

**PHASE NOISE IMPROVEMENT TECHNIQUES FOR MIXED-MODE
PHASE-LOCKED LOOPS IN NANOMETER CMOS**

A Thesis
Presented to
The Academic Faculty

by

Juan Pablo Caram Wigdorsky

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Electrical and Computer Engineering

Georgia Institute of Technology
December 2018

Copyright © 2018 by Juan Pablo Caram Wigdorsky

PHASE NOISE IMPROVEMENT TECHNIQUES FOR MIXED-MODE
PHASE-LOCKED LOOPS IN NANOMETER CMOS

Approved by:

Professor J. Stevenson Kenney, Advisor
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Professor Stephen E. Ralph
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Professor Gregory D. Durgin
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Professor Richard T. Causey
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Professor Paul A. Kohl
School of Chemical and Biomolecular
Engineering
Georgia Institute of Technology

Date Approved: 26 September 2018

*For my wife Brenna
and our children,
Tesla and Lenny.*

ACKNOWLEDGEMENTS

I would like to thank Silicon Creations LLC for their contributions to this work. They not only supported it financially by sponsoring my research at Georgia Tech, but also provided laboratory equipment, design and simulation software, and integrated circuit fabrication. Most importantly, I must thank the company founders, Jeff Galloway and Randy Caplan, as well as many other colleagues, with whom I held multiple technical discussions. These became the foundation for many original ideas contained in this writing.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
LIST OF TABLES	viii
LIST OF FIGURES	ix
SUMMARY	xii
I INTRODUCTION	1
1.1 Motivation	1
1.1.1 Timing Accuracy	1
1.1.2 The Analog CMOS IC Challenge	3
1.1.3 Timing Circuits	3
1.2 Thesis Overview	4
II OSCILLATORS AND PHASE NOISE	6
2.1 Introduction	6
2.2 Phase Noise in Signal Sources	6
2.3 Integrated Voltage-Controlled Oscillators	10
2.3.1 Ring VCOs	12
2.3.2 LC VCOs	13
2.4 Phase Noise Analysis Techniques for Oscillators	15
2.4.1 Linear Model	15
2.4.2 Hajimiri Model	16
2.5 Phase Noise in Ring Oscillators	17
2.5.1 Inverter Delay Uncertainty	17
2.5.2 Linear Model	18
2.5.3 Hajimiri Model	19
III FREQUENCY SYNTHESIS BY PHASE LOCK	20
3.1 Introduction	20
3.2 The Charge-Pump PLL	22
3.3 PLL Noise Dynamics	23

3.4	PLL Figure of Merit	27
3.5	Bang-Bang Phase-Locked Loops	27
3.6	Digital Phase-Locked Loops	29
3.6.1	Basic Concepts	29
3.6.2	Phase-Domain DPLL	30
3.7	Other PLL Topics	31
IV	TIME-TO-DIGITAL CONVERTERS	33
4.1	Introduction	33
4.2	Quantization Noise	34
4.3	Overview of TDC Techniques	36
4.3.1	Basic Concepts	36
4.3.2	Delay Line TDCs	38
4.3.3	Noise Shaping TDCs	40
4.3.4	Other TDC Techniques	41
V	DESIGN OF HYBRID DIGITAL-ANALOG PHASE-LOCKED LOOPS	43
5.1	Introduction	43
5.2	Proposed Architecture	43
5.3	Implementation	45
5.4	Applications and Results	47
5.4.1	Bang-Bang Behavior	47
5.4.2	Gapped Clock	47
5.4.3	Spread Spectrum Demodulation	49
5.4.4	Large Multiplication Ratios	49
5.5	Conclusions and Proposed Research	50
VI	DESIGN OF A RING VCO WITH IMPROVED FOM BY ADDING AN INDUCTIVE LOAD	54
6.1	Introduction	54
6.2	Proposed Architecture	54
6.3	Frequency of Oscillation	55
6.4	Frequency Tuning Range	56
6.5	Phase Noise and Jitter	57

6.6	Prototype and Measurement Results	58
6.7	Conclusions and Proposed Research	61
VII DESIGN OF A TIME-TO-DIGITAL CONVERTER WITH SAMPLE-AND-HOLD AND QUANTIZATION NOISE SCRAMBLING USING HARMONICS IN RING OSCILLATORS		63
7.1	Introduction	63
7.1.1	Bandwidth in Noise-Shaping TDCs	63
7.1.2	Harmonics in Ring Oscillators	63
7.2	Proposed Architecture	64
7.2.1	Sample-and-Hold in the Phase Domain	64
7.2.2	Number of Harmonics	65
7.2.3	Measurement Process	67
7.2.4	Noise Scrambling	68
7.3	Noise Analysis	68
7.3.1	Quantization Noise	69
7.3.2	Device and Circuit Noise	70
7.4	Implementation	76
7.5	Measurement Techniques and Results	78
7.5.1	Input Signals	78
7.5.2	Test Setup	79
7.5.3	Characterization	79
7.5.4	Proposed Figure of Merit	82
7.6	Conclusions and Proposed Research	83
VIII CONCLUSIONS, CONTRIBUTIONS AND FUTURE RESEARCH		85
8.1	Conclusions	85
8.2	Contributions	88
8.3	Future Research	88
REFERENCES		90

LIST OF TABLES

1	List of parameters at which LC and ring VCOs outperform each other. . . .	15
2	Performance summary and comparison between recent ring-oscillator-based noise-shaping TDCs and the proposed TDC.	83

LIST OF FIGURES

1	Examples of analog signals: (a) Continuous voltage amplitude and continuous time. (b) Continuous time amplitude and discrete time.	1
2	The PSD of phase noise and the PSD of the phase-noise-modulated carrier.	7
3	Phase and phase noise.	8
4	Jitter accumulation.	8
5	Jitter accumulation due to correlated and uncorrelated sources.	9
6	Noise spectrum: (a) Thermal and flicker noise, (b) Integration of noise in an oscillator, and (c) Phase noise.	10
7	Reduction of phase noise by increasing power.	11
8	Ring voltage-controlled oscillators: (a) Single-ended, supply-driven, (b) Differential with quadrature outputs.	13
9	Cross-coupled-pair LC voltage-controlled oscillator.	14
10	Linear model of an oscillator with a noise source.	15
11	Approximate waveform and ISF of a ring oscillator [22].	17
12	Model for current noise to delay noise conversion in a CMOS inverter [1].	18
13	Linear model of a three-stage ring oscillator with a noise current source [37].	18
14	Reference, feedback (or VCO) and error signals in a PLL.	21
15	Feedback control system.	21
16	Charge-pump PLL.	22
17	Linear model of a CPPLL.	24
18	Noise transfer functions in a PLL.	25
19	PLL phase noise contributors and output.	26
20	Jitter accumulation in a PLL.	26
21	Phase tracking of a bang-bang PLL in the time domain.	28
22	Frequency response to reference noise in a bang-bang PLL.	29
23	Implementation of a DPLL by using an ADC and a DAC around a digital loop filter.	30
24	Synchronous phase-domain all-digital PLL-based transmitter [47].	31
25	First-order $\Delta\Sigma$ modulator.	35
26	Noise transfer function of a first-order $\Delta\Sigma$ modulator.	36

27	Effective quantization noise in $\Delta\Sigma$ converters of different order (L) as a function of oversampling ratio (OSR).	37
28	TDC Input Signals: (a) Event signals and (b) Reference clock.	37
29	Delay-line-based TDCs. (a) Flash TDC and (b) Vernier TDC.	38
30	Ring-oscillator-based noise-shaping TDCs. (a) Gated ring oscillator (GRO) TDC, (b) Switched ring oscillator (SRO) TDC.	40
31	PLL noise transfer functions optimized for (a) VCO noise and (b) reference noise.	44
32	DPLL using a fractional-N CPPLL as a DCO.	44
33	Output phase noise contributors in the proposed hybrid PLL.	45
34	Test setup to determine the response of the hybrid PLL to input phase noise.	47
35	Effect of loop-filter gain and modulation amplitude on the frequency response of the DPLL.	48
36	FIFO buffer for two merging networks with “gapped” <i>write</i> clock and recovered <i>read</i> clock.	48
37	Measured output phase noise of the proposed DPLL with gapped-clock input.	49
38	Spread-spectrum demodulation with the proposed DPLL. Spectrum of the input and output signals.	50
39	Measured output phase noise for the proposed DPLL with large multiplication ratios.	51
40	Proposed hybrid DPLL with an alternative set of reference signals.	52
41	Output phase noise contributors in hybrid PLL with LC and crystal references.	52
42	Proposed ring VCO with an inductive load.	54
43	Linear model of a ring oscillator with an inductive load and noise current source.	55
44	Simplified schematics of the PLL for testing the proposed VCO prototype.	58
45	Photograph of the PLL for testing the proposed VCO prototype implemented on a PCB.	59
46	Frequency and current versus voltage for different inductance values.	60
47	PN and FOM comparison between the ring oscillator without inductor and with different inductor values when $V_{\text{ring}} = 4.5\text{V}$	60
48	Measurements and predictions summary for the proposed VCO when $V_{\text{ring}} = 4.5\text{V}$	61
49	Ring oscillator operating in second harmonic mode.	64
50	Differential sample-and-hold in the phase domain.	66

51	Conceptual representation of the HRO TDC during its three measurement states: (a) idle, (b) sample and (c) hold.	66
52	Measurement process and definitions for the HRO TDC.	67
53	Model of the HRO TDC for noise analysis.	69
54	Comparison of quantization noise versus the oversampling ratio of $\Sigma\Delta$ noise-shaping TDCs and the HRO TDC.	71
55	Impact of noise on the sample-and-hold process. (a) Fully correlated noise and (b) uncorrelated noise.	72
56	Simulation setup to compute the ISF at the supply voltage of a ring oscillator.	74
57	Simulation results from the setup in Figure 56: (a) phase error as a function of the angle at which the charge was injected and (b) voltage at each node in the ring.	75
58	Schematics of the prototype HRO TDC.	76
59	HRO TDC chip photograph (a), layout and dimensions (b).	78
60	HRO TDC test setup. Phase shifting signal B using a fractional-N $\Delta\Sigma$ feedback divider PLL.	79
61	Measured TDC response for an input of two 100 MHz signals swept in relative phase for 360 degrees in steps of 12.5 ps.	80
62	Measured integral nonlinearity of the HRO TDC.	81
63	Measured precision of the HRO TDC versus the number of samples per measurement for constant measurement time.	81
64	Measured precision of the HRO TDC versus the number of samples per measurement for constant sampling frequency.	82

SUMMARY

In this research, I present solutions for frequency synthesis and other time-domain signal processing problems. The current trends of CMOS technology motivate these solutions and their approach. In particular, I address the degradation in analog performance, resulting from a lower supply voltage, higher mismatch and higher leakage. I also embrace the opportunities arising from faster transistors and faster logic gates.

I begin by extracting from the available literature a collection of essential tools that describe circuit performance in the time domain, particularly about the effects of noise. This review is divided into three chapters. The first explores how phase noise and jitter are modeled and the mechanisms that cause phase noise and jitter in oscillators. The second examines how phase-locked loop architecture affects their ability to process phase noise. In particular, I explore the charge-pump PLL, the bang-bang PLL and digital PLLs. Finally, I delve into the theory and techniques behind time-to-digital converters. These circuits are a fundamental building block for measuring time intervals and for implementing phase detectors in digital PLLs.

Building upon this knowledge, I develop techniques, in particular, new circuit architectures, that improve the performance of frequency synthesizers.

The first proposed technique responds to the limitation of traditional PLLs, which filter reference noise and VCO noise complementarily. I present an architecture that combines digital and analog PLL loops which proves to have the ability to filter both sources of noise simultaneously. The second technique is a VCO architecture that eases tradeoffs between the convenience and flexibility of ring oscillators and the low phase noise of LC oscillators. It is based on a ring oscillator which has every node in the ring attached to an inductive load structure. It achieves a well-characterized improvement in phase noise and power consumption. The third and final contribution is a time-to-digital converter architecture which builds upon the most recent TDC developments and presents for the

first time a sample-and-hold mechanism. This technique allows for significant improvement in measurement bandwidth.

For each of these techniques, I have collected substantiating data through prototypes in either integrated circuit form, PCB form, or both.

CHAPTER I

INTRODUCTION

1.1 Motivation

1.1.1 Timing Accuracy

The signal quantities that may be found in electronic systems can be classified as time-domain quantities (such as period, frequency, phase, pulse width or duty cycle), or as non-time-domain quantities (such as voltage, current, charge or even a numerical digital value). When we generate, transmit, measure or process any signal we introduce errors. This is especially true when we convert signals between domains. The magnitude of the errors depends on the choice of quantities in each domain and the conversion mechanism between these domains.

As an example, consider the choice that a circuit designer must make in order to transmit a signal over a long distance across an electronic circuit that has a noisy power supply. Two possible methods of encoding the information in the signal are shown in Figure 1. Signal (a) encodes it as a voltage, while signal (b) encodes it in the width of the pulses. How would these two signals respond to noise in the circuit's power supply? Long distances can result in large electromagnetic coupling between the signal and the supply rails. In such a case, the noise can easily distort the amplitude of signal (a), while signal (b) would remain




Figure 1: Examples of analog signals: (a) Continuous voltage amplitude and continuous time. (b) Continuous time amplitude and discrete time.

less sensitive to it. However, long distances may also require introducing buffers along the line, and the variations in the supply can modulate the delay across them. In this case, the noise translates directly into the time-domain making signal (b) highly vulnerable.

Furthermore, we must evaluate the distortion introduced during the generation, measurement and processing of these signals. These three operations are often used cyclically, where a signal is generated and then measured against a reference to compute the error. This error is then processed in order to apply a correction to the generation mechanism, yielding an adjusted signal which can be processed again. In such a sequence, the accuracy of the reference signal is critical. Any signal that is generated with or measured against a reference can only be as accurate as the reference.

In this work, I shall provide particular attention to time-domain reference signals and the circuits that generate them. This focus materializes specifically in the subjects of phase noise, oscillators and phase-locked loops (PLLs).

Timing inaccuracies, whether systematic or random, limit the performance of multiple kinds of electronic systems. For example, in both wireless and wire-line communications, timing accuracy determines how much data per unit time can be transmitted through a given channel, how efficiently channels can be shared, and the integrity of the data in those channels. Furthermore, it determines the complexity and cost of transmitters and receivers.

Another example is data converters, which include analog-to-digital (ADC) and digital-to-analog (DAC) converters. A noisy sampling clock (i.e., a clock that exhibits variations in its period) reduces the signal-to-noise ratio (SNR) of the converted signal.

One other example is synchronous digital logic, where uncertainty in the clock period limits the maximum clock frequency. The period of the clock must never be shorter than the time it takes for the logic to settle. Therefore, the nominal clock period must be set to a larger value to ensure that a period that is shorter than the settling time never occurs.

The dual of this problem is when there is uncertainty about the propagation delay through the logic. This uncertainty enforces setting the clock to have a period larger than the propagation delay that may ever be encountered. Propagation delay uncertainty is not just a concern in digital design but an overarching problem in timing accuracy.

1.1.2 The Analog CMOS IC Challenge

The electronic circuit platform of choice today is the silicon complementary metal-oxide-semiconductor (CMOS) integrated circuit (IC). This technology has seen extraordinary progress as predicted by Moore in 1965, resulting in smaller and faster transistors, smaller and faster logic gates, and lower power consumption.

However, smaller transistors are typically associated with inferior analog performance. They exhibit higher leakage, which means that the drain-source resistance in their off state is lower. They also display worse matching, which is the result of larger manufacturing imperfections relative to their size. Lastly, they operate at a lower power supply voltage, which lowers signal amplitude, leading to a lower SNR.

The circuit designer faces the problem of having to achieve continually more demanding system requirements while designing for a platform with an analog performance that has and may continue to degrade.

1.1.3 Timing Circuits

The family of timing circuits is composed of all those circuits that generate, measure or process signals carrying information in the time domain. The PLL is arguably the most ubiquitous timing circuit. Since it entails sub-circuits that carry out all three of the previously mentioned operations, it is perhaps the best representation of the entire timing circuit family.

PLLs can accurately establish the frequency and phase of a tunable oscillator. They achieve this by “locking” the phase of the oscillator to the phase of a reference oscillator. This method has proven to be very effective and reliable, and it is well understood. These characteristics have ultimately made the PLL the preferred frequency synthesis technique. Moreover, it has been so for many decades, throughout which it has been continuously refined. It has matured into what we call today the charge-pump fractional-N PLL.

The trends in CMOS, however, have been driving the interest to implement many sub-circuits of the PLL, a predominantly analog technique, in digital form. Converting as many

PLL subcircuits to digital form may seem an obvious path to follow, especially when considering the benefits of the loop filter. In digital form, it requires a much lower area, it is insensitive to process and environmental variations, and it is much more flexible and reconfigurable than its analog counterpart. PLLs that incorporate a digital loop-filter are referred to as digital PLLs (DPLLs).

Despite the name, PLLs cannot be entirely digital. The phase error detector and the tunable oscillator remain unavoidably mixed-signal circuits. Mixed-signal circuits are analog circuits that incorporate digital logic (or vice-versa). The phase detector of a DPLL involves some form of (analog-) time-to-digital converter (TDC), and the oscillator involves some form of digital-to-(analog-)frequency conversion. Such an oscillator is commonly referred to as a digitally-controlled oscillator (DCO). TDCs and DCOs belong to a category of mixed-signal circuits called data converters. This category groups together all ADCs and DACs.

Since TDCs and DCOs are inevitably analog circuits and are therefore affected by the degrading analog performance in CMOS, designers and researchers need new approaches to continue to improve their performance.

The challenges and opportunities described above have guided the proposed research, which narrows down to how oscillators, PLLs and TDCs affect timing accuracy and how we may improve them.

1.2 Thesis Overview

I explore three essential topics in electronic frequency synthesis and time-domain signal processing.

The first topic examines oscillators and phase noise. In Chapter 2, I provide an introduction to phase noise and jitter. These two terms are terms used when referring to timing uncertainty in the signal generated by an oscillator. Noise in oscillators is the primary source of timing uncertainty in a frequency synthesizer. In this same chapter, I provide an introduction to oscillators for integrated circuits. In Chapter 6, I propose a new architecture of voltage-controlled oscillator (VCO) based on a ring oscillator and inductive loading

which provides better phase noise than traditional ring oscillators without consuming more power.

The second topic is frequency synthesis using PLLs. PLLs are feedback circuits that synchronize a VCO with a reference signal to reduce the VCO's phase noise and establish an arbitrary output frequency. In Chapter 3, I provide an introduction to PLLs, their architecture and how they shape the noise spectrum of the VCO. I also provide a brief introduction to digital PLLs (DPLLs), which are a recent trend in response to the evolution of CMOS technology. In Chapter 5, I propose an architectural technique that combines analog and digital PLLs to overcome a fundamental limitation of traditional PLLs: the complementary filtering of VCO and reference noise.

The third and final topic investigates TDCs. These data converters are an active research topic responding to the increasing interest in time-domain signal processing. They are also a fundamental building-block in DPLLs. In Chapter 7, I propose a ring-oscillator-based TDC that provides the first reported differential sample-and-hold mechanism in the phase-domain, allowing for much higher bandwidth than recent TDCs.

I conclude in Chapter 8 with overarching observations about my contributions to the field of frequency synthesis and time-domain signal processing, and highlight some opportunities that may arise from further research.

CHAPTER II

OSCILLATORS AND PHASE NOISE

2.1 Introduction

Phase noise and jitter are frequent terms used to refer to the fluctuations, or the uncertainty, in a signal's time-domain properties, typically the phase and frequency of a periodic signal. These variations introduce uncertainty that limits the performance of communication systems, digital systems, and data converters, among others. Understanding the nature and behavior of phase noise is fundamental to designing circuits that minimize it or that are highly insensitive to it.

2.2 Phase Noise in Signal Sources

Periodic signals can be represented mathematically by

$$x(t) = A(t) \cdot g(\omega_0 t + \phi_n(t)) \quad (1)$$

where g is a function that is periodic in 2π . Its argument, $\omega_0 t + \phi_n(t) = \phi(t)$, is the phase and ω_0 is the frequency. $A(t)$ is a stationary random process, with an average A_0 , which represents perturbations in the amplitude. The zero-mean random process $\phi_n(t)$ represents the deviations of $\phi(t)$ from the ideal $\phi_i(t) = \omega_0 t$ and is known as phase noise.

The phase (and therefore the phase noise) of g is not directly observable from the values of $x(t)$, in part, because it is not possible to separate the contributions of $A(t)$ and $\phi_n(t)$ to the deviations of $x(t)$ from $A_0 \cdot g(\phi_i(t))$. However, if we consider the case where $g(\phi) = \cos(\phi)$, and assume that $A(t) = A_0$ and $\phi_n(t) \ll 1$ radian,

$$x(t) = A_0 \cos(\omega_0 t + \phi_n(t)) \quad (2)$$

$$\approx A_0 \cos(\omega_0 t) - A_0 \phi_n(t) \sin(\omega_0 t), \quad (3)$$




Figure 2: The PSD of phase noise and the PSD of the phase-noise-modulated carrier.

This expression indicates that the power spectral density (PSD) of $x(t)$, S_x , is the sum of the PSD of the noiseless $x(t)$ and S_ϕ^1 , the PSD of $\phi_n(t)$, translated by ω_0 and scaled by A_0 (Figure 2). Consequently, we can measure S_ϕ by directly measuring S_x .

Another practical approach to estimating the phase noise is by sampling the times when $x(t)$ crosses a specified threshold. One should choose the threshold such that the sensitivity of the measurement to variations in $A(t)$ is minimal, for example, when $dg/d\phi$ is maximum. The deviations from the times $x(t)$ is expected to cross the threshold is a discrete set of values known as jitter. Figure 3 illustrates how phase, phase noise, and jitter are related.

