

ADAPTIVE RECEIVERS FOR HIGH-SPEED WIRELINE LINKS

by

Dustin Dunwell

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Electrical and Computer Engineering
University of Toronto

Copyright © 2012 by Dustin Dunwell

ADAPTIVE RECEIVERS FOR HIGH-SPEED WIRELINE LINKS

Dustin Dunwell

Doctor of Philosophy, 2012
Graduate Department of Electrical and Computer Engineering
University of Toronto

Abstract

This thesis examines the design of high-speed wireline receivers that can be adapted to a variety of operating conditions. In particular, the ability to adapt to varying received signal strengths, channel losses and operating frequencies is explored. In order to achieve this flexibility, this thesis examines several key components of such a receiver.

First, a 15 Gb/s preamplifier with 10-dB gain control for the input stage of an analog front end (AFE) is presented that automatically adjusts its power consumption to suit the gain and linearity requirements of the AFE for various received signal strengths. The gain of this preamplifier, along with the amount of peaking delivered by a linear equalizer in the AFE are controlled using a new adaptation technique, which adds only a small amount of overhead to the receiver. This adaptation scheme is able to sense changes in the received signal conditions and automatically adjust the equalization and gain of the AFE in order to optimize the vertical opening of the received eye.

In addition, this thesis presents the first clock multiplier with both a wide operating frequency range and the ability to transition between completely off and fully operational modes in under 10 cycles of the reference clock. This multiplier relies on the careful use of several injection-locked oscillators (ILOs) with an aggregate lock range of 55.7% of the 3.16-GHz centre frequency. The design of these ILOs was facilitated by the use of a new method for modeling the injection locking behaviour of oscillators. This model differs from existing techniques in the way that it relies on the simulated response of an oscillator to injected stimuli, instead of complex equations using quasi-physical parameters, to predict the behaviour of an ILO.

Acknowledgements

This section is to acknowledge and thank the many parties that generously assisted in the completion of this thesis.

Firstly, the interest and enthusiasm shown by Dr. Anthony Chan Carusone in his teaching has played a large part in sparking my own interest in very high speed analog circuit design. His continuing guidance and support, not to mention his exceptional problem solving skills, have been instrumental in the successful completion of my graduate studies.

Secondly, thanks is owed to the other graduate students in the electronics group at the University of Toronto who were never too busy to lend a helping hand and whose company during long hours at the lab was a comfort during the most stressful times. I sincerely hope that the relationships that I have made with my peers will last well beyond my graduation.

Finally, I would like to thank my wife and best friend for her patience and unwavering support of my studies during the better part of the past decade. She is the best mother in the world and with her by my side, I know that we can handle any challenge. I look forward to starting the next chapter of our lives together.

Contents

1	Introduction	1
1.1	Motivation	1
1.1.1	Cable TV and internet	1
1.1.2	Ethernet	2
1.1.3	Memory links	2
1.2	Thesis Overview and Major Contributions	3
1.2.1	Analog Front-End Gain	3
1.2.2	AFE Adaptation	4
1.2.3	Adaptive Clocking	5
1.2.4	ILO Modeling	7
2	Analog Front-Ends	8
2.1	Variable Gain Amplifiers	9
2.2	Equalization	13
2.2.1	Continuous Time Linear Equalizers	13
2.3	Analog Front-End Adaptation	16
2.3.1	Eye Monitor-Based Adaptation	17
2.4	Summary	20
3	Analog Front-End Adaptation	21
3.1	PDF-Based Adaptation	22

