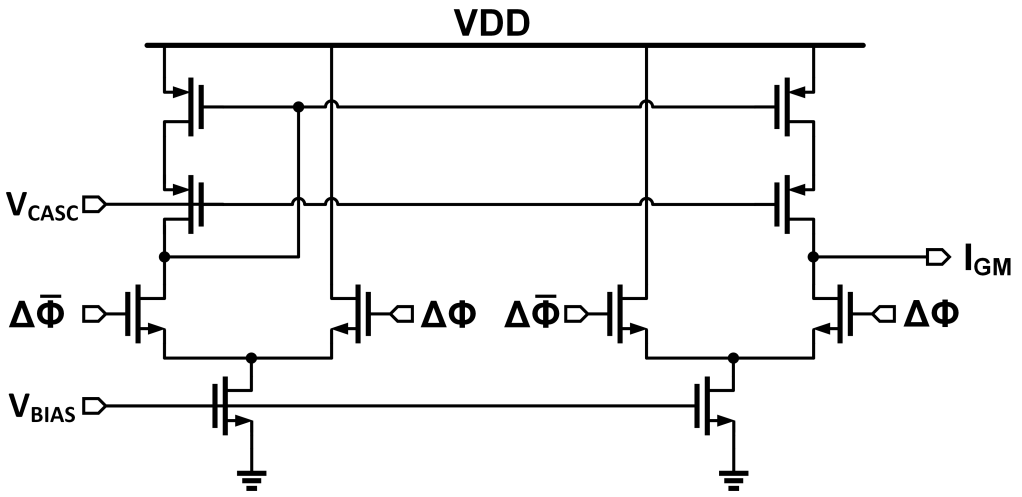

Figure 3.19: Switched G_M schematic.

together to form a current summer. The NMOS current sources were sized to 5 μm/1 μm and the input differential pair were sized to 2 μm/ 60 nm. The high-side PMOS were sized to 2 μm/80 nm, and the low-side cascode devices were sized the same.

The cascode mirror architecture has several benefits over the traditional on-off current-steering design. First, by mirroring the currents between both pairs, the matching between the source and sink currents is greatly increased. Second, since both branches are fully differential with binary inputs, the drain node of both current sources should be at a constant voltage, since it is ideally always providing the same current. This helps minimize any current glitching during transitions. Lastly, it has an inherent internal pole due to the mirroring. This is another knob we can use to optimize the bandwidth of the transconductor, which should ideally lie between the input frequency and the VCO frequency, as discussed in 3.1. In order to adapt this circuit to an even lower supply, one would likely need to reduce the headroom by removing one of the PMOS cascode devices.

3.3.3 PLL

The schematic for the PLL is shown in Figure 3.20. A type-I PLL was chosen for several reasons. First, we wanted to minimize the area penalty as much as possible. While a type-II PLL would give better performance, it comes at the cost of using large capacitors in the loop filter, negating any

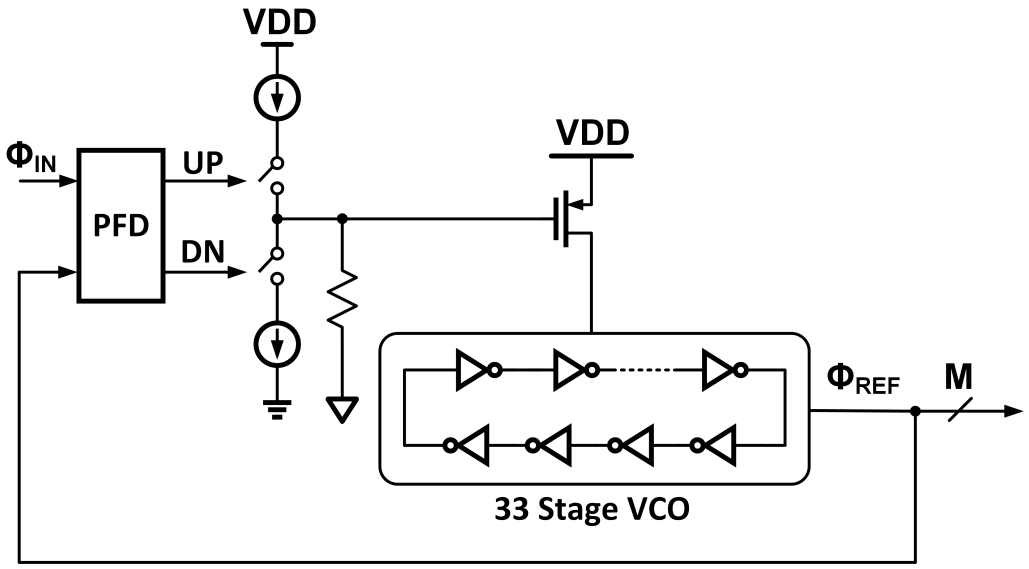

Figure 3.20: PLL schematic.