Jitter is most often specified by its standard deviation, $\sigma_{\Delta T}$, where ΔT is the time of observation. Since jitter is a random-walk process, $\sigma_{\Delta T}$ increases with ΔT as illustrated in Figure 4. We can calculate $\sigma_{\Delta T}$ from S_ϕ (and vice versa) from the relationship [1]

$$\sigma_{\Delta T}^2 = \int_0^\infty S_{\Delta T}(f) df = \int_0^\infty S_\phi(f) \frac{\sin^2(\pi f \Delta T)}{(\pi f_0)^2} df. \quad (4)$$

The specific rate at which jitter accumulates depends on the correlation of the contributing noise sources (Figure 5). Uncorrelated sources dominate from small values of ΔT up to some time t_c . In this region $\sigma_{\Delta T} = \kappa\sqrt{\Delta T}$ [34]. For larger intervals, correlated sources

¹Since the PSD of ϕ and ϕ_n are identical, most authors simply use S_ϕ to refer to the PSD of phase noise. Some authors also refer to phase noise by its single side-band PSD, $\mathcal{L}(f) = S_\phi(f)/2$. The notation $S_\phi(\Delta\omega)$ is also often found in the literature where $\Delta\omega$ is the “offset from the carrier”, since it is the equivalent to observing $S_x(\omega_0 + \Delta\omega)$.




Figure 3: Phase and phase noise.




Figure 4: Jitter accumulation.

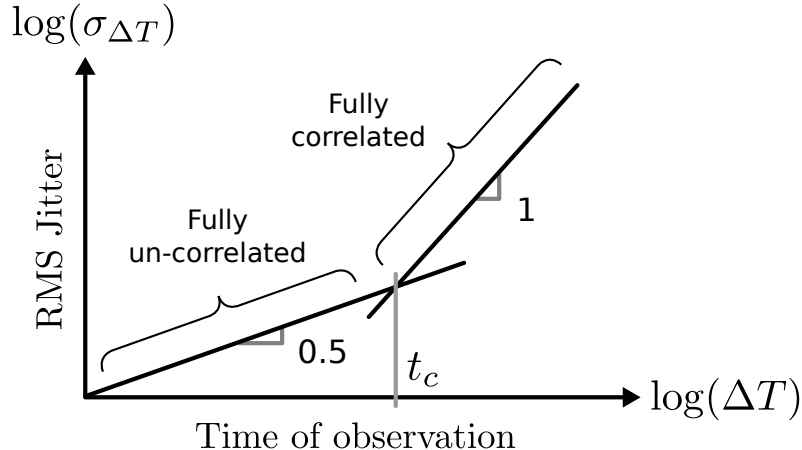


Figure 5: Jitter accumulation due to correlated and uncorrelated sources.

dominate, and $\sigma_{\Delta T} = \zeta \Delta T$ [22]. The constants κ and ζ are noise figures-of-merit specific to a given signal source.

Noise from two different devices, from one device and the power supply, or from one device and another circuit, are all examples of pairs of uncorrelated noise sources. Such combinations of noise lead to jitter accumulation in proportion to $\sqrt{\Delta T}$. However, it is possible to have a single noise source affecting the phase of an oscillator in a correlated manner. This situation occurs when the effect of noise across consecutive cycles of the signal is correlated. For this to happen, the noise source must have a non-zero autocorrelation at time offsets greater than zero [13]. This condition also manifests as an uneven or “colored” PSD. Examples of such noise sources are low-frequency power-supply noise and flicker noise. Conversely, a signal with a flat PSD, like thermal noise, for example, is referred to as being “white” and has zero autocorrelation everywhere except at zero time offset.

The voltage (or current) spectrum S_{V_n} of combined thermal and flicker noise sources is shown in Figure 6(a). Flicker noise exhibits a $1/f$ profile and dominates below some frequency $f_{1/f}$. This “corner” frequency depends on the relative power of thermal and flicker noise. The process by which these forms of noise translate to phase noise depends on many factors, such as where, within the oscillator circuit, the noise originates, the specific oscillator topology and its specific implementation.

However, free-running oscillators do not have a mechanism that will restore the phase




Figure 6: Noise spectrum: (a) Thermal and flicker noise, (b) Integration of noise in an oscillator, and (c) Phase noise.

to its ideal value after it has been disturbed by noise. All perturbations are accumulated indefinitely. In other words, noise is integrated into phase noise. Therefore, the conversion from voltage/current noise to phase noise must exhibit a $1/s$ term (Figure 6(b)). Consequently, the $1/f$ and the flat profiles of flicker and thermal noise, appear with $1/f^3$ and $1/f^2$ profiles respectively in the phase-noise PSD (Figure 6(c)).

The f_{1/f^3} corner frequency is lower than $f_{1/f}$ as I show later with Eq. 14, and can be minimized by ensuring waveform symmetry in the oscillator [22]. It can also be related to t_c (Figure 5) by the approximation [31]

$$t_c \approx \frac{1}{25f_{1/f^3}}. \quad (5)$$

2.3 Integrated Voltage-Controlled Oscillators

Voltage-controlled oscillators (VCOs) are commonly compared using the figure-of-merit (FOM)

$$\text{FOM}_1 = \frac{f_0^2}{P \times \text{PN} \times \Delta f^2} \quad (6)$$

where f_0 is the frequency of oscillation, P is the power consumption and PN is the phase noise measured at $f = f_0 + \Delta f$ [38].

FOMs, in general, indicate direct design tradeoffs. First, a given value of FOM_1 indicates

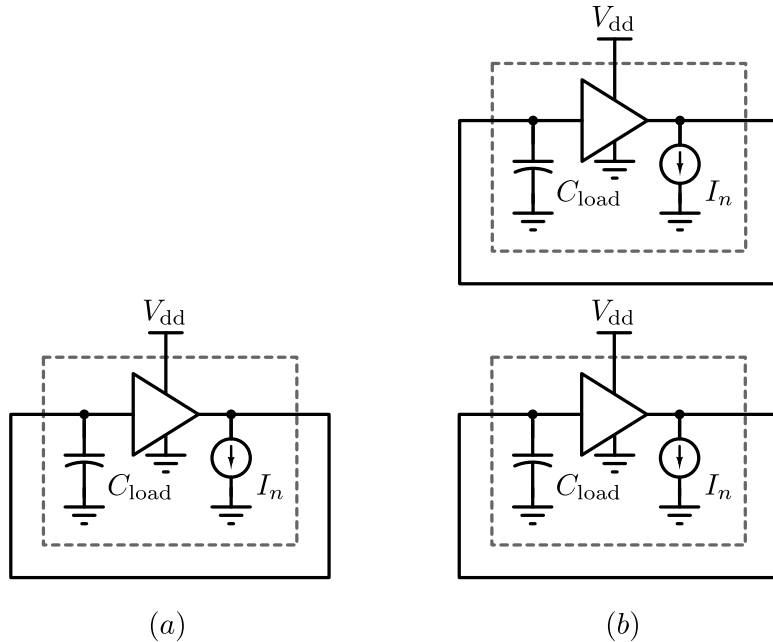


Figure 7: Reduction of phase noise by increasing power.

that PN is inversely proportional to Δf^2 . This, as shown in Section 2.2, assumes that PN is being measured at $\Delta f > f_1/f^3$. Second, that PN is proportional to f^2 . This can also be concluded from Section 2.2. It is reasonable to assume that jitter is independent of frequency, therefore, as a fraction of the oscillation period, it is proportionally larger as the frequency increases. Lastly, that PN can be reduced proportionally by using more power.

We can explain the PN-versus- P tradeoff as follows. Consider the hypothetical oscillator circuit in Figure 7(a) which consumes a power P . It consists of an amplifier connected in feedback with an input capacitance C_{load} and output-referred current noise of magnitude I_n . I_n is converted to a voltage V_n as it charges C_{load} . Now, consider two such oscillators connected in parallel as shown in Figure 7(b). Since they are identical, no current flows through the connection. Therefore, no change in frequency or power consumption in each oscillator would occur. However, the total power is now $2P$. In the presence of the two uncorrelated (since they come from different devices) noise sources, their powers add, i.e., the total noise current amplitude is $\sqrt{2}I_n$. Since there is now twice as much load capacitance, the total noise voltage becomes $\sqrt{2}V_n/2 = V_n/\sqrt{2}$. Therefore, twice as much power results

in $1/\sqrt{2}$ as much noise amplitude and half as much noise power. Of course, we can extend this to any number of oscillators in parallel.

FOM_1 is only a first-order quantification of the “value” (merit) of a VCO. Further characteristics, such as tuning range, silicon area, the linearity of the frequency as a function of the control voltage/current, power supply sensitivity, portability across technology nodes, and availability of differential, quadrature, or multi-phase outputs, are essential. These, as well as FOM_1 , determine if the requirements of a system can be met with a given choice of VCO architecture.

Two architectures are common in modern CMOS ICs: ring oscillators and cross-coupled-pair LC tank oscillators. Basic schematics for both approaches are shown in Figures 8 and 9 respectively.

2.3.1 Ring VCOs

Ring VCOs can be implemented by cascading inverting logic gates as shown in Figure 8(a). With an odd number of single-ended inverters, the loop is always unstable and therefore oscillates. The frequency is tuned by the supply voltage, although it is also possible to control it with a similar performance by using a current source. Other methods of tuning the frequency are possible but not as common.

The frequency of a ring oscillator is given by

$$f_{\text{ring}} = \frac{1}{2N\tau} \quad (7)$$

where N is the number of inverting stages, and τ is the propagation delay of each stage.

The propagation delay of a CMOS inverter is not easily captured in a closed-form equation since it depends on the non-linear dynamics of the transistors in the gate, the input waveform and possibly a nonlinear load (such as another CMOS gate). However, by assuming an instantaneous logic transition at the input and a constant current at the output, τ can be approximated by

$$\tau = \frac{V_{\text{ring}}}{2} \times \frac{C_{\text{gate}}}{I_{\text{sat}}} \quad (8)$$




Figure 8: Ring voltage-controlled oscillators: (a) Single-ended, supply-driven, (b) Differential with quadrature outputs.

where V_{ring} is the supply voltage, I_{sat} is the saturation current of either the PMOS or NMOS transistor in the inverter, and C_{gate} is the input capacitance of the next inverter in the ring, i.e., the load [1].

Ring VCOs can also be designed using differential or pseudo-differential gates. It is possible to use an even number of such gates since an odd number of inversions can be achieved by inverting the polarity at one of the stages. Furthermore, this allows for an even number of equally spaced phases to be generated from the VCO as shown in Figure 8(b).

In general, except for the metrics represented by FOM_1 , ring VCOs are more convenient, versatile and superior to LC VCOs. Table 1 summarizes the benefits of ring VCOs when compared to LC VCOs.

2.3.2 LC VCOs

LC VCOs are based on an LC resonator (or tank). The frequency is that at which the inductor and capacitor resonate,




Figure 9: Cross-coupled-pair LC voltage-controlled oscillator.

$$f_{LC} = \frac{1}{2\pi\sqrt{LC}}, \quad (9)$$

and is commonly tuned by varying the tank's capacitance through the use of varactors.

The predominant architecture for integrated LC VCOs, known as the cross-coupled-pair LC VCO, is shown in Figure 9. The cross-coupled pair of MOSFETs, M1 and M2, act as a negative resistance amplifier that compensates for the losses in the tank, maintaining the oscillation amplitude.

LC VCOs exhibit significantly better phase noise for a given power consumption (FOM_1) than ring VCOs. This is why they are typically chosen for the most demanding low phase-noise applications such as local oscillators for wireless radios. It is uncommon, however, to use an LC VCO when the FOM_1 requirements of a specific application allow for a ring VCO. LC VCOs are costly and inconvenient when compared to ring VCOs as summarized in Table 1.

Table 1: List of parameters at which LC and ring VCOs outperform each other.

LC VCOs	Ring VCOs
FOM ₁	Tuning Range
	Area
	Linearity
	Portability
	Startup Gain
	Quadrature/multiphase




Figure 10: Linear model of an oscillator with a noise source.

2.4 Phase Noise Analysis Techniques for Oscillators

2.4.1 Linear Model

A theoretical oscillator can be constructed by cascading and feeding back a series of linear amplifiers with an overall response $H(j\omega)$. The conditions for such oscillator to oscillate at a specific frequency ω_0 are that the open-loop gain must be unity and that the phase shift must be a multiple of 2π . These conditions are known as the Barkhausen criterion [38].

A methodology for analyzing the noise of oscillators by modeling them as linear circuits was developed in [37]. It considers a feedback oscillator with an open-loop transfer function $H(j\omega)$, an output voltage V and a noise source V_n , as shown in Figure 10.

The noise-to-output transfer function is then

$$\frac{V}{V_n}(j\omega) = \frac{H(j\omega)}{1 + H(j\omega)}. \quad (10)$$

At a small $\Delta\omega$ offset away from the oscillation frequency ω_0 , i.e., $\omega = \omega_0 + \Delta\omega$,

$$H(j\omega) \approx H(j\omega_0) + \Delta\omega \frac{dH}{d\omega}. \quad (11)$$

Since the Barkhausen criterion requires that $H(j\omega_0) = -1$, and in general, $|\Delta\omega dH/d\omega| \ll$

1,

$$\frac{V}{V_n} [j(\omega_0 + \Delta\omega)] \approx \frac{-1}{\Delta\omega \frac{dH}{d\omega}} \quad (12)$$

and

$$\left| \frac{V}{V_n} [j(\omega_0 + \Delta\omega)] \right|^2 = \frac{1}{\Delta\omega^2 \left| \frac{dH}{d\omega} \right|^2}. \quad (13)$$

This expression confirms the $1/s^2$ shaping of the baseband noise PSD as it upconverts around the carrier frequency as I informally proposed in Section 2.2 and Figure 6, and establishes a gain factor of $|dH/d\omega|^{-2}$.

2.4.2 Hajimiri Model

Named informally after the author, the theory in [21] builds upon two fundamental observations. It first recognizes that the response of a VCO to noise is time-variant. I.e., the phase disturbance in an oscillator depends on the time at which certain noise transient occurs. Moreover, it points out that at any given time (with respect to the period of oscillation), if a fixed amount of charge is instantaneously injected into the oscillator circuit, the resulting phase disturbance is permanent and its magnitude is in linear proportion to the amount of injected charge.

The proportionality between the amount of injected charge and the amount of resulting phase shift is a time-varying proportionality constant, periodic in 2π , known as *impulse sensitivity function* (ISF), and is denoted by $\Gamma(x)$. Figure 11 shows $\Gamma(x)$ along the waveform $g(x)$ for one node of a ring oscillator [22].

By calculating $\Gamma(x)$ for a specific oscillator, this methodology allows for accurate computation of the upconversion of baseband current or voltage noise occurring at any node of the oscillator circuit.

One remarkable and applicable finding in [21] is that

$$f_{1/f^3} = f_{1/f} \frac{\Gamma_{dc}^2}{\Gamma_{rms}^2}, \quad (14)$$




Figure 11: Approximate waveform and ISF of a ring oscillator [22].

where Γ_{dc} and Γ_{rms} are the DC and RMS values of $\Gamma(x)$ respectively. This expression quantifies the upconversion of flicker noise and other low-frequency noise.

2.5 Phase Noise in Ring Oscillators

Multiple phase noise analysis techniques such as the general ones presented in Section 2.4 are readily applicable to ring oscillators. We can achieve further insight by studying specific ring oscillator subcircuits.

Since a ring oscillator is often built with a series of inverter logic gates connected in feedback, it is instrumental to understand how each gate contributes to phase noise and jitter.

2.5.1 Inverter Delay Uncertainty

The analysis presented in [1] derives the following expression for the uncertainty, i.e., the variance, in propagation delay τ of a single-ended CMOS inverter due to thermal noise (Figure 12):

$$\sigma_{\tau}^2 = \frac{4kT\gamma\tau}{I_{\text{sat}}(V_{\text{dd}} - V_t)} + \frac{kTC_{\text{gate}}}{I_{\text{sat}}^2}. \quad (15)$$

In Eq. 15, γ and V_t are transistor parameters. I_{sat} is the saturation current on either




Figure 12: Model for current noise to delay noise conversion in a CMOS inverter [1].




Figure 13: Linear model of a three-stage ring oscillator with a noise current source [37].

the NMOS or PMOS transistor (depending if the input to the inverter is *high* or *low*), V_{dd} is the supply voltage, C_{gate} is the load capacitance (the input capacitance of the next inverter) and k is the Boltzmann constant. This expression confirms that higher power consumption, this is, higher I_{sat} and higher V_{dd} , results in lower jitter.

2.5.2 Linear Model

The linear model for VCO noise of Section 2.4.1 applied to a three-stage ring oscillator [37] is shown in Figure 13 and yields

$$\left| \frac{V}{I_n} [j(\omega_0 + \Delta\omega)] \right|^2 = \frac{R^2}{27} \left(\frac{\omega_0}{\Delta\omega} \right)^2, \quad (16)$$

which follows after applying the Barkhausen criterion, dictating that $\omega_0 = \sqrt{3}/(RC)$ and $G_m R = 2$.

These expressions may be significantly inaccurate for computing absolute noise quantities due to the linear approximation. They can also be hard to apply since a real inverter does not have a constant output impedance R or constant transconductance G_m . However, this expression is essential when comparing the phase noise between ring oscillator topologies using the same inverter gates. It was instrumental in the development of the VCO presented in Chapter 6.

2.5.3 Hajimiri Model

The work in [21] was applied to ring oscillators by the same authors of [22]. They were able to reasonably approximate $\Gamma(x)$ for ring oscillators (Figure 11) as a function of the number of inverters N . Furthermore, they found that

$$\Gamma_{\text{rms}} = \sqrt{\frac{2\pi^2}{3\eta^3} \frac{1}{N^{1.5}}}, \quad (17)$$

where η is a proportionality constant, close to 1, between the gate delay and the inverse of the maximum slope g'_{max} of the waveform (Figure 11).

This result in combination with Eq. 14 and Figure 11, allow for two important conclusions. First, the symmetry of the oscillator's waveform determines the symmetry of $\Gamma(x)$, which consequently determines Γ_{dc} . Therefore, the designer can optimize for flicker noise upconversion by adequately sizing the PMOS and NMOS transistors in the inverters. Second, the non-linear dependence of Γ_{rms} on N leads to the finding that we can reduce f_1/f^3 by increasing the number of stages in the ring.

CHAPTER III

FREQUENCY SYNTHESIS BY PHASE LOCK

3.1 Introduction

Practical electronic circuits require some form of a timing signal. Most will need more than one, perhaps at different frequencies, with some specific phase relationship, and with the least possible phase noise. In many cases, the frequency and phase of those signals may need to be accurately adjusted and even rapidly modulated. The generation of these signals, sometimes from a common fixed-frequency reference, is often referred to as frequency synthesis.

Phase-locked loops (PLLs) are circuits that can generate signals of arbitrary frequency and phase offset from a single fixed-frequency reference. PLLs are also able to recover a “pure” tone from a modulated clock or recover a clock from a serial data stream [40].

The circuit techniques in PLLs are many and diverse. Some of these involve the processing of continuous as well as sampled signals, in both analog and digital domains, with linear and non-linear behavior, including feedback and control systems. However, these are all technology-independent techniques, which have allowed PLLs to become ubiquitous since its invention in the 1930s, pre-dating semiconductor electronics, to today’s nanometer CMOS.

PLLs must achieve, at a minimum, three major functions: phase locking, frequency multiplication, and noise filtering.

Phase locking is the action of controlling the phase of a VCO such that it “follows” the phase of a reference signal, maintaining, on average, a constant difference (or error) with the reference (Figure 14). In some cases, this difference is desired to be as close to zero as possible. A PLL is said to be “in lock” when it achieves, or it is very close to, this steady-state condition.

To achieve phase lock, the PLL operates as a feedback control system (Figure 15). In a




Figure 14: Reference, feedback (or VCO) and error signals in a PLL.

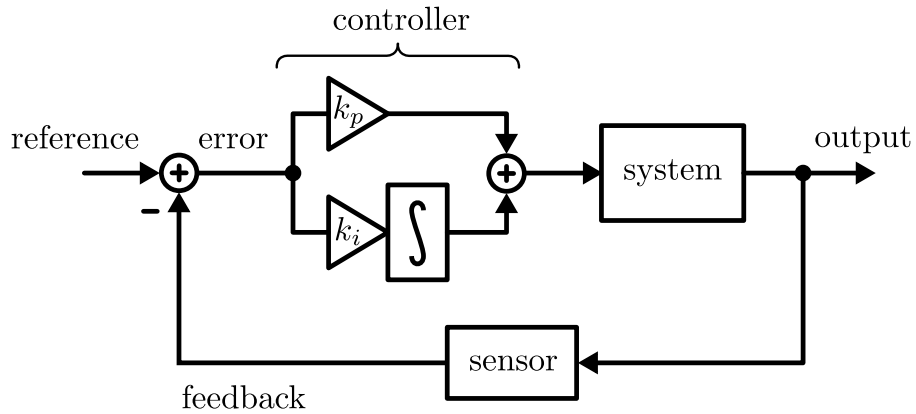


Figure 15: Feedback control system.

general control loop, the output value of a system is compared to a reference value. In some cases, the output of the system cannot be directly compared to the reference, so it is first converted to reference units by a sensor. The controller evaluates the difference (or error) in order to apply a correction to the system which shall reduce the error. One common form of controller is the proportional-integral (PI) controller. As its name implies, it applies a correction to the system that is proportional to the error as well as proportional to the integral of the error. The integral response is not strictly necessary, but it allows, in the absence of noise, to achieve zero error in steady state.

In the case of PLLs, the reference, output, and feedback are in the form of phase and the system is a voltage-controlled oscillator (VCO).