3.1.1	Measuring the Received PDF	23
3.1.2	Equalizer's Effect on the PDF	25
3.1.3	VGA's Effect on the PDF	27
3.2	Variable-Gain Preamplifier Design	29
3.2.1	Gain Variation	31
3.2.2	Output CM Control	32
3.3	Preamplifier Measurement Results	33
3.3.1	S-parameters	34
3.3.2	Time Domain	35
3.4	Adaptation Results	37
3.5	Conclusion	38
4	Injection Locked Oscillators	42
4.1	Injection Locked Oscillators	43
4.2	The Frequency Domain Model	44
4.3	The Impulse Sensitivity Function	47
4.3.1	ISF-Based ILO Modeling	48
4.3.2	ISF-Based Modeling for Strong Injection	50
4.4	Summary	52
5	Phase Transfer ILO Model	54
5.1	The Phase Transfer Characteristic	55
5.2	PTC-Based ILO Modeling	57
5.2.1	Steady-State Phase Shift	60
5.2.2	Lock Range	63
5.2.3	Lock Time	64
5.2.4	Tracking Bandwidth	67
5.2.5	Simulation Time	70

5.3	Wide Lock Range ILO Design	72
5.3.1	Injection Point Selection	74
5.3.2	Frequency Pre-Conditioning	77
5.4	Summary	79
6	A Wide Lock Range, Fast Power-On Clock Multiplier	81
6.1	MILOs with Adjacent Lock Ranges	83
6.1.1	Verifying Lock and Measuring Frequency Offset	85
6.1.2	Verifying Multiplication Factor	87
6.1.3	Power Down Unused MILOs	88
6.1.4	Creating Adjacent Lock Ranges	92
6.2	Measured Results	93
6.2.1	Free-Running Frequencies	94
6.2.2	Phase Transfer Characteristic	95
6.2.3	Lock Range	97
6.2.4	Deterministic Jitter	99
6.2.5	Random Jitter	102
6.2.6	Fast Power On/Off Transients	104
6.2.7	Power Consumption	107
6.3	Summary	111
7	Conclusion	112
7.1	Suggestions for Future Work	113
A	Summary of Contributions	115
B	Schematic Diagrams	117
	Bibliography	119

List of Tables

3.1	Component values used to implement the EQ design shown in Fig. 2.7.	26
3.2	Variable-gain broadband preamplifier performance summary and comparison.	37
5.1	Comparison of simulation time required to determine the lock range of an ILO to that required to determine the PTC, which can be used to determine lock range in addition to other ILO parameters.	72
5.2	Comparison of lock ranges calculated using the ISF model and the PTC model to those obtained directly using extensive SPICE-level simulations.	77
6.1	Comparison of measured lock range of the breakout MILO to those calculated using PSF simulations.	98
6.2	Performance comparison with other recent publications.	111
A.1	Summary of contributions made along with related locations in the thesis.	116

List of Figures

2.1	Variable gain can be implemented by using source degeneration with adjustable resistance. Resistors R_1 and R_2 can be added to increase gain control and provide a broadband impedance match.	10
2.2	Variable gain can be achieved by varying activating a variable number of unit-sized differential pairs, effectively changing the input transistor width.	10
2.3	The shunt-feedback configuration can provide a low-noise, broadband impedance match in electrical wireline applications [1].	11
2.4	Drawing bias current through R_f raises the CM level at the output and increases the voltage swing headroom for linear operation.	12
2.5	Measured losses of three coaxial cables illustrate frequency-dependent losses.	14
2.6	Channel losses can be equalized by introducing a combination of high frequency gain and low frequency gain/loss to the receiver.	14
2.7	By combining the outputs of high-pass and low-pass amplification stages and adjusting the weights of their contributions, channel losses can be equalized across all frequencies of interest.	15
2.8	Source degeneration reduces the gain of the amplifier. At high frequencies this reduction is eliminated by the short circuit created by C_s resulting in higher gain at high frequencies and equalization of the channel losses. . .	16
2.9	A typical wireline receiver with an AFE including gain and equalization control signals that should be generated automatically and intelligently. .	17