benefits of technology scaling. Second, we do not need true phase tracking that a type-II PLL provides. We are instead interested in creating a low-bandwidth frequency locked loop, which a type-I PLL does. A type-I PLL locks with some static phase error, which will create a fixed output. As mentioned previously in Section 3.2, this manifests itself as the series resistance of the inductor.

A DFF-based PFD was used with a reset time extension. The schematic is shown in Figure 3.21. Since the inductor operates at relatively low frequencies, there was no need to optimize the PFD for speed. The reset branch was designed to be weak to compensate for any dead-zone issues when locking the PLL. The charge pump circuit was the same design as the switched-G_M, shown in Figure 3.19. The only difference was the inputs: the left-side differential amplifier had inputs UP and $\overline{UP}$, and the right-side differential amplifier had inputs DN and $\overline{DN}$, respectively. The charge pump was designed to provide a 1 μA current nominally, but more current could be mirrored with external controls. The loop filter resistance was set to a fixed value of approximately 1.5 kΩ.

Since a type-I PLL was used, the loop filter consists of solely a resistor, with one end receiving the charge pump output current and the other end tied to a DC bias. The maximum tuning range of the PLL will be $i_{CP} \cdot R \cdot K_{VCO}$. In order to account for mismatches in fabrication between the reference and

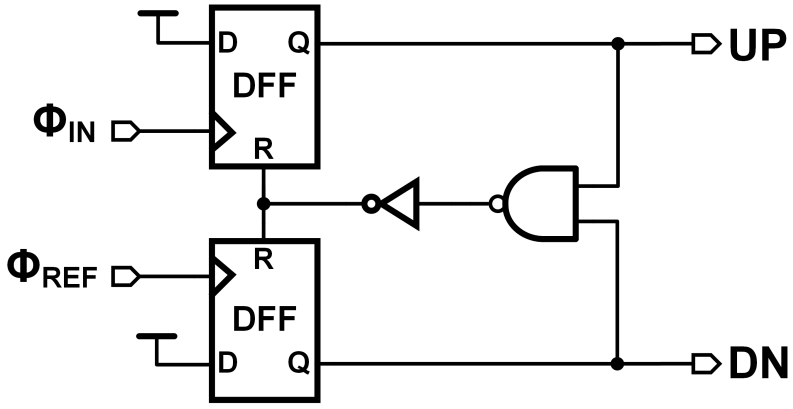

Figure 3.21: PFD schematic.

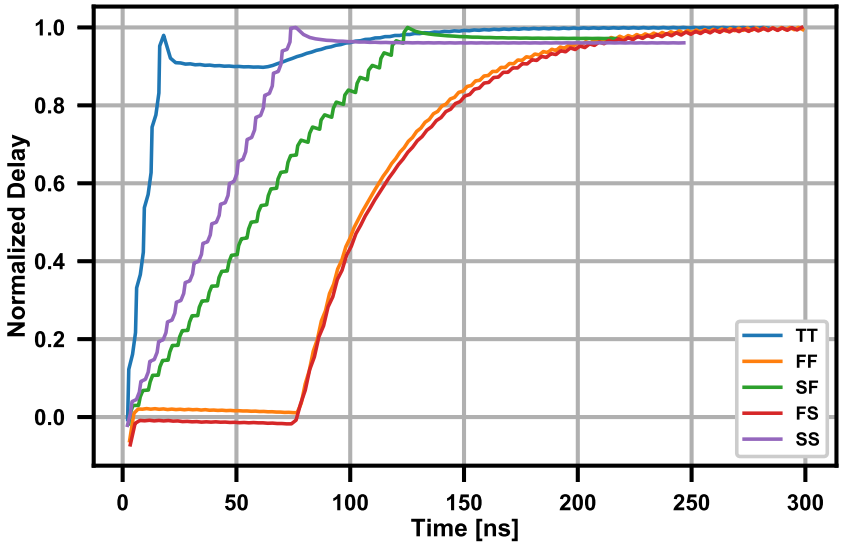

Figure 3.22: PLL step response across process corners.

input VCOs, we use an external coarse tuning to get the DC voltage where the both input and reference frequencies are close, then we activate the loop to finely calibrate any offsets.