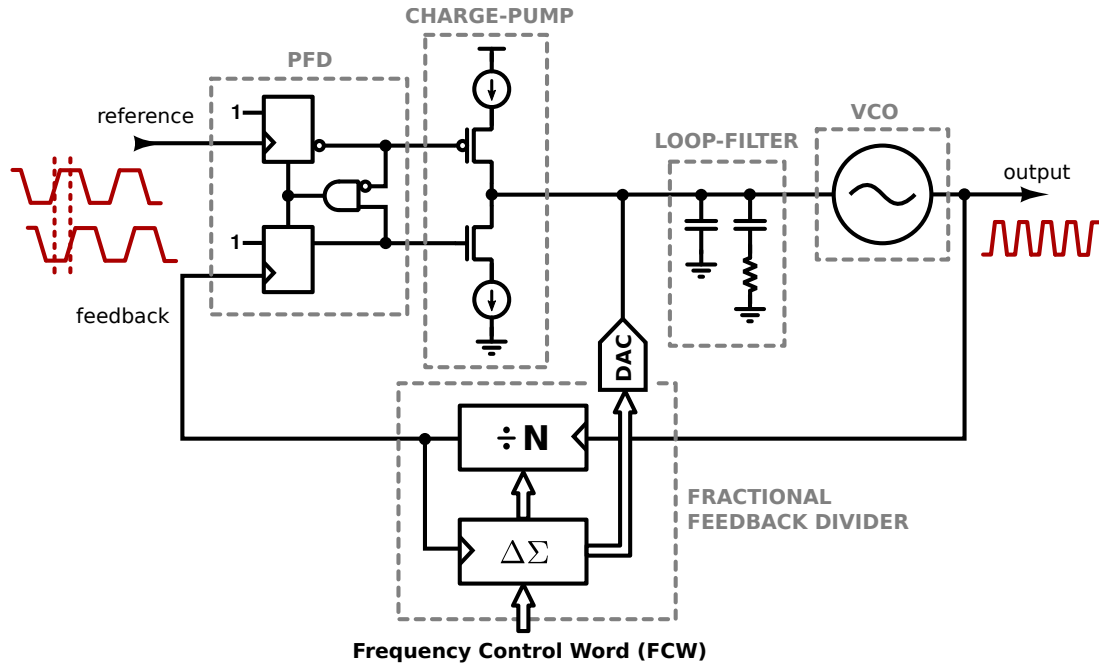


Figure 16: Charge-pump PLL.

3.2 The Charge-Pump PLL

Even though there are multiple ways to implement the control loop in a PLL, the charge-pump PLL (CPPLL) has been, and is today still, the best performing general purpose PLL architecture. Figure 16 shows the block diagram of a CPPLL.

The phase-frequency detector (PFD) and the charge pump (CP) operate together to deliver an amount of charge that is proportional to the phase error. This operation is achieved with excellent linearity, from an infinitesimally small phase error, up to $\pm\pi$. Furthermore, beyond the linear $\pm\pi$ range, the output is also proportional to the frequency difference. This behavior is essential for bringing the feedback frequency close enough to the reference frequency before phase locking can be achieved.

The loop filter is the controller in the loop. The circuit of a loop filter, shown in Figure 16, is a simple implementation that achieves a proportional as well as an integral voltage response to the charge delivered by the charge pump. Integration from current to voltage is performed by the capacitors, while the proportional response is due to the transient caused by the series RC branch.

The feedback divider is the sensor in the control loop which divides down the frequency of the VCO. This allows for the feedback signal to have the same frequency as the reference such that their phases can be compared. This also allows the VCO frequency to be a specified multiple of the reference frequency. The divider is most often implemented as a digital programmable counter, generating a rising edge at its output every N rising edges at its input. High VCO frequencies, however, may require the use of different techniques when implementing the divider, such as preceding the programmable counter by a prescaler, the use of dynamic-logic flip-flops, common-mode logic (CML), or down-conversion by mixing.

The use of a counter for frequency division is limited to integer values of N . However, a fractional N can be achieved by dithering between integer values. For example, if a division value of 9.9 is desired, N can be set to 10 for nine cycles and to 9 for one cycle. Of course, there is a quantization error, or noise, of 0.1 cycles during the first nine cycles, and a quantization error of 0.9 cycles during the last cycle. However, with proper design of the PLL's loop dynamics (Section 3.3), it is possible to significantly filter out the instantaneous error from the desired average at the output of the PLL.

Delta-Sigma ($\Delta\Sigma$) modulation is a technique commonly used to dither the divider values in PLLs [35]. I provide a brief introduction to $\Delta\Sigma$ modulation in Section 4.2. This method shifts the quantization noise power to higher frequencies [43] where it can be highly attenuated by the PLL dynamics (Section 3.3).

Further reduction of the “noise” injected by the $\Sigma\Delta$ modulator is possible by directly subtracting the error in a feed-forward manner after the phase detection. The magnitude of the error in each detection cycle is known and is available in digital form in the modulator. A digital-to-analog converter can then be used to subtract the analog equivalent from the loop filter.

3.3 PLL Noise Dynamics

Phase locking and frequency multiplication are essential requirements for a PLL to function. Its performance, however, is determined by how it processes (filters) different sources of noise.

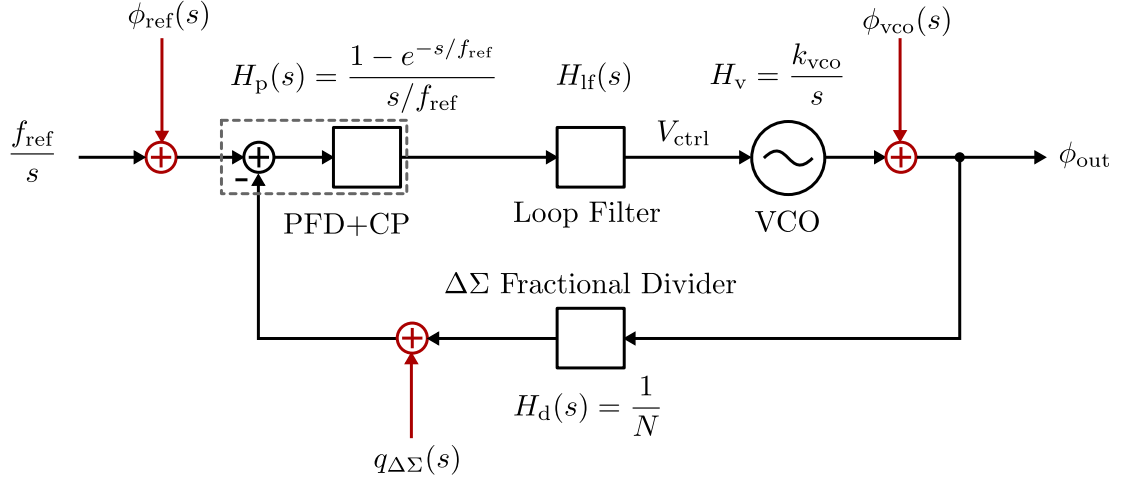


Figure 17: Linear model of a CPPLL.

Even though PLLs are highly non-linear circuits, when “in lock”, only small deviations from an average operating point occur. Small-signal approximation allows then for an accurate linear representation and modeling in the Laplace domain [35]. Figure 17 shows a linear model of a CPPLL including its primary sources of noise.

The loop filter is typically a linear time-invariant (LTI) circuit. Therefore it does not require a small signal approximation. From a control-systems perspective, a proportional-integral response of the form $H_{lf}(s) = k_p + k_i/s$ is desired. The frequency of the VCO (f_0) responds to the control voltage (V_{ctrl}) in proportion to the constant k_{vco} . Therefore, since phase is the integral of frequency, $H_v(s) = k_{vco}/s$. The feedback divider divides the frequency by N . Therefore it divides the phase by N as well. Finally, the PFD and CP are modeled as a zero-order-hold since the phase error between the reference and the feedback is periodically sampled.

If we define the forward path transfer function as $H_f(s) = H_{pfd}(s) + H_{lf}(s) + H_v(s)$, the transfer function from reference noise to PLL output is given by

$$H_{ref}(s) = \frac{\phi_{out}}{\phi_{ref}} = \frac{H_f(s)}{1 + H_f(s)/N} \quad (18)$$

and from VCO noise to PLL output by

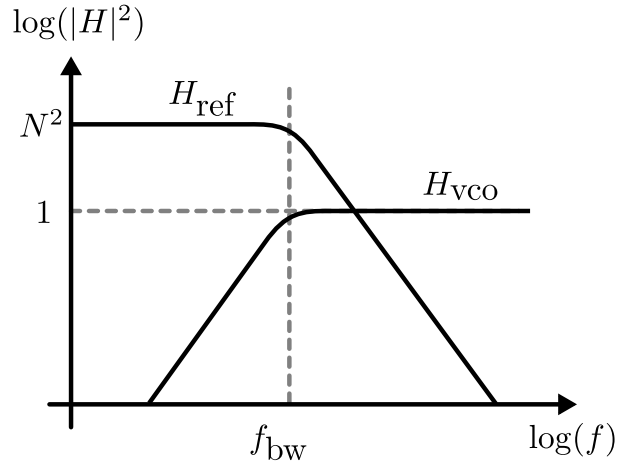


Figure 18: Noise transfer functions in a PLL.

$$H_{\text{vco}}(s) = \frac{\phi_{\text{out}}}{\phi_{\text{vco}}} = \frac{1}{1 + H_f(s)/N}. \quad (19)$$

These two transfer functions are plotted in Figure 18. H_{ref} and H_{vco} are low-pass and high-pass filters, respectively, with a common 3dB bandwidth frequency f_{bw} . Having a common f_{bw} means that at any given frequency, we can only design to suppress either the VCO's noise or the reference noise, but not both. In addition, reference noise is subject to amplification of N^2 .

Figure 19 illustrates how VCO noise and reference noise contribute towards the total output phase noise when subject to a given loop bandwidth and frequency multiplication ratio. At high frequencies, the VCO's noise dominates, reaching the output unaltered. Below f_{bw} it flattens since it is attenuated approximately as fast as the noise increases with decreasing frequency. At further lower frequency, the amplified reference noise surpasses the VCO noise and appears directly at the output.

A time-domain (i.e., jitter) equivalent of Figure 19 is shown in Figure 20. For a small observation time, below $t_{\text{bw}} \approx 1/f_{\text{bw}}$, jitter increases linearly, like it would for a free-running oscillator (Figure 5). Above t_{bw} , the PLL suppresses the VCO's jitter, i.e., no additional jitter continues to accumulate. Therefore, the curve flattens. Past some value of ΔT , accumulation of jitter from the reference appears at the output and continues to increase.

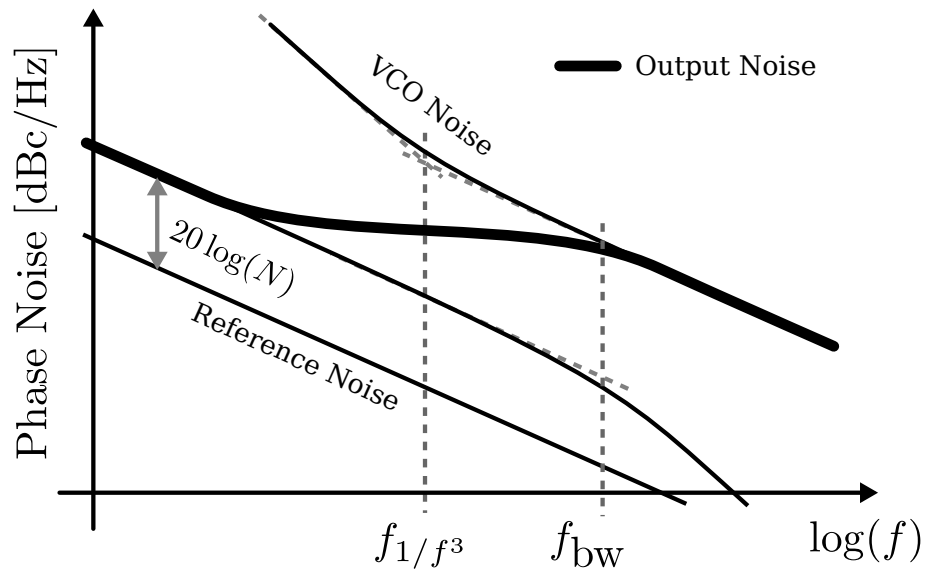


Figure 19: PLL phase noise contributors and output.




Figure 20: Jitter accumulation in a PLL.

3.4 PLL Figure of Merit

Even though the performance and suitability of a PLL for a given application can be assessed by multiple parameters, integrated (RMS) jitter σ_{rms} and power consumption P are often perceived as the most relevant. For such, the figure of merit [38]

$$\text{FOM}_2 = 10 \log_{10} \left[\left(\frac{\sigma_{\text{rms}}}{1\text{sec}} \right)^2 \times \left(\frac{P}{1\text{mW}} \right) \right] \quad (20)$$

yields about -250dB for state-of-the-art PLLs.

3.5 Bang-Bang Phase-Locked Loops

PLLs that have a phase detector that on each reference cycle only report whether the feedback was early or late (instead of indicating how early or late) with respect to the reference are called bang-bang PLLs (BBPLLs). Such a phase detector is known as a bang-bang or binary phase detector (BBPD or BPD). In contrast to the CPPLL's PFD and charge-pump, which together deliver a charge to the loop filter that is linearly proportional to the phase error, the BPD is highly non-linear.

This nonlinearity has several drawbacks. The linear Laplace analysis approach described in Section 3.3 is not applicable. It also complicates the use of traditional $\Sigma\Delta$ fractional feedback dividers. Moreover, since the BPD does not provide a response in proportion to the frequency error, an additional frequency loop is required to bring the VCO to the correct frequency before the PLL can achieve phase lock.

Furthermore, the binary output of the BPD prevents an asymptotic settling towards a zero phase error, even in the absence of random noise, and results in a limit cycle (or orbit) [12]. It settles into an oscillatory steady-state, which is the source of additional jitter.

Despite these drawbacks, a BPD is considerably simpler. It is also the simplest form of a time-to-digital converter (TDC), a 1-bit TDC, which simplifies the design of a digital loop filter and a DPLL. BPDs are also often preferred over linear PDs when implementing clock and data recovery (CDR) circuits [40].

It is instrumental to understand how the BBPLL responds to reference noise. Since the output of the BPD is ± 1 , the proportional path of the PLL switches the VCO (or DCO)




Figure 21: Phase tracking of a bang-bang PLL in the time domain.

between only two output frequencies. This behavior causes the output phase to change at a fixed rate (slew rate) of

$$\left| \frac{d\phi_{out}}{dt} \right| = k_p k_{vco}, \quad (21)$$

where k_p is the proportional gain of the loop-filter in volts and k_{vco} is the gain of the VCO in radians/sec/volt.

An input (reference) signal that is phase-modulated (noisy) by $\phi_{in} = \Theta_m \sin(2\pi f_m t)$ has a maximum slope of $|d\phi_{in}/dt|_{max} = 2\pi f_m \Theta_m$. By equating this to Eq. 21, we can solve for the largest frequency-amplitude product that the PLL can track:

$$f_m \Theta_m = \frac{k_p k_{vco}}{2\pi}. \quad (22)$$

Beyond this maximum, the input phase-noise is attenuated at a rate of 20 dB/dec. The ability of a BBPLL to track the input phase-noise depending on its amplitude and frequency is illustrated in Figure 21.

Even though the dynamics of a BBPLL are nonlinear, the behavior described above can still be represented in the frequency domain if the input noise amplitude Θ_m is a constant. Figure 22 shows the frequency response for different constant input amplitudes and a given proportional path gain, $k_p k_{vco}$.

Furthermore, it has been shown that in the presence of (small amounts) of random noise, from the input or the VCO, the limit cycles may be broken, resulting in an output from the BPD that appears like a random binary sequence which averages to the (noiseless) phase




Figure 22: Frequency response to reference noise in a bang-bang PLL.

error. The behavior of the phase detector becomes then equivalent to a 1-bit $\Sigma\Delta$ ADC [40] and makes it possible to analyze the behavior of a BBPLL with the previously presented linear approach.

3.6 Digital Phase-Locked Loops

3.6.1 Basic Concepts

PLLs that use a digital loop filter are commonly referred to as digital PLLs (DPLLs). Some authors have, at times, used the somewhat misleading term all-digital PLL (ADPLL) instead. The DPLL, as a concept, has been around since the early days of the PLL. However, there has been little motivation to replace analog PLLs in high-performance systems. The fractional-N CPPLL has been the workhorse of frequency synthesis for decades and has been optimized in phase noise, power consumption, and silicon area to such a point that any changes in architecture, including any digital architecture, have resulted in degraded performance. CMOS scaling, however, continues to favor high-speed digital performance over analog, which is not only making DPLLs more attractive but necessary.

The strongest driver for implementing PLLs in digital form has typically been the loop

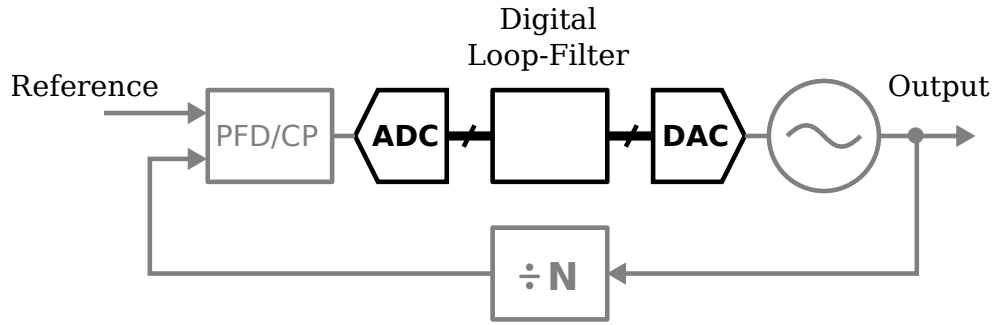


Figure 23: Implementation of a DPLL by using an ADC and a DAC around a digital loop filter.

filter. A digital loop filter can be designed for high configurability, high out-of-band rejection, adaptive spur cancellation, noise immunity, etc. Moreover, it can occupy minimal silicon area. At the same time, the very high leakage in the latest CMOS processes makes it hard to hold charge in an analog loop filter, which results in large reference spurs.

A straight-forward approach to exploit the benefits of the digital loop filter, shown in Figure 23, is by establishing analog-to-digital and digital-to-analog boundaries around the filter using traditional voltage-to-digital and digital-to-voltage converters. Naturally, this results in added complexity and high trade-off penalties related to the quantization errors, non-linearity, and high power consumption of ADCs and DACs.

Current research has focused on structuring the phase detector and the oscillator as data converters themselves. The phase detector must then include some form of a time-to-digital converter (TDC), and the oscillator becomes a digitally-controlled oscillator (DCO). High-performance TDCs and DCOs are the most challenging parts of the DPLL from a circuit design perspective. However, the DPLL's architecture has also been of great focus. It plays an essential goal in exploiting the benefits and circumventing the difficulties arising from the TDC and the DCO.

3.6.2 Phase-Domain DPLL

The most notorious advancement on DPLLs to this date was by Staszewski *et al.*, in [47]. They introduced the first DPLL ever used in a commercial application. This DPLL was used

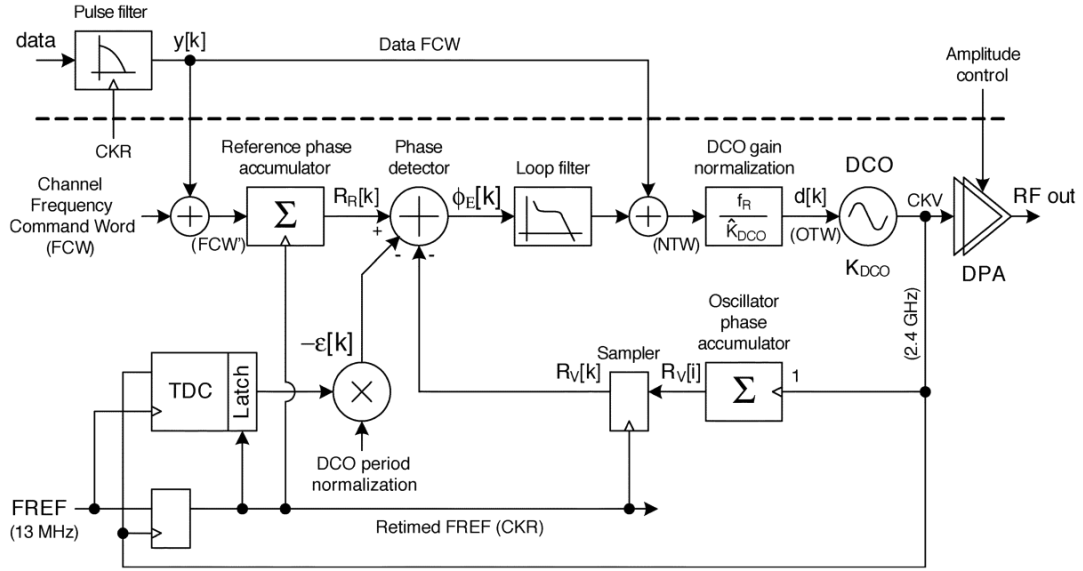


Figure 24: Synchronous phase-domain all-digital PLL-based transmitter [47].

as the local oscillator (LO) in a Bluetooth radio accompanying a digital signal processor (DSP).

Figure 24 shows the architecture of the design in [47], known as phase-domain DPLL [45]. This architecture is only possible in digital form: The DPLL achieves phase-lock by comparing absolute phase between the reference and the DCO. The phase of both signals is stored in counters (or accumulators), incrementing their value by FCW and by one on every reference and DCO cycle respectively. The two values are then subtracted to obtain the phase error. The FCW can contain a fractional part which allows for fractional frequency multiplication of $f_{dco} = FCW \times f_{ref}$ without any dithering. The subtraction of the phases takes place in a synchronous digital circuit which requires the reference or the VCO signal to be re-timed (sampled) by the other. This operation introduces a timing error. However, a TDC measures the error and which subsequently accounted for in the subtraction.

3.7 Other PLL Topics

PLLs are a very active research topic, and researchers are focusing on too many diverse aspects of their behavior and design to summarize them here fully. However, some promising techniques are currently getting significant attention such as the sub-sampling PLL

[17][18][32][10], which avoids the N^2 amplification of reference noise (See Section 3.3), and the delay-locked loop (DLL) [24][14], which can achieve higher bandwidth than PLLs at the cost of higher reference spur.

Other work that has also proven useful in this research is on bang-bang DPLLs [12], on event-driven simulation [46], and on DCO interfacing [49].

CHAPTER IV

TIME-TO-DIGITAL CONVERTERS

4.1 Introduction

Time-to-digital converters (TDCs) are a form of analog-to-digital converter (ADC), where the analog input quantity is time. This time quantity is a time interval, typically defined between two discrete events.

TDCs have found applications in multiple fields, although historically, particle physics has been the field that has most strongly driven the research on accurate electronic time-interval measurements. In principle, the use of time (and distance) to determine the speed of a particle under the influence of a magnetic field, plays a fundamental role in characterizing and identifying the kind of particle being observed [36].