2.10	EQ adaptation can be performed by analyzing the high-frequency content of the EQ output but this does not necessarily maximize the resulting eye opening.	18
2.11	Adjusting the (a) sampling threshold, (b) sampling time or (c) both produces an indication of the eye opening that can be used to adapt the AFE settings.	18
2.12	An example of a BER-based adaptation used to optimize DFE tap settings.	19
3.1	A noiseless transmitted signal encounters noise and ISI in a frequency-dependent channel, which results in spreading of the PDF of the received signal.	22
3.2	A block diagram of the VGA and EQ adaptation performed in this work. The low-speed ADC digitizes the DC (average) output of the offset slicer. This information is used by the adaptation algorithm to minimize the spreading of the received PDF.	24
3.3	The DC output of the slicer drops as V_{th} is increased. The magnitude of the slope of this decline yields the PDF of the received data.	25
3.4	Simulated gain of the EQ using Cadence Spectre with full RC-extraction of the circuit shows up to 8 dB of high-frequency peaking.	26
3.5	Measured DC output voltage is proportional to (a) the percentage of time that the output spends at logic high. The slope of this curve is then used to create (b) the PDF of the the received signal.	27
3.6	Measurement results show that (a) increasing EQ peaking narrows the PDF and increases its peak value and that (b) increasing VGA gain moves this peak towards the target maximum threshold level, in this case set to 100 mV.	28
3.7	Flow diagram showing how the equalizer and gain settings are determined.	29

3.8	A preamplifier with gain control replaces a fixed-gain preamplifier and VGA in applications requiring high linearity and wide dynamic range. . .	30
3.9	Schematic of the preamplifier introduced in this work.	31
3.10	Simulated gain of the preamplifier as V_{gain} is varied. Simulations were performed using Spectre with RC-extraction of the entire preamplifier. .	32
3.11	Die photo of the receiver fabricated in 65-nm CMOS.	34
3.12	Measured S_{11} shows broadband input matching for various preamplifier gain settings.	35
3.13	(a) 4-PAM, 146 mV _{pp} , 1.1 GS/s test signal applied to the receiver input and (b) corresponding receiver output shows little distortion of the 4-PAM eyes.	36
3.14	Receiver output for PRBS inputs at (a) 1 Gb/s with an 8 mV _{pp} eye opening and (b) 15 Gb/s with a 50 mV eye opening demonstrate receiver sensitivity and speed.	37
3.15	Schematic-level simulations of the adaptation scheme in Spectre compare equalizer output eyes to the corresponding PDF and show that by choosing the peak PDF, the algorithm settles to the correctly equalized eye. . . .	39
3.16	Measured results show that the adaptation algorithm automatically increases equalizer peaking and preamplifier gain to compensate for increasing channel losses across (a) coaxial cables and (b) PCB traces.	40
3.17	Eye diagram of the receiver output after automatic adaptation when receiving 10 Gb/s data sent across a 10 m coaxial cable.	40
4.1	Oscillation amplitude of an LC tank will decay due to resistive losses unless compensated for by the addition of a negative resistance.	43
4.2	Injection of an oscillating signal, I_{inj} , can cause the LC tank output to lock to the frequency, ω_{inj} , of this injected signal.	44

4.3	Shorting the tank during narrow injection pulses makes application of the frequency domain model difficult.	46
4.4	Equivalence between (a) the physical model of the series resistive losses in the inductance of an LC tank and (b) the parallel resistance used to develop the frequency domain model breaks down for low Q values. . . .	46
4.5	Impulses applied to an oscillator have varying impacts on the oscillator output depending on the relative phase at which they are applied.	47
4.6	Example impulse sensitivity functions for oscillators with (a) sinusoidal and (b) square wave outputs.	48
4.7	Spectre simulations of the I and Q states of a four-stage VCO show that (a) an injected impulse causes a perturbation (dotted line) from the steady state (solid line). Repeated impulses (b) can lock the VCO to a different frequency but this requires a new ISF to model the oscillator's new trajectory through state-space.	50
4.8	Injected impulses result in step changes in the oscillator output, which must be accounted for by introducing step changes to the ISF in order to model injection locking.	50
4.9	Dividing an injected signal into impulses that act immediately on the ILO output phase also shift the corresponding $\Gamma(t)$ function, allowing for injection locking to be accurately modeled by the ISF.	51
4.10	Transistor-level simulation of an ILO's response to an injected impulse (a) determines the ISF of the oscillator. This ISF can then be used to predict the ILO's sensitivity to pulses with (b) larger amplitudes or (c) wider pulse widths but fails to accurately predict large, wide pulses (d).	53
5.1	Definitions of the ILO input and output signals that will be used to develop the ILO model in the remainder of this chapter.	56