Simulated results are shown in Figure 3.22. We tested the PLL across five process corners to verify performance, with all corners tested at a nominal temperature (25° C). We then plot a normalized delay (as each corner has a different output frequency) and verify that it locks under all corners.

CHAPTER 4

MEASUREMENTS

This design was fabricated using the TSMC 65 nm CMOS LP process, and the die micrograph is shown in Figure 4.1. Two inductors were fabricated to verify differential performance, and each inductor occupies an area of 0.017 mm^2. When each VCO was operating at a frequency of 200 MHz, the total power draw was 529 μW. The majority of that biasing power comes from the VCOs and PLL, drawing 514 μW. The digital control circuitry consumes 5 μW, and the G_M cell biasing consumes 10 μW.

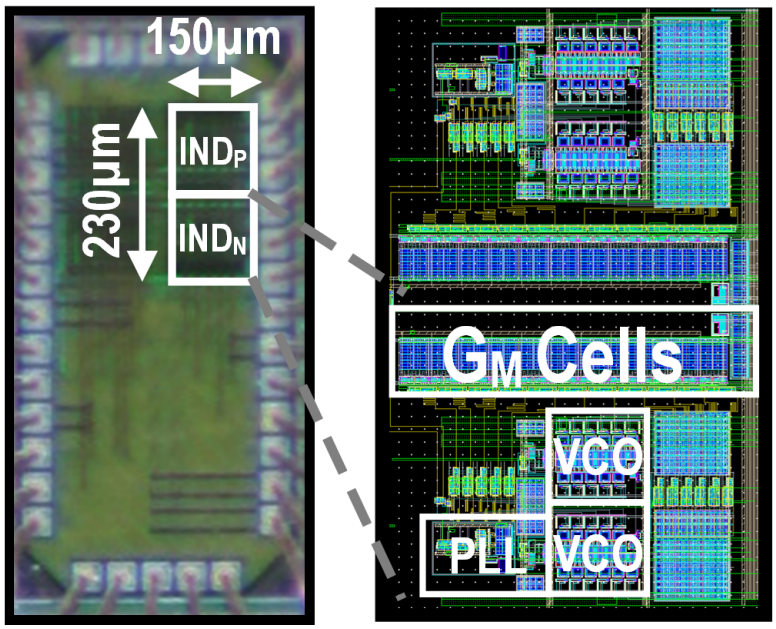

Figure 4.1: Die micrograph and approximate layout representation.

Since the inductor was designed as a one-port inductor, with one floating port and the other connected to ground, there were a limited number of test setups available. We opted to use a high-pass filter setup, as shown in Figure 4.2. Our results are shown in Figure 4.3.

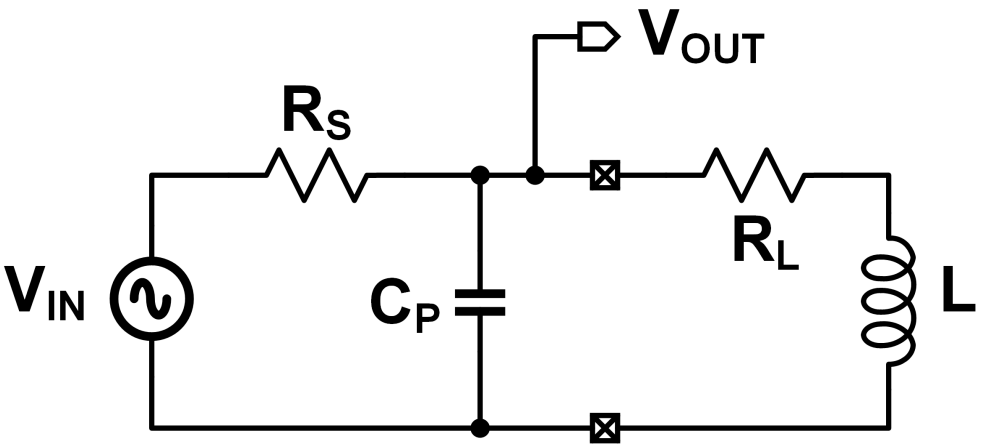

Figure 4.2: Schematic of the test setup.