The focus on TDCs has recently shifted with the increased interest in implementing most sub-circuits of a frequency synthesizer in digital form.

As for voltage or current ADCs, the performance of TDCs is primarily specified by resolution, linearity and power consumption. Contrary to other data converters, however, TDCs have no control over the sampling rate. Consider Figure 1(b) for example. The analog quantity is the width of each pulse. However, the sampling rate is the pulse rate of the signal, over which the TDC has no control.

When a TDC is used as a phase detector in a DPLL, the designer's main concern is how the TDC will impact FOM_2 (Eq. 20). The TDC will impact phase noise by introducing quantization noise which will establish an in-band phase noise floor given by

$$S_\phi = \frac{(2\pi\sigma_{\text{tdc}}f_{\text{vco}})^2}{f_{\text{ref}}} \left[\frac{\text{rad}^2}{\text{Hz}} \right] \quad (23)$$

where f_{vco} and f_{ref} are the frequencies of the VCO and the reference respectively, and σ_{tdc} is the RMS quantization error in the TDC.

Other characteristics of TDCs can further degrade the performance of a PLL. The TDC

can limit the bandwidth, as I show in Section 4.2, which reduces the PLL’s ability to suppress the VCO’s phase noise. Furthermore, nonlinearity in the TDC’s response can create mixing products that result in spurs at low-frequency offsets, which are not attenuated by the loop dynamics.

TDC techniques have suffered, historically, from low resolution, low linearity, high power consumption, high sensitivity to device and power supply noise, as well as low bandwidth. Recent developments using ring oscillators have allowed, however, for first-order [48][15], second-order [51], and even third-order [16] quantization noise shaping. This technique shifts (or shapes) the quantization noise to higher frequencies, which can then be attenuated by low-pass filtering. It has enabled a tradeoff between bandwidth and precision.

Unfortunately, bandwidth and phase detector precision are both essential to achieving low phase noise in PLLs, therefore exercising this tradeoff might provide limited improvement. This prompts for methods to increase TDC precision without reducing the bandwidth.

4.2 Quantization Noise

Since TDCs are ADCs, it is instrumental to understand quantization noise, which is the primary cause of distortion introduced by the quantization process. Quantization noise is the result of limited resolution and becomes present even before any circuit imperfections are taken into consideration.

The input T_{in} to an ideal uniform quantizer, normalized to its resolution τ , can be written in the form

$$\frac{T_{\text{in}}}{\tau} = n + v, \tag{24}$$

where n is the integer part, and $v \in [0, 1)$ is the fractional part. Then, if the quantized output is $Q\{T_{\text{in}}\} = n$, the quantization error is v .

When v is properly randomized across multiple samples such that $v \sim U(0, 1)$ [4][44], the RMS error in the quantized output, with respect to the input, is then the standard deviation of v , $\sigma_\tau = 1/\sqrt{12}$. Moreover, the error in a measurement consisting of the average




Figure 25: First-order $\Delta\Sigma$ modulator.

of m samples with the same statistical properties is then

$$\sigma_m = \frac{\sigma_\tau}{\sqrt{m}}. \quad (25)$$

This provides a straightforward way to reduce the effective quantization noise: averaging. This approach, however, requires collecting m times as many samples per unit time than the Nyquist rate. For this reason, this kind of approach is referred to as oversampling.

Through the use of oversampling and noise shaping, we can achieve further improvement than what is possible by just averaging. Delta-Sigma ($\Delta\Sigma$) modulation is a common technique to achieve quantization noise shaping. Figure 25 shows the block diagram of a first-order $\Delta\Sigma$ modulator. The Δ operation is the difference between the input $X(z) = T_{\text{in}}(z)$ and output $Y(z)$, and the Σ operation is the integration that follows. Finally, the output of the integrator is quantized by a uniform quantizer (represented by the quantization error or noise $q(z)$ that it introduces) to produce the output. By following the diagram, we can write the output as

$$Y(z) = X(z)z^{-1} + q(z)(1 - z^{-1}). \quad (26)$$

The output is a delayed version of the input plus the first difference of the quantization noise. This last term is known as the noise transfer function (NTF). By replacing z with $e^{j2\pi fT}$, where T is the sampling period, we obtain $\text{NTF}(f) = 2 \sin(\pi fT)$, which is plotted in Figure 26.

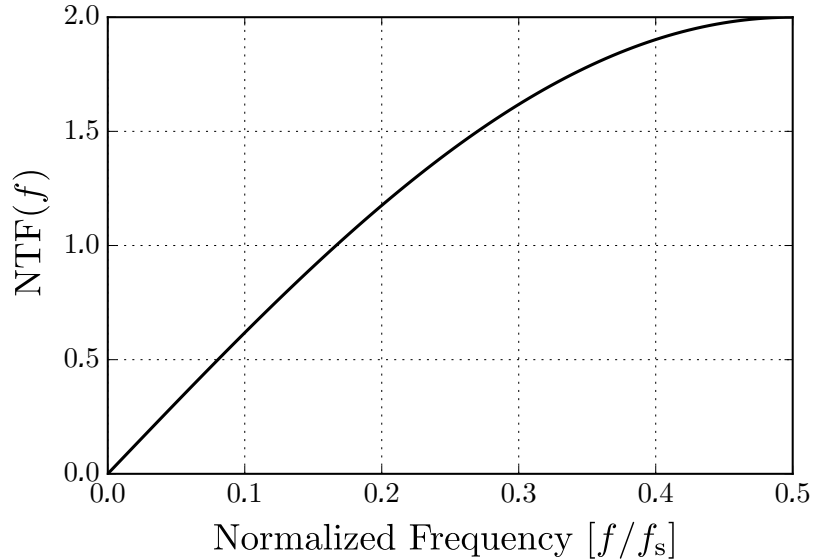


Figure 26: Noise transfer function of a first-order $\Delta\Sigma$ modulator.

It becomes clear that the quantization noise increases with frequency (relative to the sampling frequency $1/T$), therefore the signal-to-noise ratio (SNR) can be maximized by limiting the signal bandwidth or increasing the sampling rate.

The in-band quantization noise power in $\Sigma\Delta$ converters is given by [43]

$$\sigma_{\text{tdc}}^2 = \frac{\pi^{2L} \sigma_q^2}{(2L + 1)(\text{OSR})^{2L+1}} \quad (27)$$

where L is the modulation order, σ_q is the RMS value of the quantization noise q , and OSR is the oversampling ratio (Figure 27). The OSR is the sampling frequency in multiples of the Nyquist frequency for the given input signal.

4.3 Overview of TDC Techniques

4.3.1 Basic Concepts

A time interval can be described by two events respectively defining the *start* and the *stop* times of the interval (Figure 28(a)). Alternatively, it can be described by a single event happening at time t_{event} with duration T_{in} . The TDC's function is to measure T_{in} and to provide the result in digital form.

The most straightforward approach to digitally quantifying a time interval is by counting

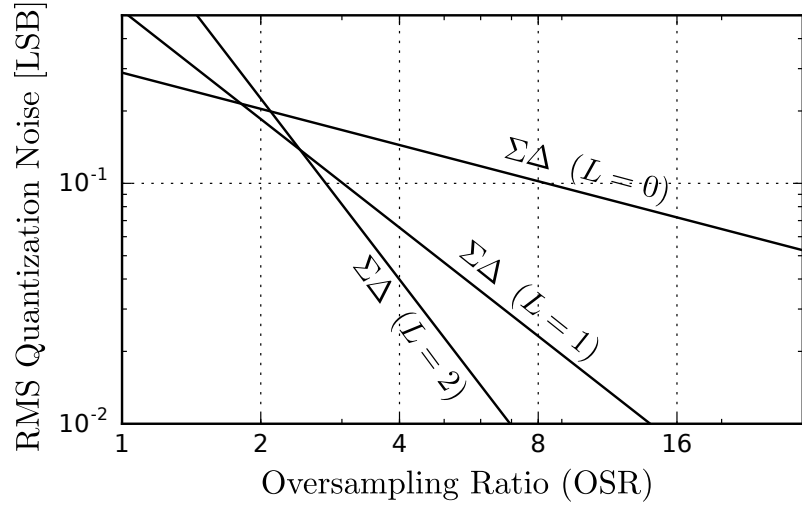


Figure 27: Effective quantization noise in $\Delta\Sigma$ converters of different order (L) as a function of oversampling ratio (OSR).

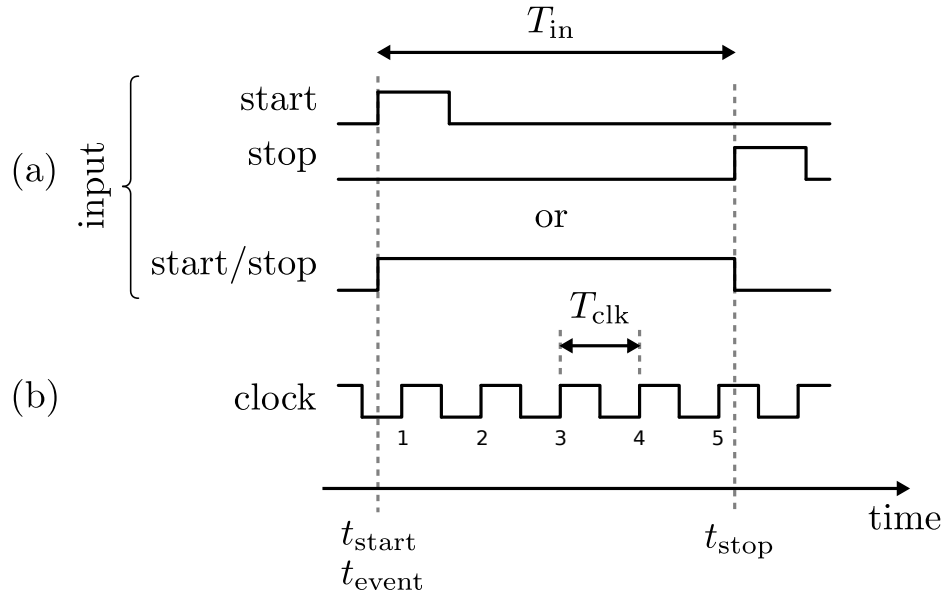


Figure 28: TDC Input Signals: (a) Event signals and (b) Reference clock.

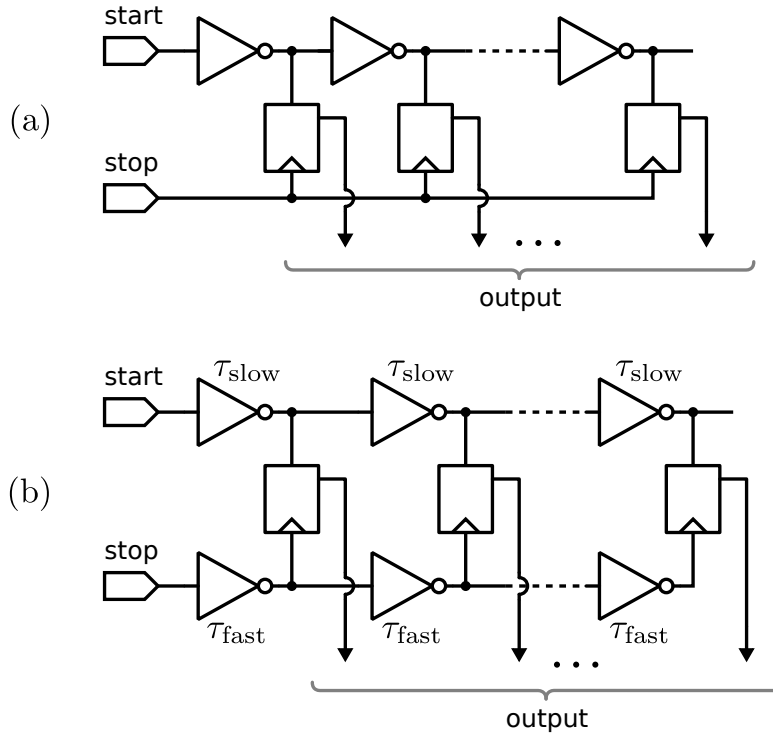


Figure 29: Delay-line-based TDCs. (a) Flash TDC and (b) Vernier TDC.

clock cycles for the duration of the interval (Figure 28(b)). The resolution of such measurement is the period of the clock, T_{clk} . It is limited only by the speed of digital logic, which is only predicted to improve. However, a counter running at an extremely high frequency will consume significant power, making this approach unattractive for high-resolution applications. Instead, digital counters are often used in TDCs to accommodate large input time intervals, and some other technique is used to achieve high resolution at the same time [36][26].

4.3.2 Delay Line TDCs

The Flash (or delay line) TDC is convenient for its simplicity and is appropriate for measuring short time intervals with better resolution than by the use of counters. Its architecture is shown in Figure 29(a).

It consists of a delay-line of logic gates (inverters, typically) into which the *start* signal is injected, and the *stop* signal is used as the clock to the flip-flops. The delay of each gate,

τ_{gate} , provides a mapping between the interval’s duration and “distance” of propagation in number of gates. This information is captured in the flip-flops. The measurement resolution of this TDC is, therefore, τ_{gate} .

The resolution will only get better with the continuous scaling of CMOS devices. However, the method introduces several challenges. The delay of an inverter (or any logic gate) is a very unpredictable quantity [1], not only across process, voltage and temperature (PVT), but also across neighboring devices (matching) in the same die [33]. Delay mismatch may result in large nonlinearity. Moreover, flash TDCs also require a large number of gates in the delay line to support large time intervals, and it is also important to highlight that this architecture cannot resolve negative time intervals (when the *start* signal arrives after the *stop* signal).

The Vernier delay-line TDC, shown in Figure 29(b), improves further on resolution. It uses two delay lines, with different nominal gate delays, τ_{fast} and τ_{slow} , for each of the input signals. The *start* signal is launched into the slower line and the *stop* signal into the faster one. Eventually, both signals will have traveled through the same number of inverters, M , i.e., the *stop* signal will “catch up” with the *start* signal, where

$$\begin{aligned} M\tau_{\text{fast}} + T_{\text{in}} &= M\tau_{\text{slow}} \\ \Rightarrow T_{\text{in}} &= M(\tau_{\text{slow}} - \tau_{\text{fast}}). \end{aligned} \tag{28}$$

Since M is captured in the flip-flops after both signals have propagated through the lines, the input interval $T_{\text{in}} = t_{\text{stop}} - t_{\text{start}}$ can be computed with a resolution of $\Delta\tau = \tau_{\text{slow}} - \tau_{\text{fast}}$.

This improvement in resolution does not help with the inherent linearity problem in delay-line based TDCs, however. Furthermore, this technique exacerbates the problem of requiring a large number of gates and increases proportionally with resolution.

The Vernier technique is not exclusive to delay lines. It can also be implemented using two counters, for example, driven by clocks with slightly different frequencies [3] and, in fact, precedes the Vernier delay-line TDC.

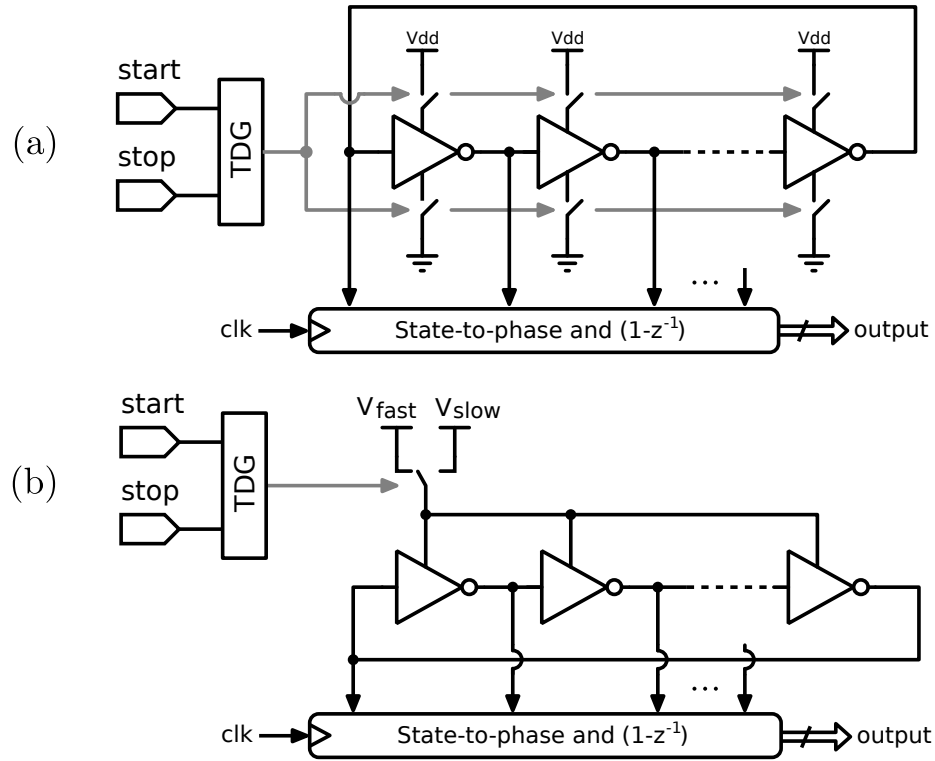


Figure 30: Ring-oscillator-based noise-shaping TDCs. (a) Gated ring oscillator (GRO) TDC, (b) Switched ring oscillator (SRO) TDC.

4.3.3 Noise Shaping TDCs

The frequency-to-phase integration nature of oscillators and the discrete phases of ring oscillator are convenient features for the implementation of the integration and quantization operations required in $\Sigma\Delta$ -based ADCs. This property has for long been exploited in conventional ADCs [25] before being adopted in TDCs, which was reported for the first time with the introduction of the Gated Ring Oscillator (GRO) TDC [48].

The GRO TDC, shown in Figure 30(a), operates by switching (gating) a ring oscillator *on* and *off*, specifically by disconnecting the inverters from their supply rails. The *on* and *off* states correspond to the logic state of a periodic input signal. Therefore, the total phase change in the oscillator during one period of the input is directly proportional to the input signal's duty cycle.

The phase accumulation during each cycle is determined by first sampling the entire

ring, capturing the logic state at each node with flip-flops, for example. Then, the phase, modulo- 2π , can be computed from the sampled state of the ring. Finally, the output is generated by subtracting the phase computed in the previous cycle from the current phase. The “raw” resolution of the output is the delay of one inverter. However, the quantization error is noise-shaped. The operation of the GRO TDC is consistent with $\Delta\Sigma$ modulation as described earlier except for the order of Δ and Σ : The ring oscillator performs the integration (Σ) first and the difference (Δ) is computed digitally afterward.

One weakness in this TDC is that it is significantly affected by leakage. When the ring is stopped, the parasitic capacitance holding the state of the ring loses charge. This loss of charge introduces an additional source of error, resulting in performance below that of first-order noise shaping.

The Switched Ring Oscillator (SRO) TDC [15], shown in Figure 30(b), mitigates the leakage problem of the GRO with a slight modification. Instead of switching between f_{ring} and $f = 0$, it never entirely stops the ring and switches between two frequencies f_{fast} and f_{slow} .

It is important to note that both architectures require a time-difference generator (TDG) to provide a single line that controls the gating or switching. This block generates a control signal with a variable duty cycle depending on the relative arrival times of the rising and falling edges at the inputs. Therefore, these TDCs, as presented in [48] and [15], require periodic input signals with known duty cycles.

4.3.4 Other TDC Techniques

An intuitive approach to improving the resolution of any data converter is to amplify the input signal before quantization. This intuition has led to the concept of time amplification. Attempts to design an analog “Time Amplifier” have been reported in the literature [29], but since they rely heavily on analog performance, they are of little practical use in modern CMOS.

An alternative to ring oscillator based TDCs providing sub-gate delay resolution and linearity that is also actively being researched is the stochastic TDC [28]. In particular, the

techniques' performance is independent of gate delay and mismatch and depends instead on the accuracy of the reference clock.

CHAPTER V

DESIGN OF HYBRID DIGITAL-ANALOG PHASE-LOCKED LOOPS

5.1 Introduction

The complementary nature of reference noise filtering (H_{ref}) and VCO noise filtering (H_{vco}) in traditional PLL architectures, as described in Section 3.3, limits the applicability of PLLs for solving certain frequency synthesis problems. One of which is when a low jitter signal with frequency f_0 is to be recovered from a signal with an average frequency f_0 and subject to frequency/phase modulation or noise.

I propose a hybrid analog and digital architecture that decouples H_{ref} and H_{vco} . This decoupling opens several possibilities. One such possibility is the implementation of a very low bandwidth PLL that filters most of the reference noise without exposing additional VCO noise.

The approach consists of two interdependent PLL loops such that one can be optimized for low bandwidth H_{ref} , which processes the reference, and the other, optimized for high bandwidth H_{vco} , which processes the VCO. This decoupled optimization is shown in Figure 31.

5.2 Proposed Architecture

The proposed architecture, first presented by us in [9], is shown in Figure 32.

An outer (or “master”) loop is a fully digital bang-bang loop. The DCO in this loop, implemented by a fractional-N CPPLL, is the inner loop. The outer loop commands the inner loop through the DCO’s digital frequency control word (FCW), which is the CPPLL’s feedback divider value, and takes the CPPLL’s VCO output as feedback.

The CPPLL is phase-locked to its own low phase-noise reference of frequency f_{xtal} . Since there is insignificant phase noise originating at the reference, H_{ref} can be sacrificed in favor of H_{vco} . The DPLL, in contrast, has a minimal phase-noise DCO (the already optimized CPPLL), so its H_{vco} can be sacrificed in favor of its H_{ref} . The H_{ref} of the DPLL and H_{vco}




Figure 31: PLL noise transfer functions optimized for (a) VCO noise and (b) reference noise.

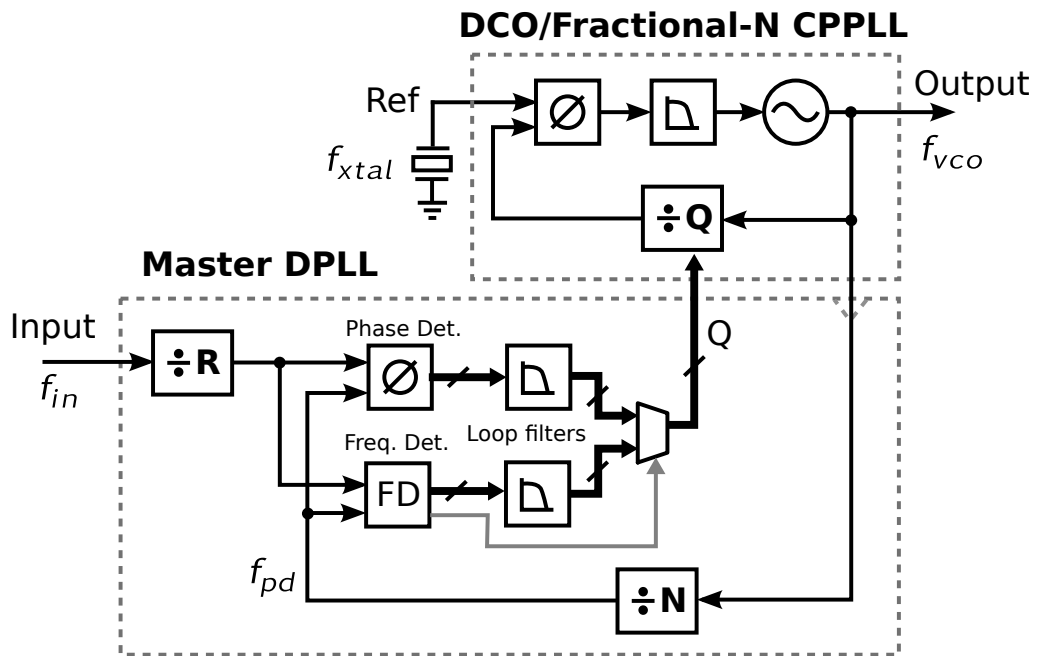


Figure 32: DPLL using a fractional-N CPPLL as a DCO.




Figure 33: Output phase noise contributors in the proposed hybrid PLL.

of the CPPLL determine the system’s noise dynamics and are independent of each other.

Figure 33 illustrates how the different phase noise sources in the overall system contribute to the output phase noise. Above the bandwidth of the CPPLL, f_{bw1} , the phase noise of the VCO is not attenuated and appears directly at the output. Below f_{bw1} the VCO’s phase noise is attenuated until the phase noise of the crystal reference dominates. At further lower frequencies, approaching the bandwidth of the DPLL, f_{bw2} , the input noise starts to dominate, appearing directly at the output.

The improvement with respect to a single CPPLL configured to a very low bandwidth f_{bw2} is the region where the output phase noise drops below the VCO noise.

5.3 Implementation

I developed the original prototype of the DPLL [9] by synthesizing the design from Verilog RTL into an FPGA. The CPPLL, in a separate IC, was controlled in real time via SPI from the FPGA. Since then, several commercial products have been implemented based on this architecture, where the DPLL has been automatically synthesized, placed and routed alongside a hard-macro of the CPPLL in the same IC.

Figure 32 shows two feedback paths (or loops) implemented in the DPPLL that connect

the VCO in the CPPLL to its feedback divider. One path is a bang-bang phase-detection loop and the other a frequency-detection loop. As I explained earlier in Section 3.5, bang-bang loops requires the difference between reference and feedback frequencies be small before locking can be achieved. For this reason, the frequency detector (FD) monitors when the frequency error falls below a given threshold, during which it switches the multiplexer allowing the phase-detection path to control the DCO.

The loop filters (to the right of the phase and frequency detectors), were implemented as $k_p + k_i/s$ stages, where k_p and k_i were programmable gains. Additional logic was also included which would step k_p and k_i from a large to a small and final value based on a lock detector (not shown) in order to speed up the settling process.

There is little concern about circuit non-idealities in this architecture, especially since the DPLL is fully digital and the CPPLL is a mature and proven circuit. In particular, when we treat the CPPLL as a DCO, its frequency resolution is determined by the number of bits of the $\Delta\Sigma$ fractional-N feedback divider. Adding resolution to the $\Delta\Sigma$ modulator is of little cost. In all of our implementations, 12 bits of fractional resolution were available in the CPPLL, which resulted in negligible quantization noise. Furthermore, due to the feedback nature of PLLs, the frequency versus feedback-divider-value response of the CPPLL can be considered perfectly linear.

Despite these benefits, the interface between the DPLL and the CPPLL must be designed carefully. In particular, the update of the feedback divider value in the CPPLL may happen asynchronously and at a different rate than provided by the DPLL. Such a problem can be addressed by using a circular buffer or FIFO. With this form of memory, the DPLL always writes to the latest position M . At the time of the *write* operation the *read* address is updated to $M - k$, where k is a small number, and the *write* address is updated to $M + 1$. The *read* address is kept in gray-code and is resampled by the CPPLL on every update. This approach ensures that the values written by the DPLL are loaded into the CPPLL monotonically and error free.

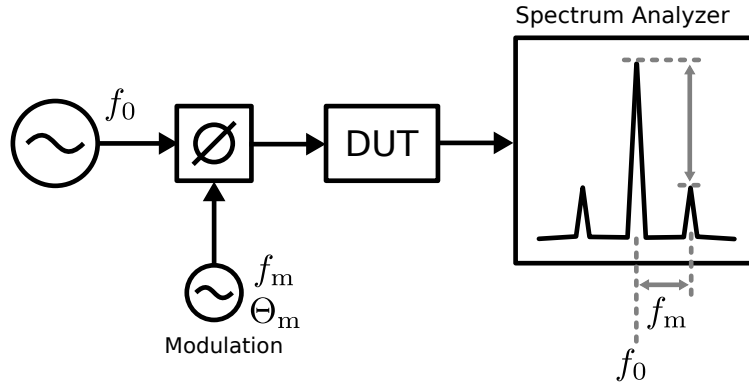


Figure 34: Test setup to determine the response of the hybrid PLL to input phase noise.

5.4 Applications and Results

The system was verified and characterized in two steps. The loop dynamics of the system were measured first. Then, the ability to recover low phase-noise unmodulated tones from modulated signals was assessed.

5.4.1 Bang-Bang Behavior

The overall response of the complete DPLL was characterized by measuring the relative amplitude of the sidebands when the input is a single-tone phase-modulated carrier (Figure 34). For every chosen modulation amplitude Θ_m , the modulation frequency f_m was swept in discrete steps.

The results for two modulation amplitudes and two different proportional gain (k_p) settings are shown in Figure 35. This confirms the bang-bang behavior as depicted in Figure 22.

5.4.2 Gapped Clock

The first case of modulated input that I studied was that of a “gapped” clock. As the name implies, it is a clock signal with gaps or missing pulses. It occurs at the merging point of two data networks, where incoming non-payload packets like error correction or routing information data shall not be transmitted to the second network. To discard them, the corresponding clock pulses on the receiving side that write the packets into a FIFO memory




Figure 35: Effect of loop-filter gain and modulation amplitude on the frequency response of the DPLL.




Figure 36: FIFO buffer for two merging networks with “gapped” *write* clock and recovered *read* clock.

are “swallowed”. On the transmitting side, the *read* clock must be a clean unmodulated signal with the same average frequency as, and recovered from, the *write* clock. The relationship between the *read* and *write* clocks for the case of one non-payload packet for every four packets is shown in Figure 36.

Gapped clocks were applied to the input of the DPLL. The average frequency was kept constant at 197 MHz and the gap rate was set to 1 in every p cycles, for $p = 5, 5.25$ and 63. A fractional rate of 5.25 is the average of 1-in-5 for one time followed by 1-in-4 for four times.

The DPLL was configured to the lowest possible stable bandwidth and the CPPLL was stabilized by a crystal reference with $f_{\text{xtal}} = 50$ MHz. An $f_{\text{out}} = 197$ MHz was achieved by dividing $f_{\text{vco}} = 2364$ MHz by 12.

Figure 37 presents the measured phase noise and integrated (RMS) jitter. These curves




Figure 37: Measured output phase noise of the proposed DPLL with gapped-clock input.

can be compared to the phase noise of the CPPLL in stand-alone mode (labeled OL), i.e. without the feedback divider value being modulated by the DPLL. The gapped clock cases have higher phase noise only at very low frequencies as predicted in Figure 33.

5.4.3 Spread Spectrum Demodulation

Spread-spectrum modulation is a technique that spreads the power of a signal across a wide frequency range. This ensures that the signal does not exceed a specified power density at any specific frequency, reducing the interference between communication channels.

The clock of a transmitter can be spread-spectrum modulated for this purpose. However, a clean, stable clock must be recovered at the receiver. Figure 38 shows the spectrum of a spread-spectrum modulated signal and the clean tone recovered by the DPLL.

5.4.4 Large Multiplication Ratios

This architecture is also effective for very large multiplication ratios. In contrast with the previous two applications, where low bandwidth is the primary requirement, the stability of the VCO or DCO (in this case, the CPPLL) is the enabling feature. Since a correction from a low-frequency reference is applied only once every many oscillator cycles, it is imperative that little drift (less than a VCO cycle) occurs between corrections. Otherwise, a “cycle




Figure 38: Spread-spectrum demodulation with the proposed DPLL. Spectrum of the input and output signals.

slip” may occur and thus the loss of phase lock.

Figure 39 shows the measured phase noise and integrated (RMS) jitter for different frequency multiplication ratios (from 10×10^3 to 30×10^3). For proper comparison, the VCO frequency was kept constant at $f_{\text{vco}} = 960$ MHz, while the reference frequency was changed.

5.5 Conclusions and Proposed Research

I have presented theory, implementation and measurement data for a technique that allows us to have independent and decoupled filtering of reference and oscillator noise in a PLL. In particular, it allows for very low bandwidth low-pass response from the input to the output and a very high bandwidth high-pass response from the VCO to the output.

This has been made possible by an approach combining digital and analog PLL loops which can be optimized and configured separately. Furthermore, this approach lends itself for easy prototyping, such as with an FPGA and a stand-alone CPPLL IC, as well as for a reliable integrated implementation on silicon, which has been proven through multiple

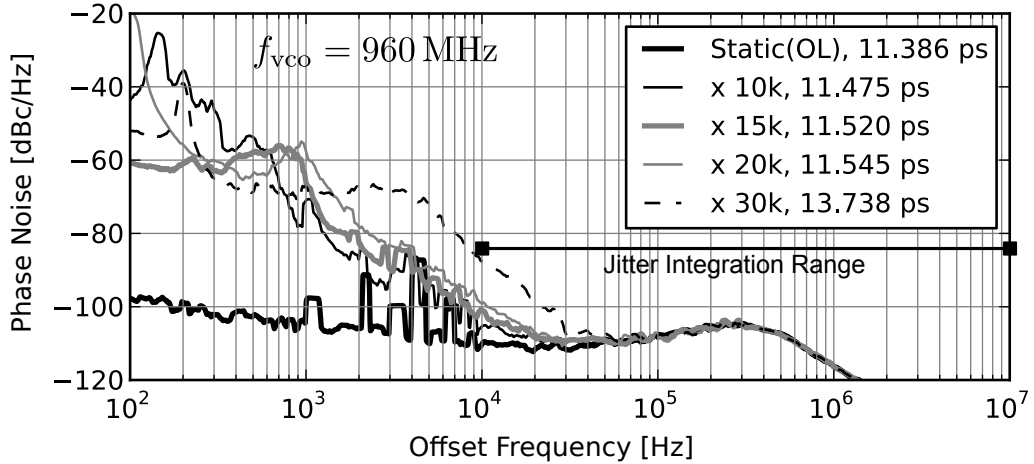


Figure 39: Measured output phase noise for the proposed DPLL with large multiplication ratios.

commercial IP implementations.

Even though this architecture was initially conceived for input noise filtering (i.e., low input-to-output bandwidth), there are further applications not yet explored. Consider the case in Figure 40 consisting of the same original architecture. However, the CPPLL is ring-oscillator-based and its reference is a free-running LC oscillator. The DPLL has a crystal oscillator reference. The benefit of this configuration may not be evident at first until we analyze how the different phase noise sources contribute to the output.

Figure 41 shows the contributors and total output phase noise for the configuration in Figure 40. We know that at very high frequencies, the phase noise of the VCO in the CPPLL appears unattenuated at the output. Below f_{bw1} , it drops until the phase noise of the reference free-running LC oscillator dominates. It is important to highlight that one can make f_{bw1} very large since the LC oscillator may run at several gigahertz. This is much faster than crystal oscillators which range typically between 25 MHz and 100 MHz.

This behavior indicates that the phase noise performance of the CPPLL, at least up to f_{bw1} , is as if it was implemented using an LC VCO. However, it is still using a ring VCO, which has multiple benefits (Table 1) such as a large tuning range.

Since the LC oscillator is free-running, the CPPLL needs to be frequency-stabilized and

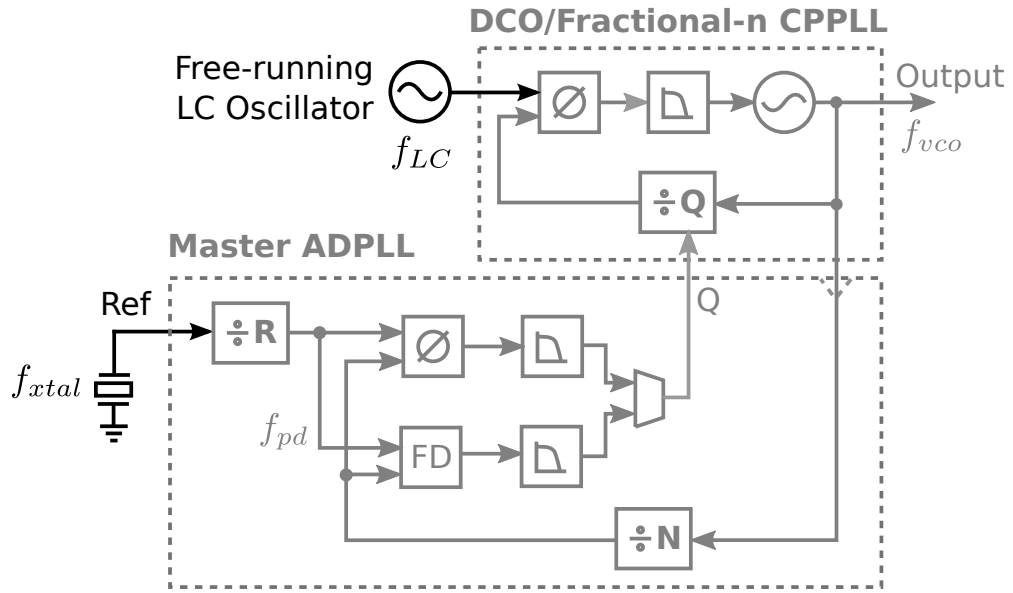


Figure 40: Proposed hybrid DPLL with an alternative set of reference signals.




Figure 41: Output phase noise contributors in hybrid PLL with LC and crystal references.

the low-frequency phase noise needs to be suppressed. These requirements are achieved by the DPLL, which in this case is configured for high bandwidth, f_{bw2} . Below this frequency, the output phase noise does not follow the phase noise of the LC oscillator and instead drops down to the phase noise of the crystal reference.

This architecture, then, does not only allow for low-bandwidth low-pass and high-bandwidth high-pass simultaneous filtering, as demonstrated by our prototype and measurement results. It may also allow for the benefits of both LC and ring VCOs at the same time.

CHAPTER VI

DESIGN OF A RING VCO WITH IMPROVED FOM BY ADDING AN INDUCTIVE LOAD

6.1 Introduction

In this chapter, I present a novel VCO architecture [6] that behaves similarly to a standard ring VCO although exhibiting better phase noise. It is a hybrid approach between ring VCOs and LC VCOs, which improves on FOM_1 (Eq. 6) by trading off area and tuning range. This solution breaks the binary choice between LC and ring VCOs, allowing for trade-offs between the properties shown in Table 1.

6.2 Proposed Architecture

The proposed VCO is shown in Figure 42. It is a regular inverter-based single-ended ring oscillator (RO) with the output of each stage connected to a common center node through identical inductors.

The inductive load structure is symmetric with respect to the three phases of the RO. This symmetry prevents DC current from flowing through the inductors and makes the center node, at a voltage $V_c \approx V_{ring}/2$, a “virtual ground” at the frequency of oscillation f_0 and its harmonics. At multiples of $3f_0$, however, it appears as an open circuit. This virtual




Figure 42: Proposed ring VCO with an inductive load.

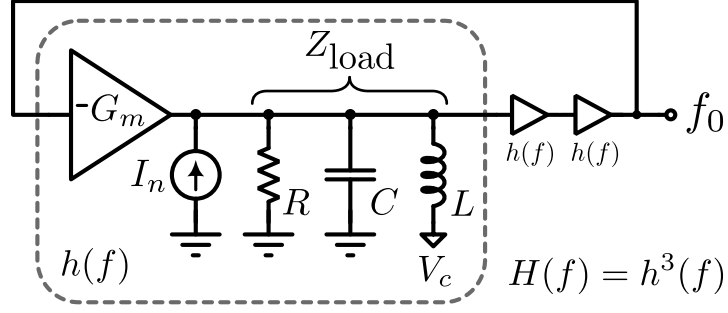


Figure 43: Linear model of a ring oscillator with an inductive load and noise current source.

ground allows for analysis of this circuit as if V_c was a physical AC ground, and to treat the whole circuit as three identical and independent cells connected as shown in Figure 43. Even though this model does not capture the non-linear nature of the CMOS inverters, it suffices for our purposes as a tool for improving the FOM of ring oscillators.

One may suspect that the inductors provide energy storage, therefore reducing power consumption and the FOM. However, the effect on PN/jitter and frequency are not immediately clear.

6.3 Frequency of Oscillation

The frequency of a standard (or “bare”) inverter-based RO (f_{ring}) is given by Eq. 7 (repeated for convenience in Eq. 29), which depends directly on the delay of each inverter (τ).

$$f_{\text{ring}} = \frac{1}{2N\tau} \quad (29)$$

This delay can be approximated by Eq. 8 (repeated for convenience in Eq. 30), which represents the time to charge a capacitance C_{gate} (the input capacitance of the next inverter) by a fixed current I_{sat} , up to a voltage of $V_{\text{ring}}/2$.

$$\tau = \frac{V_{\text{ring}}}{2} \times \frac{C_{\text{gate}}}{I_{\text{sat}}} \quad (30)$$

If an inductor is connected in parallel to the capacitive load, the inverter will “see” an impedance

$$Z_{LC}(f) = \left(j2\pi f C_{\text{gate}} + \frac{1}{j2\pi f L} \right)^{-1} \quad (31)$$

which resonates at

$$f_{\text{res}} = \frac{1}{2\pi\sqrt{LC_{\text{gate}}}}. \quad (32)$$

When $f > f_{\text{res}}$, the parallel LC tank appears purely capacitive, i.e.,

$$Z_{LC}(f) = \frac{1}{j2\pi f C_{\text{equiv}}} \quad (33)$$

where

$$C_{\text{equiv}} = C_{\text{gate}} - \frac{1}{(2\pi f)^2 L} \quad (34)$$

$$< C_{\text{gate}}. \quad (35)$$

By using C_{equiv} instead of C_{gate} in Eq. 30, we obtain a new frequency of oscillation f_0 . This can be written in terms of f_0 , f_{ring} , and f_{res} as $f_0^2 - f_0 f_{\text{ring}} - f_{\text{res}}^2 = 0$. Solving for a positive f_0 yields

$$f_0 = \frac{f_{\text{ring}} + \sqrt{f_{\text{ring}}^2 + 4f_{\text{res}}^2}}{2}. \quad (36)$$

A higher frequency of oscillation than f_{ring} and C_{equiv} lower than C_{gate} can be explained intuitively. The physical capacitance C_{gate} is now charged partially by current from the inductor. Therefore, only a fraction of the total required charge, also in a fraction of the time, needs to be supplied by the inverter for the voltage on C_{gate} to trigger the next inverter in the ring.

6.4 Frequency Tuning Range

The previous analysis yields some insight on how much f_0 changes within a given range of V_{ring} as a function of L . In an inductor-less ring oscillator, the amount of charge that an inverter needs to deliver to C_{gate} to trigger the next inverter is more or less constant and

proportional to C_{gate} . A change in V_{ring} modulates the current and therefore the time it takes to deliver this charge. With the inductive load present, however, each inverter is now delivering charge to an equivalent C_{equiv} , which is only a fraction of the physical capacitance C_{gate} . Therefore, the gate-delay τ can only be modulated in proportion to C_{equiv} . Since C_{equiv} decreases as L decreases, so does the tuning range.

From Eq. 34 we can calculate the absolute minimum inductance where both C_{equiv} and the tuning range drop to zero:

$$L_{\text{min}} = \frac{1}{(2\pi f)^2 C_{\text{gate}}} \quad (37)$$

6.5 Phase Noise and Jitter

Eq. 15 hints at how the proposed architecture may improve jitter and PN. Faster overall gates (smaller τ) and lower gate capacitance (lower C_{equiv}) are possible based on Eqs. 30 and 34. Larger I_{sat} also decreases σ_τ but increases power consumption which may reduce the overall FOM.

In our architecture, the current that is charging/discharging C_{gate} is not only I_{sat} but the current from the inductors as well. Therefore, the total current charging C_{gate} can be larger even as I_{sat} becomes lower. This reduces σ_τ and P at the same time, which points to a lower FOM.

To quantify the improvement in PN by adding the inductors, we can resort to the phase noise analysis by linear modeling described in Section 2.4.1 and is applied to ring oscillators in Section 2.5. It follows from Eqs. 13 and 16 that the attenuation factor of noise being up-converted from baseband to $f_0 + \Delta f$ in a 3-stage “bare” ring oscillator is

$$m_{\text{bare}}(f_0) = \left| \frac{dH}{df} \right|_{f=f_0}^2 = \frac{27}{16\pi^2 f_0^2}. \quad (38)$$

For the 3-stage RO with inductors, applying the same criterion results in

$$m_{\text{ind}}(f_0) = \left| \frac{dH}{df} \right|_{f=f_0}^2 = \frac{27}{16\pi^2 f_0^2} \left[\frac{\left(\frac{f_0}{f_{\text{res}}}\right)^2 + 1}{\left(\frac{f_0}{f_{\text{res}}}\right)^2 - 1} \right]^2. \quad (39)$$




Figure 44: Simplified schematics of the PLL for testing the proposed VCO prototype.