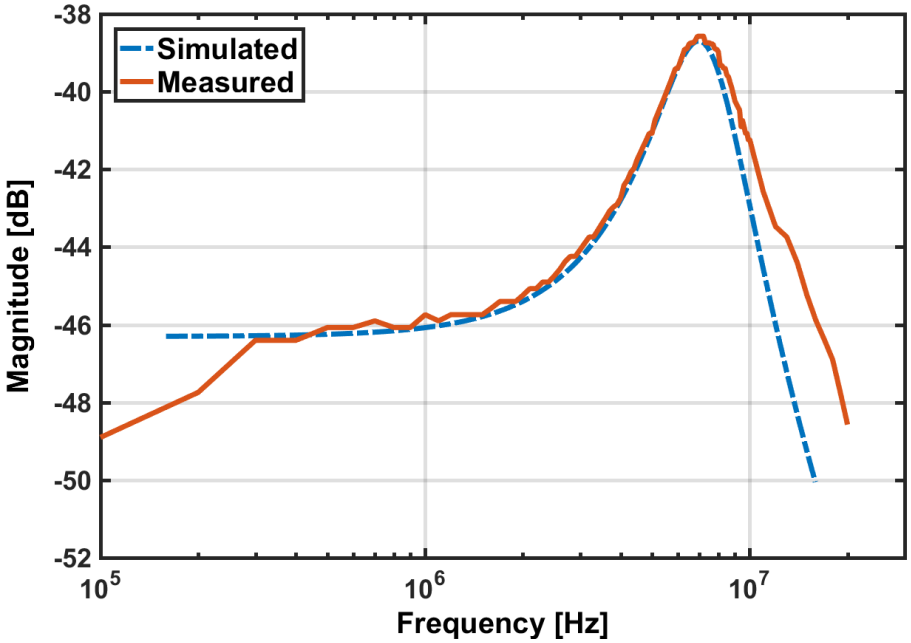

Figure 4.3: Inductor measurement results.

In our test setup, with no parasitic capacitance, we would ideally have:

$$A_v(s) = \frac{V_o}{V_i} = \frac{R_L + sL}{R_S + R_L + sL} \tag{4.1}$$

This is a high-pass system, with one zero at $f_z = (R_L + R_s)/L$ and one pole at $f_p = R_L/L$. Of course, parasitic capacitance exists in our system, so when we add it to our model, we get the following transfer function:

$$A_v(s) = \frac{sL + R_L}{s^2 LCR_s + s(L + R_L R_S C) + R_L} \tag{4.2}$$

This introduces another pole at frequency $f_{p2} = 1/R_S C$. Our estimates for the parasitic capacitance were around $C_p \approx 100$ fF, which put the second pole high enough to view the entire bandpass response. However, we did not account for the much larger parasitic capacitance of the testing setup and the PCB, which we now estimate to be approximately 15 pF. This means the pole occurs far earlier, which limits the maximum inductive behavior we could see.

Additionally, the series resistance was far higher than expected, as our analysis calculated $R_S = 1.5$ kΩ. This resulted in $Q = 1.6$ when calculated at the maximum inductive point. We estimated our R_S when simulating around low K_{VCO} operation. We believe this effect was ultimately due to larger than expected mismatch between the input and reference VCOs, which resulted in requiring a significantly larger than expected PLL bandwidth to maintain stable behavior.

CHAPTER 5

CONCLUSION

A highly tunable, highly digital active inductor is proposed that employs time-domain signal processing techniques. By using a ring oscillator to integrate the input voltage and a switched transconductor to inject current into the input node, the proposed time-domain gyrator achieves inductive input impedance without using either large resistors or capacitors. The inductor consumes 528 μW of power while operating at a frequency of 180 MHz and realizing an inductance of 500 μH. Due to the fully digital nature, the inductor supply voltage can scale while the input voltage swing remains a significant percentage of the supply. Additionally, realizing the gyrator in this manner makes it significantly more amenable for technology scaling.

5.1 Future Work

Of course, no useful research should ever be considered complete. There are many possible areas of interest that can extend from this work, and we would like to detail a few of them here. The most interesting possibility is to implement a differential active inductor. Most gyrator topologies simply duplicate a single gyrator and reverse it to provide a pseudo-differential operation. This is extremely wasteful both area- and power-wise as opposed to a true differential topology. Due to the digital output of the time-based inductor, it is simple to provide both inverting and non-inverting outputs. The ideal implementation would have some sort of differential input to the VCO, so only duplicating the VCOs would not be necessary.