We can verify that when $f_{\text{res}} \rightarrow 0$, which is equivalent to removing the inductors, $m_{\text{ind}} \rightarrow m_{\text{bare}}$.

The reduction of noise by adding the inductors is then

$$k = \frac{m_{\text{bare}}}{m_{\text{ind}}} = \left[\frac{\left(\frac{f_0}{f_{\text{res}}}\right)^2 - 1}{\left(\frac{f_0}{f_{\text{res}}}\right)^2 + 1} \right]^2. \quad (40)$$

Since f_0 and f_{res} are not independent, it is convenient to present k in terms of $f_{\text{ring}}/f_{\text{res}}$ where

$$k \approx \frac{1}{4} \left(\frac{f_{\text{ring}}}{f_{\text{res}}} \right)^2. \quad (41)$$

Therefore, since a lower inductance will yield a higher f_{res} , it will result in a lower k , lower PN and higher FOM.

6.6 Prototype and Measurement Results

The proof-of-concept prototype consisted of a 3-stage ring VCO with individual single-chip unbuffered CMOS inverters and discrete surface-mount inductors. In order to measure PN, the VCO was stabilized by a PLL. The PLL bandwidth was intentionally kept very low in order to expose a large frequency offset (Δf) range of the VCO's PN.

The simplified schematics of the prototype are shown in Figure 44 and include the VCO, the PLL IC and all other discrete components. A photograph of the prototype is shown in Figure 45.




Figure 45: Photograph of the PLL for testing the proposed VCO prototype implemented on a PCB.

The prototype was characterized by determining f_0 and I_{ring} versus V_{ring} of the free-running VCO (Figure 46) and then measuring PN in closed-loop. This was repeated for a VCO without inductors and then with different inductor values.

The PN normalized to the frequency and power of the inductor-less case along with the corresponding FOM are shown in Figure 47. All cases in this figure are for constant $V_{\text{ring}} = 4.5\text{V}$, which is also constant power.

To compute the measured improvement in the FOM, ΔFOM_M , for a given inductance, I took the difference (in dB) between the FOM for that L and the inductor-less (bare). Then I averaged across Δf between 10 kHz and 1 MHz. This band is the least affected by the PLL dynamics and the instrument's noise floor.

To compute the predicted improvement in FOM, ΔFOM_P , for a given inductance, I first computed the predicted reduction in PN, k . For this, in addition to the data in Figure 46, we only need to know f_{res} , which in turn requires knowing C_{gate} . C_{gate} can be directly approximated (ignoring crowbar current) from the data of the inductor-less case in Figure 46 by

$$C_{\text{gate}} = \frac{I_{\text{ring}}}{N V_{\text{ring}} f_0}, \quad (42)$$

where N is the number of inverters in the ring.

Finally,




Figure 46: Frequency and current versus voltage for different inductance values.




Figure 47: PN and FOM comparison between the ring oscillator without inductor and with different inductor values when $V_{\text{ring}} = 4.5\text{V}$.




Figure 48: Measurements and predictions summary for the proposed VCO when $V_{\text{ring}} = 4.5\text{V}$.

$$\Delta FOM_P = \frac{FOM_P}{FOM_{\text{bare}}} = \left(\frac{f_0}{f_{\text{ring}}} \right)^2 \times \frac{1}{k}, \quad (43)$$

where FOM_P is the predicted FOM (which we do not need to calculate directly) and FOM_{bare} is the FOM of the bare RO. The $(f_0/f_{\text{ring}})^2$ term is the normalization for the change in frequency.

Figure 48 shows the summary of measurements versus predictions for frequency and improvement in FOM across inductance values.

6.7 Conclusions and Proposed Research

I have proposed a VCO architecture that allows for better PN-power FOM than that of a traditional inverter-based ring VCO. I developed a prototype exhibiting 8 dB of FOM improvement over a simple ring VCO. I also accurately predicted this improvement as well as the increase in frequency.

The VCO retains the supply-voltage/current-control and high linearity of a traditional ring VCO. These properties allow for straightforward performance improvement in a ring-VCO-based PLL. Designers may use the accurate and concise expressions that I have developed to quickly calculate the expected performance before fine tuning with expensive

simulation. These benefits, nonetheless, come at a trade-off cost of additional area for the inductors and smaller tuning range.

Certain aspects of the proposed architecture were not explored and provide an opportunity for further research. One of them is the assumption that the noise from the inverters always dominates while the noise originating in the resistance of the inductors (i.e., due to their quality factor Q) is negligible. Naturally, the phase noise due to the inductors' Q is a limiting noise floor. However, two questions arise: How close can we get to this theoretical limit, and how does the phase noise due to limited Q in the proposed star-shaped structure compare to a simple LC tank? Conversely, it is also of interest to answer the question of how physically small can the inductors be before their Q becomes the limiting factor.

Another question that may arise for the circuit designer is how to deal with the increased frequency as a result of lower inductance, or, more specifically, whether there is still a FOM improvement if the frequency is kept constant by increasing the load capacitance as the inductance is reduced.

Finally, possible variants to the proposed architecture could provide further improvement. Two of these have been considered but not yet studied: the use of more than just three phases and the use of coupling between the inductors in different phases. These changes may affect the FOM as well as tuning range.

CHAPTER VII

DESIGN OF A TIME-TO-DIGITAL CONVERTER WITH SAMPLE-AND-HOLD AND QUANTIZATION NOISE SCRAMBLING USING HARMONICS IN RING OSCILLATORS

7.1 *Introduction*

7.1.1 Bandwidth in Noise-Shaping TDCs

Oversampling and noise shaping by $\Delta\Sigma$ modulation, as I described in Section 4.3.3, provided TDCs with higher resolution and higher linearity than ever before. Unfortunately, this is achieved at the expense of bandwidth. This is problematic in DPLLs since high bandwidth is essential for suppressing VCO noise. Furthermore, there are cases in which oversampling is not possible at all, for example, when a single event, i.e., a single sample, must be measured with high resolution.

Sample-and-hold is a common technique used in voltage-to-digital converters. It is based on storing a sample of the input quantity that can be reused. This reuse can enable oversampling without the need to take more samples from the input. However, this has not been achieved thus far in TDCs.

I propose a TDC technique that achieves sample-and-hold for the first time, employing harmonics in ring oscillators.

7.1.2 Harmonics in Ring Oscillators

The fundamental frequency of a ring oscillator with N inverting stages, each with a propagation delay τ_{gate} , is $f_{\text{ring}} = (2N\tau_{\text{gate}})^{-1}$. It is possible, however, to make it oscillate at an integer multiple of f_{ring} by applying specific initial conditions.

Consider the ideal ring oscillator circuit in Figure 49 consisting of six inverters. Two of these can be held in a reset state which forces their outputs to logic *high*. When *reset* is released, the ring is placed in an unstable state, causing logic transitions (edges) to propagate. This ring oscillator is operating in its second harmonic.

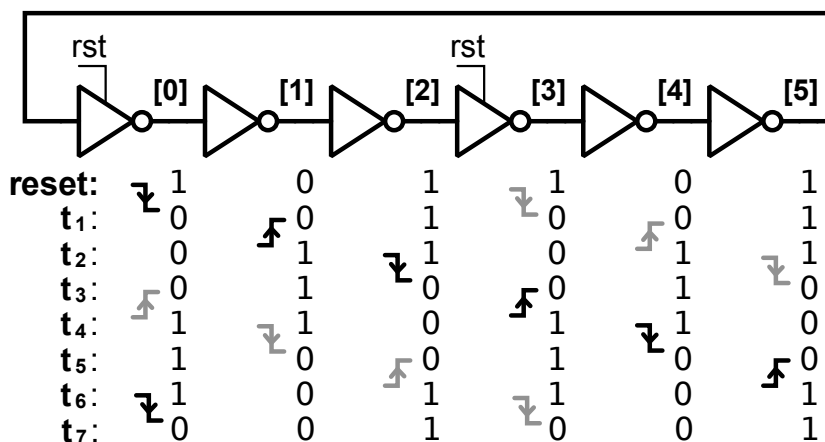


Figure 49: Ring oscillator operating in second harmonic mode.

Harmonics in ring oscillators have seldom been studied, and it is usually with the intention of avoiding them [2][5][11][42]. The exceptions that I have found report the use of second-harmonic ring oscillators for a carrier-and-data-recovery (CDR) circuit [19], and for a physically unclonable function (PUF) for chip authentication [50].

7.2 Proposed Architecture

Since the proposed TDC is primarily based on the use of harmonics in ring oscillators, I refer to it as the Harmonic Ring Oscillator (HRO) TDC.

7.2.1 Sample-and-Hold in the Phase Domain

The proposed TDC exploits the possibility of releasing the reset condition of the inverters in Figure 49, by two independent input signals A and B, at different times. Releasing the two reset signals at different times changes the relative phase of the two edges propagating in the ring.

Figure 50 illustrates the relationship between the arrival times of A and B, their absolute and relative phases in the ring, and the ring’s fundamental frequency. The relative phase of the two propagating edges ideally remains constant since both accumulate at the constant rate of f_{ring} . Therefore, this circuit behaves as a differential sample-and-hold in the phase domain.

7.2.2 Number of Harmonics

Unfortunately, in a real circuit, there will always be random mismatch. This will cause the two edges to propagate at different frequencies in an even-stage ring. Note how in Figure 49, at any given node, if one of the edges is a rising edge, then it will always be a rising edge when going through that same node again. The other edge will always be a falling edge when passing through that node (Compare times t_1 and t_7 in Figure 49). We can write the period of both edges around the ring as

$$T_{\text{edge1}} = \sum_{i=0}^{N/2-1} \tau_f[2i] + \tau_r[2i + 1], \quad (44)$$

$$T_{\text{edge2}} = \sum_{i=0}^{N/2-1} \tau_f[2i + 1] + \tau_r[2i], \quad (45)$$

where $\tau_r[i]$ and $\tau_f[i]$ are the propagation delays of each gate when its output at node i is respectively a rising or a falling edge. The rise and fall times of a gate can be designed to match, but in practice, they will never be equal because of random mismatch. Therefore, in terms of frequency of oscillation, it is as if the edges were propagating on different rings. The angle between both edges quickly grows or shrinks. Eventually, the two edges collapse, and oscillation ceases. This places severe limitations for the implementation of the proposed TDC using an even-stage ring oscillator.

In an odd-stage ring oscillator carrying any number of edges, if an edge is a rising edge at a given stage at a given time, then the next time around (after propagating through an odd number of stages) it will become a falling edge. Therefore, each edge is affected by both the rise and fall times of every gate in the ring, and their periods (average after two cycles) are equal and given by

$$T_{\text{odd}} = \frac{1}{2} \sum_{i=0}^{N-1} \tau_r[i] + \tau_f[i]. \quad (46)$$

Since an odd-stage ring oscillator overcomes the aforementioned problem in even-stage rings, it provides a better platform for the analog storage mechanism of the TDC.




Figure 50: Differential sample-and-hold in the phase domain.

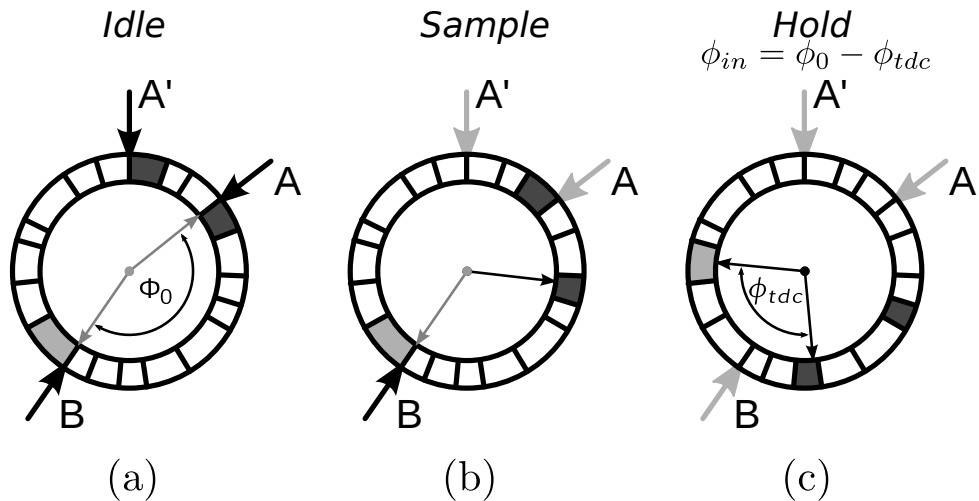


Figure 51: Conceptual representation of the HRO TDC during its three measurement states: (a) idle, (b) sample and (c) hold.

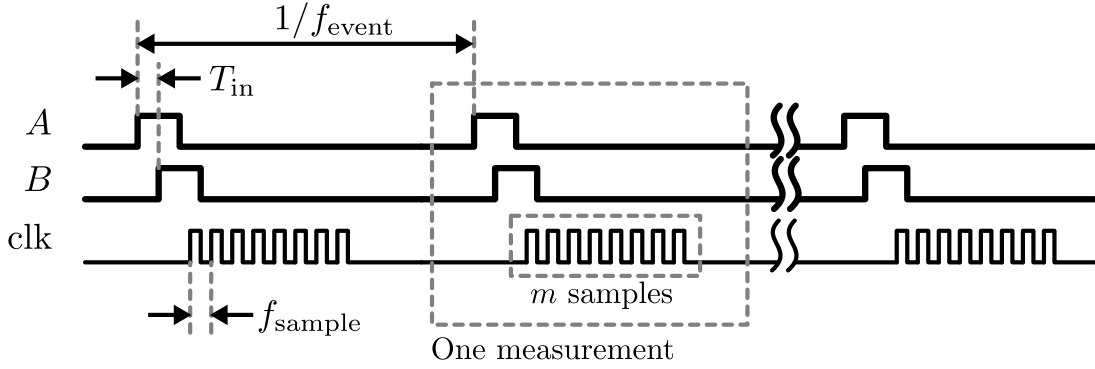


Figure 52: Measurement process and definitions for the HRO TDC.

Only two edges are required to encode the time information. However, an odd number of edges must be present in an odd-stage ring. A third edge can aid in identifying which edge, in any given sample of the ring, came from which input signal. As illustrated in Figure 51, A is injected into the ring at two locations simultaneously, close to each other, k gates apart. Signal B is injected at the maximum possible distance from the two A inputs. Then, it will be known at any time, that the two edges that are closer together correspond to signal A , and that the third one corresponds to signal B .

The above is true as long as the B edge maintains a distance greater than k gates from any of the two A edges. At any time, if the A edges are at gates i and $i+k$, then the B edge may not occupy any gates between $i-k$ and $i+2k$, which is a restriction totaling $3k+1$ gates. This restriction determines the maximum measurement range of $(N-3k-1) \times \tau_{\text{gate}}$.

7.2.3 Measurement Process

Figure 51 shows a conceptual third-harmonic HRO. The delay elements are drawn intentionally with unequal sizes in the disc to represent the unequal propagation delays. It shows the ring in three different states: (a) Idle: Neither A nor B have arrived, (b) Sampling: A has arrived, and the two corresponding edges start to propagate, and (c) Hold: B has arrived, and now the angle $\phi_{\text{tdc}} = \phi_0 - \phi_{\text{in}}$ stores the desired quantity.

The objective at this point is to recover ϕ_{tdc} with a resolution better than the average gate delay τ_{gate} and to scramble the nonlinearity that arises from the variation of individual

gate delays. For this purpose, m samples of ϕ_{tdc} are captured during the *hold* state. The complete sequence is shown in Figure 52.

At this point, I must emphasize the distinction between a *measurement* and a *sample* since this is what sets the HRO TDC apart from other TDCs. One *measurement* is the process of taking one *input-sample* and storing it in the ring, followed taking m *samples* of the constant quantity ϕ_{tdc} , which are then averaged to generate the measurement result. For other TDCs there is no distinction between *measurement* and *sample* since without a sample-and-hold mechanism, only one sample can be captured in every measurement.

7.2.4 Noise Scrambling

As long as the ring oscillation and the ring sampling remain asynchronous (which is highly likely unless the free-running ring becomes injection-locked to the sampling clock), every time the ring is sampled, each edge will be located at a different delay stage. This means each sample is subject to the mismatch of a different set of delay elements. The nonlinearity introduced by the mismatch is therefore spread or scrambled throughout multiple samples, resembling a zero-mean noise-like signal. This is known as Dynamic Element Matching (DEM).

Furthermore, the position of each phase relative to the quantization thresholds at the time of sampling will be different for each sample. Thus, the quantization noise becomes scrambled as well.

The scrambling of nonlinearities and quantization errors is essential for increasing accuracy by oversampling as I described in Section 4.2.

7.3 Noise Analysis

The diagram in Figure 53 illustrates the signal flow for different types of noise and their sources. It also represents the common-mode and differential effects by separating the signal path of the two phases/edges that jointly store the time sample. The sources can be grouped into quantization and device/circuit noise.




Figure 53: Model of the HRO TDC for noise analysis.

7.3.1 Quantization Noise

In Section 4.2 we saw how quantization noise, if uniformly distributed, has an RMS value, normalized to the resolution τ , of $\sigma_\tau = 1/\sqrt{12}$ and that the average of m quantized samples has an RMS quantization noise of $\sigma_m = \sigma_\tau/\sqrt{m}$ (Eq. 25).

In the HRO TDC, however, $T_{\text{in}}/\tau = n + v$ stored in the ring is a constant during all m samples. Therefore, v , the quantization error, is not randomized and the necessary conditions for uniformly distributed quantization noise [4][44] may not be met. However, the quantization errors are scrambled throughout samples as a random dither w is inherently introduced. As long as the ring's frequency and the sampling frequency remain asynchronous, we can consider $w \sim U(0, 1)$, and Eq. 25 still holds.

In the absence of delay mismatch, and due to the random dither, the quantized samples are a random variable given by

$$Q\{T_{\text{in}}\} = Q\{\tau \times (n + v + w)\} \quad (47)$$

$$= \begin{cases} n, & p_n = 1 - v, \\ n + 1, & p_{n+1} = v \end{cases} \quad (48)$$

where p_x is the probability of obtaining the value x .

The RMS quantization noise is, therefore, a function of v and given by $\sigma_\tau = \sqrt{v(1-v)}$. On average throughout v , $\overline{\sigma_\tau} = \pi/8 \approx 1/\sqrt{6}$. Delay mismatch in the inverters introduces further scrambling on σ_τ , reducing its dependence on v .

Further scrambling can come from phase noise in the sampling clock. Additional, even intentional, phase noise in the clock may also lower the chances for the ring to become injection locked. Contrary to most other TDC architectures, the time computation is not referenced to the circuit's clock.

The quantization noise is not only scrambled across ring samples but also throughout measurements as long as the input signals are not synchronous to the sampling clock. Quantization noise scrambling across measurements occurs because the ring samples from one measurement happen at different locations (phases) in the ring from those in other measurements.

The HRO TDC is a particular case of $\Sigma\Delta$ noise-shaping converter, carrying out zeroth-order ($L = 0$) noise shaping (Eq. 27). In this case, the quantization noise is not shifted to higher frequencies but just spread out across the entire measurement bandwidth.

It is evident from Eq. 27 that any converter with $L = 1$ or higher outperforms an $L = 0$ converter as OSR increases. However, the HRO TDC has $\sigma_q = \sigma_m = \sigma_\tau/\sqrt{m}$, while converters without sample-and-hold have $\sigma_q = \sigma_\tau$. This comparison is shown in Figure 54.

The overall impact of this property is that, for the same raw resolution or unit delay, $\tau = \tau_{\text{gate}}$, and event rate, f_{event} , the precision of the measurements taken with the HRO TDC depends on the value of m . The achievable m depends on the available time between events, $1/f_{\text{event}}$, and on the frequency of the sampling clock, f_{sample} , which we can design for. Other converters, for a given precision, have to wait for a minimum number of input events, which, more often than not, happen at a frequency outside the control of the designer.

7.3.2 Device and Circuit Noise

Figures 55(a) and 55(b) are modified versions of Figure 50 with additional information to aid in the understanding of noise in the HRO TDC. In both, the dotted lines represent




Figure 54: Comparison of quantization noise versus the oversampling ratio of $\Sigma\Delta$ noise-shaping TDCs and the HRO TDC.

the phase trajectory that the stored edges would follow for an input time interval T_{in} if the TDC were noise-free. In (a), a noise source $e_f(t)$ (Figure 53) affecting the overall $f_r(t) = f_{ring} + e_f(t)$ is introduced which deflects the ideal phase curves. At time t_B we can see that ϕ_{in} already contains an error, and remains unaltered thereafter. A noise source with these characteristics can introduce a permanent error during the sampling period but does not affect the stored quantity ϕ_{in} after time t_B .

In Figure, 55(b) a noise source only affecting the phase of the edge originating from signal B has been added. It represents fully-uncorrelated noise. Such noise, as opposed to that in Figure 55(a), can affect the stored value of ϕ_{in} after time t_B and the initial value of ϕ_{in} .

To quantitatively understand the effect of noise, we need to isolate the nature, source, and sensitivity to the different types of noise.

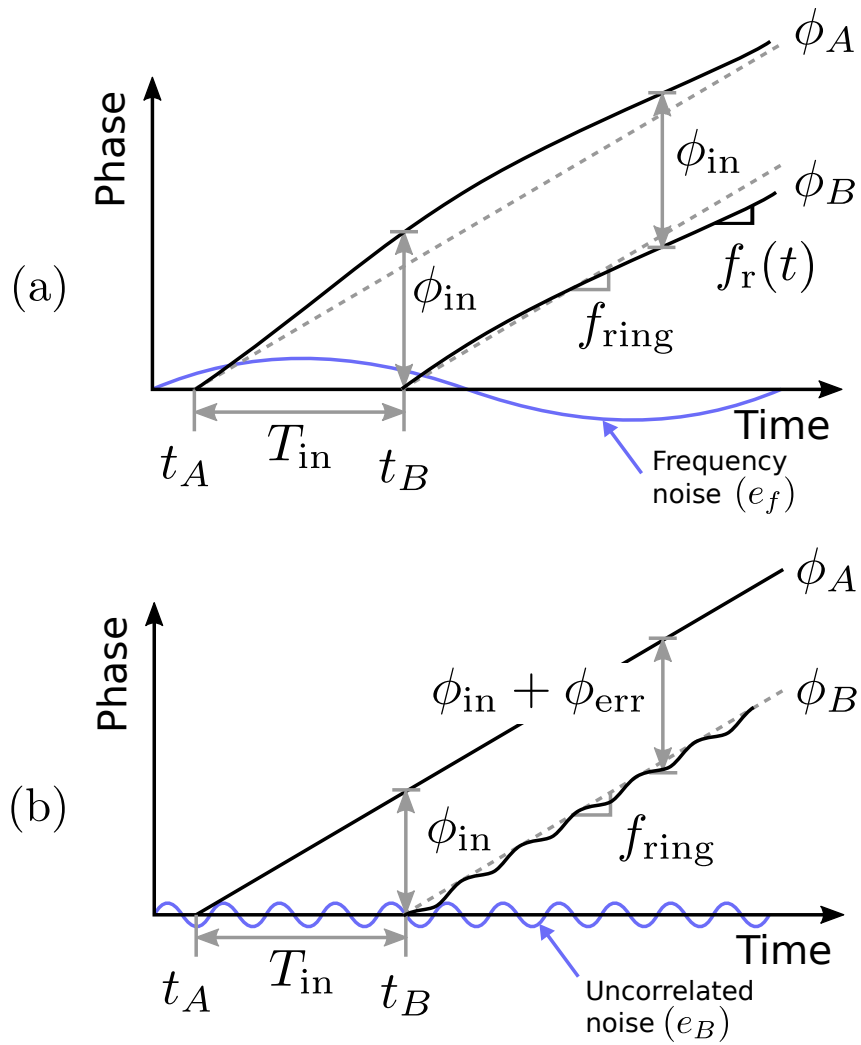


Figure 55: Impact of noise on the sample-and-hold process. (a) Fully correlated noise and (b) uncorrelated noise.

7.3.2.1 Thermal Noise

Noise originating at different inverters in a ring oscillator is uncorrelated. Also, the autocorrelation of thermal (white) noise is zero at any time offset other than zero. Therefore, the effect of thermal noise from the inverters on different edges propagating in the ring is fully uncorrelated. It corresponds to sources e_A and e_B in Figure 53 and is represented in Figure 55(b).

Since the power of the sum of uncorrelated signals is the sum of their powers, the timing uncertainty of a signal, after propagating through M gates, is then

$$\sigma_M = \sqrt{M}\sigma_{\tau_{\text{gate}}}. \quad (49)$$

This corresponds to the section of the curve with 0.5 slope in Figure 5.

The designer can use Eqs. 25 and 49 to estimate the maximum practical number of ring-samples per measurement taken at a given sampling rate by observing when the accumulated thermal noise overcomes the quantization noise.

7.3.2.2 Flicker Noise

There is no clear consensus about how Flicker noise should be modeled. We know, however, that most of its power concentrates at low frequencies and becomes apparent, i.e., dominates over thermal noise, below some frequency $f_{1/f}$. Flicker noise is still uncorrelated between inverters, but since flicker noise is “colored”, its autocorrelation $R(\Delta t)$ has non-zero values away from zero time offset, $\Delta t = 0$ [13]. The consequence is that on the same inverter, flicker noise will affect the phase similarly, in direction and magnitude, for small values of Δt .

This will occur in a general oscillator when the period is very small, i.e., the ring has few delay elements. For the HRO specifically, the effect on the phase of two propagating edges has a higher correlation when the two edges have a small phase offset, this is, when both edges cross the same inverter in a very short time. If the ring has many delay stages, and the two edges storing the time quantity are far apart, then the impact of flicker noise becomes less correlated.




Figure 56: Simulation setup to compute the ISF at the supply voltage of a ring oscillator.

We can conclude that flicker noise is partially correlated and will affect ϕ_{tdc} depending on the value of ϕ_{tdc} and the size of the ring. It is at its minimum when the ring has the most delay stages and $|T_{\text{in}}|$ is smallest.

7.3.2.3 Power Supply Noise

For noise originating at the power supply of the ring, we can hypothesize that phase noise will be fully correlated throughout all the inverters, therefore affecting just the frequency of the ring. Hence, it should not affect the relative angle of multiple simultaneously propagating edges. However, I shall show this is only partially true, using simulation and an impulse-sensitivity-function (ISF) [21] approach.

To construct the ISF, $\Gamma(\phi)$, of a ring oscillator at the power supply, we can use the simulation setup in Figure 56. A narrow current pulse is injected into C_{ring} (approximating an instantaneous charge) at uniform time offsets covering one oscillation period. For each offset, the phase deviation (from nominal) is measured several periods later. The graph in Figure 57(a) shows the observed time deviation for a current pulse injected at the specified time (or phase/state) of the oscillator (Figure 57(b)). This curve is the ISF¹. The exact values used in this simplified simulation setup are not important and in practical designs will vary. What is most relevant is that the ISF has a large DC component plus a small

¹The ISF, as defined in [21], is dimensionless.




Figure 57: Simulation results from the setup in Figure 56: (a) phase error as a function of the angle at which the charge was injected and (b) voltage at each node in the ring.

variation across one period: $\Gamma(\phi) = \Gamma_{DC} + \gamma(\phi)$.

In a ring carrying edges A and B , each has its own ISF. They are offset versions of one another, i.e., $\Gamma_B(\phi) = \Gamma_A(\phi + \phi_{tdc})$. With this information, we can construct a differential ISF for ϕ_{tdc} (the stored angle):

$$\begin{aligned}
 \Gamma_{AB}(\phi, \phi_{tdc}) &= \Gamma_B(\phi) - \Gamma_A(\phi) \\
 &= \Gamma_A(\phi + \phi_{tdc}) - \Gamma_A(\phi) \\
 &= \gamma_A(\phi + \phi_{tdc}) - \gamma_A(\phi).
 \end{aligned} \tag{50}$$

This function represents how much the relative phase between two simultaneously propagating edges changes in response to a charge injected into C_{ring} at a specific time (or phase/state) of the ring.

Since Γ_{DC} vanishes in Γ_{AB} , we can define a common-mode noise rejection ratio (CMRR) or power supply rejection ratio (PSRR) for ϕ_{tdc} as



Figure 58: Schematics of the prototype HRO TDC.

$$\text{PSRR}_{\phi_{\text{tdc}}} = \frac{\gamma_{\text{RMS}}}{\Gamma_{\text{DC}}}. \quad (51)$$

Most of the noise originating at the supply is fully correlated and common-mode, quantified by Γ_{DC} , and represented by e_f in Figure 53. However, some sensitivity is uncorrelated between edges and is also dependent on ϕ_{tdc} . It is quantified by $\gamma(\phi)$ and represented by e_A and e_B in Figure 53.

7.4 Implementation

The HRO TDC concept is general enough that it can be implemented in several different forms and optimized for different conditions of operation.

The fundamental properties of the HRO TDC were initially demonstrated with a prototype built using an FPGA [7], while the results shown here are based on a 28 nm CMOS IC prototype. It was designed in a way that would demonstrate the viability of the technique in a complex CMOS SOC. Moreover, in order to emphasize the portability of the technique across technologies, I chose to implement the HRO TDC using only standard-cell logic gates.

A simplified schematic is shown in Figure 58. The inverting stages have been chosen to be NAND gates such that external inputs could be accommodated on three of them. The rest have their second input tied to logic *high*, thus operating as inverters. This choice maintains uniform delay across gates.

Each node is sampled by D-type flip-flops (FF), and the samples of adjacent nodes are compared with XNOR gates. At the stages where transitions have not fully completed (where a transition is happening, i.e., where an edge is located), the values at the input and the output of those stages are registered as the same, therefore yielding a logic *high* at the XNOR's output. A logic *low* is observed everywhere else.

After sampling, the objective is to determine the number of stages between the edges corresponding to inputs *A* and *B*. This number is the quantized version of ϕ_{tdc} . Even though there are many ways to implement this logic function, our implementation focused on maximizing the ring-sampling frequency. This requires the logic complexity per clock cycle to be minimized. Therefore, the decoder was implemented with a pipeline structure, i.e., where a complex computation is broken down into simpler computations throughout multiple clock cycles. This approach, although resulting in higher latency, allows sampling on every clock cycle.

Figure 59 shows a photograph of the prototype chip and the layout of the ring oscillator. The signal input stages, 0, 5, and 34, have been labeled. These same stages can be identified in the schematics in Figure 58. Special attention was given to the layout of the clock signals to ensure equal propagation delay from the clock buffer (right edge in the layout) to all FFs. The decoder is not shown in the figure since it was automatically synthesized, placed and routed along with other logic unrelated to the TDC.




Figure 59: HRO TDC chip photograph (a), layout and dimensions (b).

7.5 Measurement Techniques and Results

7.5.1 Input Signals

Generating input signals with accurate time delay is typically as complicated as measuring such a delay with comparable accuracy.

For low-resolution applications, in the order of tens of picoseconds, an arbitrary waveform generator (AWG) is capable of generating the A and B signals and sweeping the time delay between the two. In [48] I used a direct-digital synthesizer (DDS) to generate two sine waves with a relative delay resolution close to 1 ps, although with a relative RMS jitter of about 10 ps.

Smaller time delays with lower noise can be achieved by using a voltage-controlled delay line [15], which is usually implemented by a cascade of inverter gates whose propagation delay is controlled by their supply voltage.

In this work, I used a different technique that allows for the adjustment of the output phase in a fractional-N PLL with close to perfect linearity and resolution only limited by the number of fractional bits. The basic principle behind it is that the $\Sigma\Delta$ modulator in the feedback divider acts as a phase accumulator. Adding an offset to the accumulator results in a phase offset in the PLL's output. To add a phase offset, we only need an

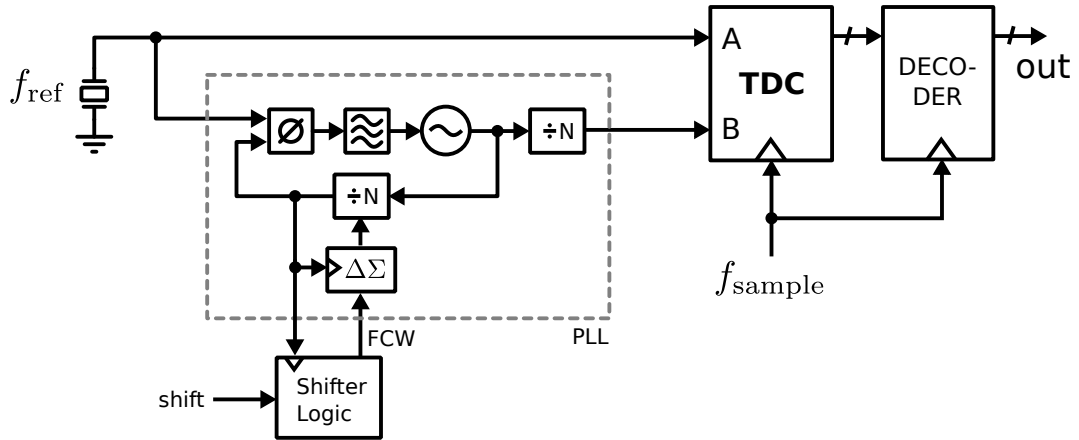


Figure 60: HRO TDC test setup. Phase shifting signal B using a fractional-N $\Delta\Sigma$ feedback divider PLL.

external circuit that switches the feedback-divider value N (fractional) from N to $N + \Delta N$ in one reference cycle and back to N on the next. The change in N adds a fixed phase offset because the amount of phase accumulated during the reference period is the result of integrating a slightly different frequency, $f + \Delta f$. The details of the technique can be found in [20].

7.5.2 Test Setup

The test setup shown in Figure 60 was fully integrated on-chip. One of the input signals to the TDC is from a reference crystal, and the other is shifted in phase by a fractional-N PLL.

7.5.3 Characterization

The first step in the characterization is to determine the average τ_{gate} . The ring oscillator can be made to oscillate in its first harmonic by supplying a rising edge on A , and then on B past the measurement range of the TDC. Edges A and A' will vanish when they reach the input stage of B . Then the B edge will be the sole edge propagating in the ring. Direct measurement of the frequency at any node yields then $f_{\text{ring}} = (2N\tau_{\text{gate}})^{-1}$. On a typical chip at room temperature, $f_{\text{ring}} = 143$ MHz was observed, therefore $\tau_{\text{gate}} = 55.5$ ps.

In normal operation, the delay between inputs A and B can be swept a whole period of




Figure 61: Measured TDC response for an input of two 100 MHz signals swept in relative phase for 360 degrees in steps of 12.5 ps.

f_{ref} . In this case, f_{ref} was 100 MHz and the 10 ns period was swept in steps of 12.5 ps. Figure 61 shows the response of the TDC using $m = 8$ averaged samples per measurement, taken at 100 Msamples/s. The range was confirmed to be 2.25 ns, slightly below the theoretical maximum of 2.6 ns. Outside the range of operation, the ring carries a single edge. Therefore, the decoded values do not have any meaning and depend on the implementation of the decoder.

The combination of precision and linearity determine measurement accuracy. Precision and linearity can be computed by collecting a large number of measurements for each input value and determining first-order statistics. The average measurements in the linear portion of the response in Figure 61 subtracted from a linear fit are shown in Figure 62. This is known as integral nonlinearity.

The standard deviation of the measurements is the RMS precision or error, σ_{tdc} . Figure 63 plots σ_{tdc} versus the number of averaged ring-samples, m . Each marker, computed from 200 measurements, is for a different input value, covering the whole input range. The $1/\sqrt{m}$ profile confirms the quantization-noise scrambling as predicted.

The measurements shown in Figure 63 have been carried out at a sampling rate proportional to the number of ring-samples, such that the total time the quantity is held in the ring is constant regardless of the number of samples per measurement. This ensures that

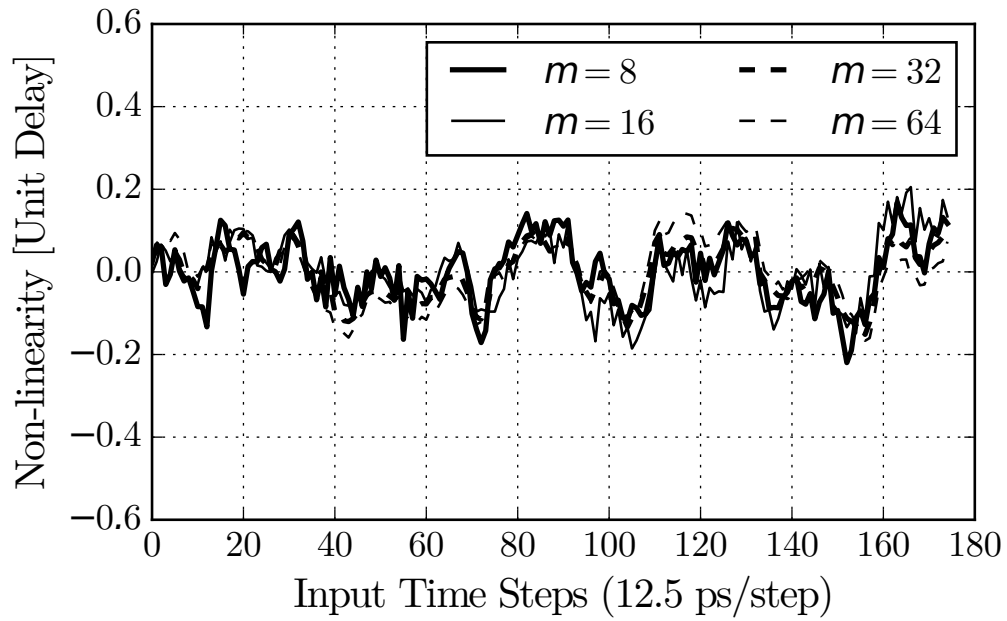


Figure 62: Measured integral nonlinearity of the HRO TDC.

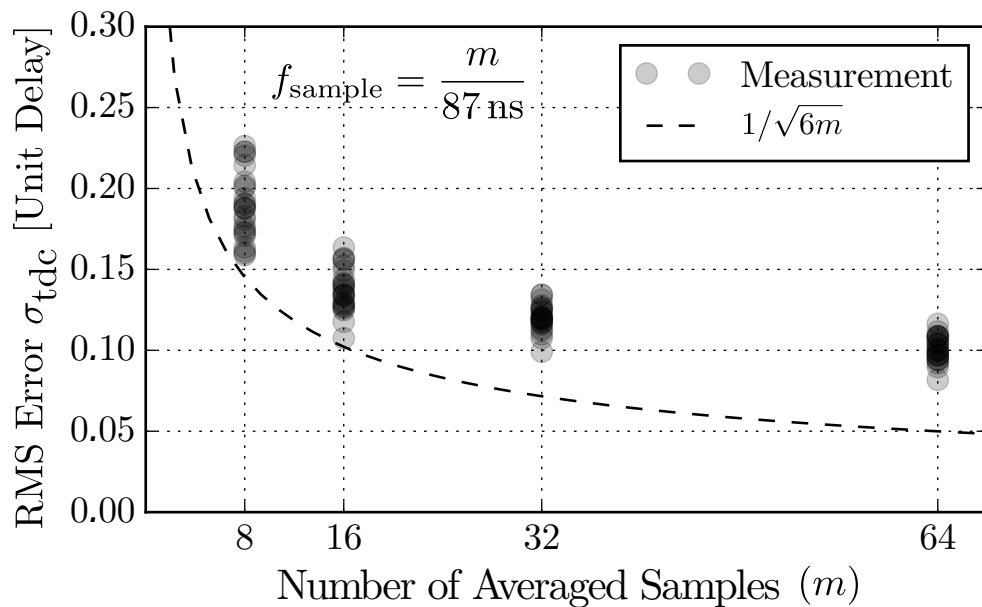


Figure 63: Measured precision of the HRO TDC versus the number of samples per measurement for constant measurement time.

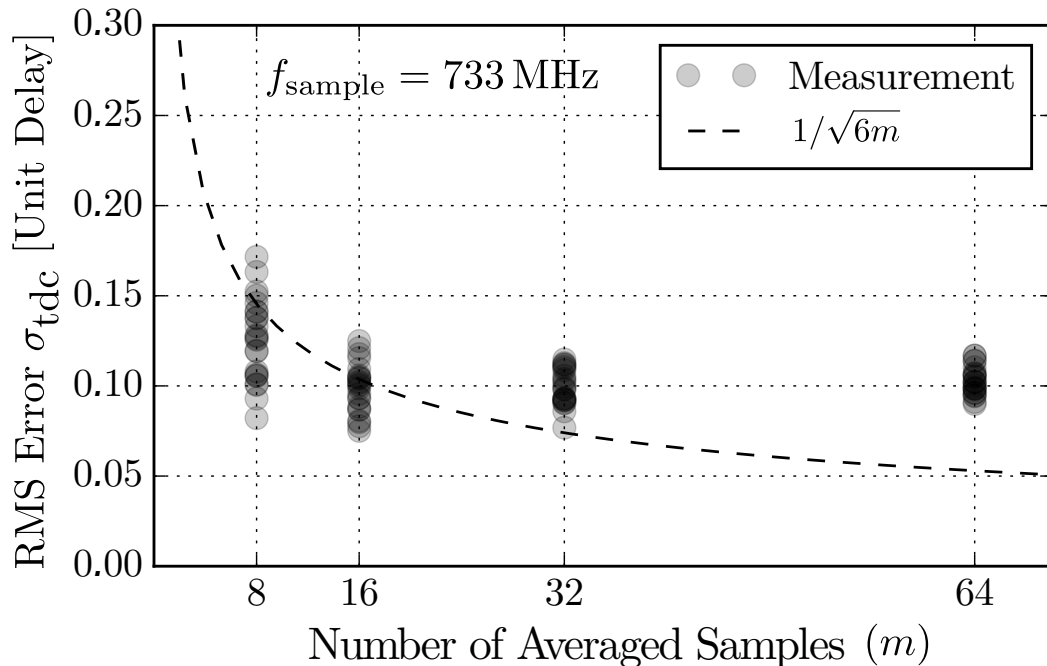


Figure 64: Measured precision of the HRO TDC versus the number of samples per measurement for constant sampling frequency.

the device and circuit noise accumulated during the measurement process stays constant. In contrast, those in Figure 64 have all been collected at 733 MHz. Degradation is observed for a large number of samples due to the increased time for noise to accumulate.

In Figure 64 the noise flattens out (around $0.1\tau_{\text{gate}}$ after $m \approx 16$) as m increases. Therefore, the additional noise (above quantization noise) must accumulate in proportion to \sqrt{m} . In other words, it must be uncorrelated (i.e., thermal).

7.5.4 Proposed Figure of Merit

Considering that power and quantization noise impact the performance of a DPLL directly and that a bandwidth-versus-quantization-noise trade-off can be leveraged, I propose the following figure-of-merit (FOM):

$$\text{FOM} = 10 \log_{10} \left(\frac{P \times \sigma_{\text{tdc}}}{\text{BW}_{\text{norm}}} \right), \quad (52)$$

where $\text{BW}_{\text{norm}} = \text{BW}/f_{\text{event}}$ is the bandwidth normalized to the event or measurement rate.

Table 2: Performance summary and comparison between recent ring-oscillator-based noise-shaping TDCs and the proposed TDC.

	This Work	[48]	[15]	[51]
Technique	HRO	GRO	SRO	GSRO
τ_{gate} [ps]	55.5	6	(1)	(1)
σ_{tdc} [ps]	5.5	0.08	0.315	0.148
f_{event} [MHz]	50	50	80	200
f_{sample} [MHz]	733	50	500	400
BW [MHz]	25	1.0	1.0	4.0
BW_{norm}	0.5	0.02	0.0125	0.02
Range [ns]	2.25	20	32	4.0
Power [mW]	3.2	21	1.5	4.0
Area [mm ²]	0.02	0.04	0.02	0.05
Noise Shaping	0th-order	1st-order	1st-order	2nd-order
FOM ⁽²⁾	-74.5	-70.8	-74.2	-73.0

(1) Not reported. (2) Eq. 52.