However, that may not provide the true suppression of even order harmonics that is expected in differential systems. The primary source of nonlinearity in our system is the VCO, which ends up being a single-ended system in the previous implementation. If one switched to a fully differential system with

two VCOs, one would need to investigate a common-mode feedback network, as shown in [3]. In order to further reduce the non-linearities, deterministic background calibration techniques such as [17] could be investigated as well.

Lastly, the power-area trade-off when attempting to use an analog type-I PLL was a very difficult design decision. By switching to a digital type-II PLL, hopefully both area and power can be minimized. Furthermore, this should naturally lend itself well to ultra-low bandwidth applications, as the bandwidth can continue to be reduced in a more area-efficient manner compared to an analog type-II PLL.

REFERENCES

[1] T. Charania, A. Opal, and M. Sachdev, "Analysis and Design of On-Chip Decoupling Capacitors," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 21, no. 4, pp. 648–658, Apr. 2013.

[2] W. M. C. Sansen and P. M. V. Peteghem, "An area-efficient approach to the design of very-large time constants in switched-capacitor integrators," *IEEE Journal of Solid-State Circuits*, vol. 19, no. 5, pp. 772–780, Oct. 1984.

[3] B. Drost, M. Talegaonkar, and P. K. Hanumolu, "Analog Filter Design Using Ring Oscillator Integrators," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 12, pp. 3120–3129, Dec. 2012.

[4] M. Z. Straayer and M. H. Perrott, "A 12-Bit, 10-MHz Bandwidth, Continuous-Time ADC With a 5-Bit, 950-MS/s VCO-Based Quantizer," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 4, pp. 805–814, Apr. 2008.

[5] B. Vigraham, J. Kuppambatti, and P. R. Kinget, "Switched-Mode Operational Amplifiers and Their Application to Continuous-Time Filters in Nanoscale CMOS," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 12, pp. 2758–2772, Dec. 2014.

[6] E. Rodriguez-Villegas, A. Yufera, and A. Rueda, "A 1.25-V Micropower Gm-C Filter Based on FGMOS Transistors Operating in Weak Inversion," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 1, pp. 100–111, Jan. 2004.

[7] O. Omeni, E. Rodriguez-Villegas, and C. Toumazou, "A micropower CMOS continuous-time filter with on-chip automatic tuning," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 52, no. 4, pp. 695–705, Apr. 2005.

[8] L. Pylarinos and K. Phang, "Low-voltage programmable g m-C filter for hearing aids using dynamic gate biasing," in *2005 IEEE International Symposium on Circuits and Systems*. IEEE, 2005, pp. 1984–1987.

[9] B. Rumberg and D. W. Graham, "A Low-Power and High-Precision Programmable Analog Filter Bank," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 59, no. 4, pp. 234–238, Apr. 2012.

[10] S.-C. Lee, S.-Y. Lee, and C.-H. Chiang, "0.9 V low-power switched-opamp switched-capacitor bandpass filter for electroneurography acquisition systems," *IET Circuits, Devices & Systems*, vol. 2, no. 2, p. 257, 2008.

[11] T. N. Rao, P. Gary, and R. W. Newcomb, "Equivalent Inductance and Q of a Capacitor-Loaded Gyrator," *IEEE Journal of Solid-State Circuits*, vol. 2, no. 1, pp. 32–33, Mar. 1967.

[12] F. Yuan, *CMOS Active Inductors and Transformers: Principle, Implementation, and Applications*. New York: Springer, 2008, oCLC: ocn213479589.

[13] J. Burghartz and B. Rejaei, "On the design of RF spiral inductors on silicon," *IEEE Transactions on Electron Devices*, vol. 50, no. 3, pp. 718–729, Mar. 2003.

[14] W. Sansen, "Distortion in elementary transistor circuits," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 46, no. 3, pp. 315–325, Mar. 1999.

[15] B. P. McGrath and D. G. Holmes, "An analytical technique for the determination of spectral components of multilevel carrier-based PWM methods," *IEEE Transactions on Industrial Electronics*, vol. 49, no. 4, pp. 847–857, Aug. 2002.

[16] B. McGrath and H. d. T. Mouton, "One-Dimensional Spectral Analysis Techniques for Multilevel PWM Strategies," *IEEE Transactions on Power Electronics*, vol. 31, no. 10, pp. 6910–6919, Oct. 2016.

[17] S. Rao, K. Reddy, B. Young, and P. K. Hanumolu, "A Deterministic Digital Background Calibration Technique for VCO-Based ADCs," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 4, pp. 950–960, Apr. 2014.