A summary and comparison with state-of-the-art ring-oscillator-based TDCs are presented in Table 2.

7.6 Conclusions and Proposed Research

Time-to-digital converters face the particular challenge of measuring discrete events instead of continuous signals. This limits when and how often they can carry out a measurement. Since oversampling determines precision, the inability to control the measurement (or event) rate sets a bandwidth-precision trade-off. By introducing sample-and-hold, this trade-off is broken since precision is no longer dependent on event rate.

The HRO TDC provides a reliable analog storage mechanism suitable for a time-sample-and-hold in deep submicron CMOS, where charge storage is becoming less dependable due to leakage. It also takes advantage of properties of ring-oscillator-based TDCs that have recently been identified in similar converters, such as intrinsic dynamic element matching and quantization noise scrambling. Finally, it requires no assumption about the periodicity or duty cycle of the input signals, and its accuracy is independent of reference clock phase

noise.

Opportunities for further research on the proposed architecture relate to the decoding mechanism, which was not thoroughly studied. One such possibility is the implementation of a higher order noise-shaping version of the HRO TDC. This appears possible, since each propagating edge can be treated independently, and their absolute phase can be tracked. Tracking the absolute phase of each edge would preserve the quantization error between samples yielding first-order noise shaping simply by a different implementation of the decoding logic. Furthermore, a large proportion of the power consumed by the HRO TDC was consumed by the decoder. Investigating better decoding algorithms that reduce power consumption as well as quantization error will directly impact most TDC FOMs.

CHAPTER VIII

CONCLUSIONS, CONTRIBUTIONS AND FUTURE RESEARCH

8.1 *Conclusions*

I have proposed throughout this thesis that the evolution of CMOS has motivated and shaped our research as well as the work of many other authors. In particular, the changes that are most influential to the circuit designer are those related to the degradation of analog performance as well as to the lower power consumption and faster logic gates. This has guided certain trends. One such trend is the strong interest in DPLLs during the past decade. In 2004, Staszewski *et al.* reported the first commercial DPLL [47], and by 2016, the ISSCC conference had a dedicated section for DPLLs. Another trend is the recent burst in publications reporting the use of time-based circuits for traditional analog techniques ([23][27][30][41]). These trends suggest that analog signal processing in the time domain may eventually perform better than in the voltage or current domain in modern CMOS.

In the field of frequency synthesis, the transition from the CPPLL to the DPLL is a clear example of adaptation to CMOS trends. In the past, slower, higher-power digital logic hindered the implementation of low-latency feedback loops of PLLs. Furthermore, large gate delays translated into large quantization noise when digitally measuring the phase error and when establishing the frequency of a DCO. Small gate delays in today's CMOS, however, are comparable to the jitter of a PLL's output. Low power logic also allows us to implement complex noise-shaping modulators that reduce the effective quantization noise to a small fraction of a gate delay.

From a circuit design perspective, the techniques that I have proposed in this thesis follow a similar approach; to overcome the analog design limitations and embrace the digital and time-domain improvements. From a scientific perspective, however, it is insufficient to just adapt to the trends. I have strived to contribute techniques that are independent of a specific implementation or choice of platform, and therefore to advance the knowledge in

the field in a fundamental way.

The proposed hybrid digital-analog PLL architecture makes use of the benefits afforded by DPLLs to overcome a fundamental limitation of traditional PLL architectures. It allows for a low-bandwidth low-pass response to input phase noise, and a high-bandwidth high-pass response to the VCO's phase noise. The technique involves fairly low complexity and provides ease of integration in SoCs since the digital portion can be automatically synthesized, placed and routed, and requires minimal effort to integrate with a standard fractional-N CPPLL. Even though this design calls for the use of a CPPLL, there are no specific requirements for it other than a high-resolution fractional feedback divider. Since existing designs of already optimized general purpose CPPLLs can be reused, the benefits of the proposed circuit may be afforded without any analog design at all. Overall, the proposed solution facilitates reuse and portability across processes, while synthesizable logic minimizes risk and reduces implementation time with respect to fully analog solutions.

The research towards the inductor-loaded ring VCO addressed a fundamental source of timing inaccuracy: the phase noise in oscillators. Specifically, I explored the difficulty of achieving trade-offs between two polarized architectures: the convenient, flexible and reliable ring oscillator and the low phase noise LC oscillator. A simple modification in a single-ended CMOS ring VCO, consisting of just a simple attached load, allows for a predictable reduction in phase noise. The trade-offs are well understood: The improvement in phase noise results in a smaller tuning range, and the use of inductors requires large silicon area. However, since low Q is not a limiting factor, the inductors can occupy a minimal area by using multiple turns and multiple metal layers. This can lead to inductors that are an order of magnitude smaller than those in LC oscillators.

The improvement in phase noise by the use of this approach is well-characterized through simple mathematical expressions and based on only a few circuit parameters. The trade-offs are easily quantified as well. Overall, the concise toolset that I provide allows designers to use this technique with confidence, by taking incremental steps, rather than switching oscillator architectures to achieve system specifications.

The approach towards the HRO TDC considered a multitude of factors: the imperfections of CMOS that result in nonlinearity, the fundamental speed limitation of logic gates leading to limited resolution, noise sensitivity of single-ended techniques, and the lack of alternatives to oversampling and $\Delta\Sigma$ noise shaping, which limit the bandwidth in favor of resolution. Furthermore, it targeted the specific use in phase detectors for DPLLs which are typically more demanding in terms of resolution, linearity and bandwidth than in other applications.

The outcome of these challenges were first the positive use of harmonics, or multiple propagating edges, in ring oscillators. Harmonics have been observed in ring oscillators before, but have been considered an undesirable behavior and, so far, the objective has been to suppress them. For perhaps the first time being reported, harmonics in ring oscillators have been understood well enough to use them as a feature and achieve improvement in circuit performance. Specifically, this has allowed for a sample-and-hold mechanism for TDCs, which is differential, i.e., immune to certain forms of common-mode noise, and allows it to achieve higher resolution by oversampling without a bandwidth penalty.

The research presented in this thesis has been carried out keeping in mind certain aspects of circuit design and engineering in general that are sometimes undervalued or misunderstood. These are simplicity, risk and convenience. The three are intimately tied together and are hardly ever quantified in figures of merit. However, in complex SoCs, there is a significant chance that individual failures will affect the overall SoC functionality. For this reason, we want the individual components to be predictable under changing circumstances, like PVT and mismatch variations, and to be characterizable with as few parameters as possible. These concerns often carry more significant consequences than traditional FOM metrics.

This perspective led us to develop a PLL technique that combines highly predictable digital logic with a highly mature technique like the CPPLL, while avoiding any analog customization. It also prompted us to find an improvement to ring VCOs that can be quantified by simple equations and achieved by adding only three identical passive components to a well-known (and also the most common) VCO in PLLs. The same applies to the HRO

TDC. I achieved this design with only standard digital gates, while reducing the dependence of the accuracy on manufacturing variations and interference from other circuits that may be present in a SoC.

8.2 Contributions

The contributions in this research to the field of electronic circuits for timing generation and measurement have materialized through journal publications, conference publications and conference presentations. These include:

- A time to digital converter that uses harmonics in ring oscillators to achieve a sample-and-hold mechanism that is differential and highly insensitive to noise. It builds on previous techniques like oversampling and dynamic element matching. It is also built entirely with digital gates, making it suitable for multiple platforms including nanometer CMOS and FPGAs. It has been proven on both platforms to achieve sub-gate-delay resolution and linearity [7][8].
- A technique that combines a digital PLL loop with an analog CPPLL to overcome the limitation of simple PLL loops of being able to either suppress phase noise originating at the VCO or the reference. It decouples the two filtering mechanisms allowing for a low-bandwidth low-pass response to reference noise and high-bandwidth high-pass response to VCO noise [9].
- A VCO architecture that eases trade-offs in ring oscillators allowing for lower phase noise in exchange for tuning bandwidth and area. The improvements are accurately quantified with simple mathematical expressions involving only a few circuit parameters, and it allows for a low-risk path to phase noise improvement in ring-oscillator based PLLs [6].

8.3 Future Research

At the end of Chapters 5, 6 and 7, I discussed possible further research for each topic. However, it is essential to consider how these contributions enhance the performance of larger

systems. Moreover, to consider if larger systems can benefit synergistically by integrating the proposed techniques together.

One possible integration is a DPLL system. The proposed HRO TDC may be used as the phase detector directly. This TDC is easy to integrate since it already has two inputs to accommodate the reference and feedback signals. The high-frequency VCO output may be used to oversample the held quantity much faster than the reference rate. The inductor-loaded ring VCO may be the core of the DCO. The proposed ring VCO can be used in conjunction with a current DAC to implement a DCO. Current DACs can be designed to be very fast and very linear [39]. This may be a significant advantage when compared to an LC-based DCO, which is tuned by varactors or switched-capacitor banks. Varactors need to be driven by a voltage DAC, which is generally slower and less linear than a current DAC. Switching capacitors, on the other hand, may lead to high nonlinearity due to mismatch and noisy switching transients.

PLLs are, however, just one out of many applications that may benefit from better TDCs and ring VCOs.

REFERENCES

- [1] ABIDI, A. A., “Phase noise and jitter in CMOS ring oscillators,” *IEEE Journal of Solid-State Circuits*, vol. 41, pp. 1803–1816, Aug 2006.
- [2] AGUSTIN, J., LOPEZ-VALLEJO, M. L., SORIANO, C. G., CHOLBI, P., MASSENGILL, L. W., and CHEN, Y. P., “Efficient mitigation of set induced harmonic errors in ring oscillators,” *IEEE Transactions on Nuclear Science*, vol. 62, pp. 3049–3056, Dec 2015.
- [3] BARON, R. G., “The Vernier time-measuring technique,” *Proceedings of the IRE*, vol. 45, pp. 21–30, Jan 1957.
- [4] BENNETT, W. R., “Spectra of quantized signals,” *The Bell System Technical Journal*, vol. 27, pp. 446–472, July 1948.
- [5] BHUSHAN, M. and KETCHEN, M. B., “Generation, elimination and utilization of harmonics in ring oscillators,” in *2010 International Conference on Microelectronic Test Structures (ICMTS)*, pp. 108–113, March 2010.
- [6] CARAM, J. P., GALLOWAY, J., and KENNEY, J. S., “Voltage-controlled ring oscillator with FOM improvement by inductive loading,” *IEEE Microwave and Wireless Component Letters*, Nov 2008 (Submitted).
- [7] CARAM, J. P., GALLOWAY, J., and KENNEY, J. S., “Harmonic ring oscillator time-to-digital converter,” in *Circuits and Systems (ISCAS), 2015 IEEE International Symposium on*, pp. 161–164, May 2015.
- [8] CARAM, J. P., GALLOWAY, J., and KENNEY, J. S., “Time-to-digital converter with sample-and-hold and quantization noise scrambling using harmonics in ring oscillators,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 65, pp. 74–83, Jan 2018.
- [9] CARAM, J. P., KENNEY, J. S., and GALLOWAY, J., “A phase locked loop used as a digitally-controlled oscillator for flexible frequency synthesis,” in *Wireless and Microwave Technology Conference (WAMICON), 2015 IEEE 16th Annual*, pp. 1–3, April 2015.
- [10] CHANG, W. S., HUANG, P. C., and LEE, T. C., “A fractional-n divider-less phase-locked loop with a subsampling phase detector,” *IEEE Journal of Solid-State Circuits*, vol. 49, pp. 2964–2975, Dec 2014.
- [11] CHEN, Y. P., LOVELESS, T. D., MAILLARD, P., GASPARD, N. J., JAGANNATHAN, S., STERNBERG, A. L., ZHANG, E. X., WITULSKI, A. F., BHUVA, B. L., HOLMAN, T. W., and MASSENGILL, L. W., “Single-event transient induced harmonic errors in digitally controlled ring oscillators,” *IEEE Transactions on Nuclear Science*, vol. 61, pp. 3163–3170, Dec 2014.
- [12] DA DALT, N., “A design-oriented study of the nonlinear dynamics of digital bang-bang PLLs,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 52, pp. 21–31, Jan 2005.

- [13] DEMIR, A., “Phase noise and timing jitter in oscillators with colored-noise sources,” *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 49, pp. 1782–1791, Dec 2002.
- [14] ELSHAZLY, A., INTI, R., YOUNG, B., and HANUMOLU, P. K., “Clock multiplication techniques using digital multiplying delay-locked loops,” *IEEE Journal of Solid-State Circuits*, vol. 48, pp. 1416–1428, June 2013.
- [15] ELSHAZLY, A., RAO, S., YOUNG, B., and HANUMOLU, P. K., “A noise-shaping time-to-digital converter using switched-ring oscillators; analysis, design, and measurement techniques,” *IEEE Journal of Solid-State Circuits*, vol. 49, pp. 1184–1197, May 2014.
- [16] GANDE, M., MAGHARI, N., OH, T., and MOON, U. K., “A 71dB dynamic range third-order $\Delta\Sigma$ TDC using charge-pump,” in *2012 Symposium on VLSI Circuits (VLSIC)*, pp. 168–169, June 2012.
- [17] GAO, X., KLUMPERINK, E. A. M., BOHSALI, M., and NAUTA, B., “A low noise sub-sampling PLL in which divider noise is eliminated and PD/CP noise is not multiplied by N^2 ,” *IEEE Journal of Solid-State Circuits*, vol. 44, pp. 3253–3263, Dec 2009.
- [18] GAO, X., KLUMPERINK, E. A. M., SOCCI, G., BOHSALI, M., and NAUTA, B., “Spur reduction techniques for phase-locked loops exploiting a sub-sampling phase detector,” *IEEE Journal of Solid-State Circuits*, vol. 45, pp. 1809–1821, Sept 2010.
- [19] GIERKINK, S. L. J., “A 2.5 Gb/s run-length-tolerant burst-mode CDR based on a 1/8th-rate dual pulse ring oscillator,” *IEEE Journal of Solid-State Circuits*, vol. 43, pp. 1763–1771, Aug 2008.
- [20] GRAY, B., MASOOD, M., GALLOWAY, J., CAPLAN, R., and KENNEY, J. S., “Microdegree frequency and phase difference control using fractional-n PLL synthesizers,” in *2012 IEEE/MTT-S International Microwave Symposium Digest*, pp. 1–3, June 2012.
- [21] HAJIMIRI, A. and LEE, T. H., “A general theory of phase noise in electrical oscillators,” *IEEE Journal of Solid-State Circuits*, vol. 33, pp. 179–194, Feb 1998.
- [22] HAJIMIRI, A., LIMOTYRAKIS, S., and LEE, T. H., “Jitter and phase noise in ring oscillators,” *IEEE Journal of Solid-State Circuits*, vol. 34, pp. 790–804, Jun 1999.
- [23] HANUMOLU, P. K., “Time-based $\Sigma\Delta$ ADCs,” in *2017 IEEE Custom Integrated Circuits Conference (CICC)*, pp. 1–44, April 2017.
- [24] HELAL, B. M., STRAAYER, M. Z., WEI, G., and PERROTT, M. H., “A highly digital mdll-based clock multiplier that leverages a self-scrambling time-to-digital converter to achieve subpicosecond jitter performance,” *IEEE Journal of Solid-State Circuits*, vol. 43, pp. 855–863, April 2008.
- [25] IWATA, A., SAKIMURA, N., NAGATA, M., and MORIE, T., “The architecture of delta sigma analog-to-digital converters using a voltage-controlled oscillator as a multibit quantizer,” *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 46, pp. 941–945, July 1999.

- [26] KERANEN, P. and KOSTAMOVAARA, J., “A wide range, 4.2 ps(rms) precision CMOS TDC with cyclic interpolators based on switched-frequency ring oscillators,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 62, pp. 2795–2805, Dec 2015.
- [27] KIM, S. J., CHOI, W., PILAWA-PODGURSKI, R., and HANUMOLU, P. K., “A 10-MHz 2800-mA 0.51.5-V 90% peak efficiency time-based buck converter with seamless transition between PWM/PFM modes,” *IEEE Journal of Solid-State Circuits*, vol. 53, pp. 814–824, March 2018.
- [28] KIM, S. J., KIM, W., SONG, M., KIM, J., KIM, T., and PARK, H., “15.5 A 0.6V 1.17ps PVT-tolerant and synthesizable time-to-digital converter using stochastic phase interpolation with 16x spatial redundancy in 14nm FinFET technology,” in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, pp. 1–3, Feb 2015.
- [29] LEE, M. and ABIDI, A. A., “A 9 b, 1.25 ps resolution coarse fine time-to-digital converter in 90 nm CMOS that amplifies a time residue,” *IEEE Journal of Solid-State Circuits*, vol. 43, pp. 769–777, April 2008.
- [30] LEENE, L. B. and CONSTANDINOU, T. G., “Time domain processing techniques using ring oscillator-based filter structures,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 64, pp. 3003–3012, Dec 2017.
- [31] LIU, C. and MCNEILL, J. A., “Jitter in oscillators with 1/f noise sources,” in *Circuits and Systems, 2004. ISCAS '04. Proceedings of the 2004 International Symposium on*, vol. 1, pp. I-773–6 Vol.1, May 2004.
- [32] LIU, Y. H., HEUVEL, J. V. D., KURAMOCHI, T., BUSZE, B., MATEMAN, P., CHILLARA, V. K., WANG, B., STASZEWSKI, R. B., and PHILIPS, K., “An ultra-low power 1.7-2.7 GHz fractional-n sub-sampling digital frequency synthesizer and modulator for IoT applications in 40 nm CMOS,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 64, pp. 1094–1105, May 2017.
- [33] MAHMOODI, H., MUKHOPADHYAY, S., and ROY, K., “Estimation of delay variations due to random-dopant fluctuations in nanoscale CMOS circuits,” *IEEE Journal of Solid-State Circuits*, vol. 40, pp. 1787–1796, Sept 2005.
- [34] MCNEILL, J., “Jitter in ring oscillators,” in *Proceedings of IEEE International Symposium on Circuits and Systems - ISCAS '94*, vol. 6, pp. 201–204 vol.6, May 1994.
- [35] PERROTT, M. H., TROTT, M. D., and SODINI, C. G., “A modeling approach for sigma;- delta; fractional-n frequency synthesizers allowing straightforward noise analysis,” *IEEE Journal of Solid-State Circuits*, vol. 37, pp. 1028–1038, Aug 2002.
- [36] PORAT, D. I., “Review of sub-nanosecond time-interval measurements,” *IEEE Transactions on Nuclear Science*, vol. 20, pp. 36–51, Oct 1973.
- [37] RAZAVI, B., “A study of phase noise in CMOS oscillators,” *IEEE Journal of Solid-State Circuits*, vol. 31, pp. 331–343, Mar 1996.
- [38] RAZAVI, B., *RF Microelectronics*. Upper Saddle River, NJ, USA: Prentice Hall Press, 2nd ed., 2011.

- [39] RAZAVI, B., “The current-steering DAC [A circuit for all seasons],” *IEEE Solid-State Circuits Magazine*, vol. 10, pp. 11–15, winter 2018.
- [40] RAZAVI, B., *Designing BangBang PLLs for Clock and Data Recovery in Serial Data Transmission Systems*, pp. 736–. Wiley-IEEE Press, 2003.
- [41] SALZ, B., TALEGAONKAR, M., SHU, G., ELMALLAH, A., NANDWANA, R., SAHOO, B., and HANUMOLU, P. K., “A 0.7V time-based inductor for fully integrated low bandwidth filter applications,” in *2017 IEEE Custom Integrated Circuits Conference (CICC)*, pp. 1–4, April 2017.
- [42] SASAKI, N., “Higher harmonic generation in CMOS/SOS ring oscillators,” *IEEE Transactions on Electron Devices*, vol. 29, pp. 280–283, Feb 1982.
- [43] SCHREIRER, R. and TEMES, G. C., *Understanding Delta-Sigma Data Converters*, chapter Introduction, pp. 1–20. Wiley-IEEE Press, 2005.
- [44] SRIPAD, A. and SNYDER, D., “A necessary and sufficient condition for quantization errors to be uniform and white,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 25, pp. 442–448, Oct 1977.
- [45] STASZEWSKI, R. B. and BALSARA, P. T., “Phase-domain all-digital phase-locked loop,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 52, pp. 159–163, March 2005.
- [46] STASZEWSKI, R. B., FERNANDO, C., and BALSARA, P. T., “Event-driven simulation and modeling of phase noise of an RF oscillator,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 52, pp. 723–733, April 2005.
- [47] STASZEWSKI, R. B., MUHAMMAD, K., LEIPOLD, D., HUNG, C.-M., HO, Y.-C., WALLBERG, J. L., FERNANDO, C., MAGGIO, K., STASZEWSKI, R., JUNG, T., KOH, J., JOHN, S., DENG, I. Y., SARDA, V., MOREIRA-TAMAYO, O., MAYEGA, V., KATZ, R., FRIEDMAN, O., ELIEZER, O. E., DE OBALDIA, E., and BALSARA, P. T., “All-digital TX frequency synthesizer and discrete-time receiver for bluetooth radio in 130-nm CMOS,” *IEEE Journal of Solid-State Circuits*, vol. 39, pp. 2278–2291, Dec 2004.
- [48] STRAAYER, M. Z. and PERROTT, M. H., “A multi-path gated ring oscillator TDC with first-order noise shaping,” *IEEE Journal of Solid-State Circuits*, vol. 44, pp. 1089–1098, April 2009.
- [49] VERCESI, L., FANORI, L., BERNARDINIS, F. D., LISCIDINI, A., and CASTELLO, R., “A dither-less all digital PLL for cellular transmitters,” *IEEE Journal of Solid-State Circuits*, vol. 47, pp. 1908–1920, Aug 2012.
- [50] YANG, K., DONG, Q., BLAAUW, D., and SYLVESTER, D., “14.2 A physically unclonable function with BER $< 10^{-8}$ for robust chip authentication using oscillator collapse in 40nm CMOS,” in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, pp. 1–3, Feb 2015.
- [51] YU, W., KIM, K., and CHO, S., “A $148f_{s_{rms}}$ integrated noise 4 MHz bandwidth second-order $\Delta\Sigma$ time-to-digital converter with gated switched-ring oscillator,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 61, pp. 2281–2289, Aug 2014.