5.2	By comparing the zero crossing times of the ILO output to that of an unperturbed copy, the phase change created by one period of the injected signal can be determined through transistor-level simulation in Spectre.	57
5.3	The PTC, $P(\phi)$, is determined through simulation by applying one period of the injected signal, $V_{inj}(t)$, at different phases, ϕ_n , relative to the oscillator's output signal, V_{out} , and observing the resulting change in output phase, P	58
5.4	Model representing nonlinear the phase relationship between the injected signal and the ILO output.	60
5.5	Transistor-level simulation results (a) comparing ILO output to an ideal reference signal at ω_{lock} shows that (b) the phase change produced by an injection event is sufficient to cancel the phase drift resulting from the difference between ω_{lock} and ω_0 . This allows the ILO to settle to the steady-state phase behaviour plotted in (c).	62
5.6	Steady-state phase relationships are determined by the difference between ω_{lock} and ω_0	62
5.7	An ILO's lock time depends on the initial phase difference, $\Delta\phi$, between the injected signal ϕ_0 and the steady-state phase difference, ϕ_{ss} , required by the frequency of the injected signal.	65
5.8	Spectre simulations of the (a) PTC and (b) transient phase response of a 4-stage ring ILO. When injection begins far from ϕ_{ss} at ϕ_{01} the lock time is significantly longer than when it begins at ϕ_{02}	66
5.9	Lock time varies greatly depending on the phase at which the injected signal begins. This effect is seen in both SPICE-level simulation simulation and the PTC-based lock time model.	67

5.10	The jitter tracking bandwidth of an ILO can be determined by applying a step change to the phase of the injected signal and observing the resulting change in the output phase. The displayed results are from Simulink simulations of the proposed model, performed using the PTC of the ILO as determined from Spectre simulations.	68
5.11	Comparison with direct transistor-level simulation using Spectre shows that the 3-dB jitter tracking bandwidth can be accurately predicted over a range of injected frequencies using the PTC model.	69
5.12	Spectre transient simulations of an ILO settle to a constant phase if the ILO is correctly locked (dotted line). If the injected signal is beyond the ILO's lock range then this can be identified by slipping of the output phase (solid line), which may not become apparent until the simulation has been run for many output clock cycles.	71
5.13	One stage of the four-stage CML injection locked ring oscillator. Applying the injected signal to a secondary differential pair provides a strong injection strength.	73
5.14	The PTC is determined by measuring the output phase change created when a single pulse is applied to the MILO at different phases relative to the oscillator's output signal.	74
5.15	Injection into multiple ILO locations can increase locking range if the injected signal experiences a delay that is equal to that created by each stage of the ring oscillator.	75
5.16	Spectre simulation results show that the peak-to-peak amplitude of the PTC increases as more injection sites are added.	76
5.17	Creating pulses at the reference clock edges emphasizes the desirable harmonic of the input thereby improving the lock range of the MILO.	78

5.18	The addition of a second edge detector with wide pulse widths further emphasizes the desired harmonic of the input signal.	78
5.19	Spectre simulations determine the PTC for ILOs using (a) one edge detector, (b) two edge detectors and (c) two edge detectors used to produce a sinusoidal (NRZ) injection signal.	79
6.1	The design of a frequency agile clock multiplier that is suitable for fast power cycling can achieve link flexibility and power savings in DVFS applications.	82
6.2	This chapter presents the first clock multiplier that is frequency-agile and has the ability to be powered on in under 10 cycles of the reference signal.	83
6.3	Four MILOs with adjacent lock ranges can cover an aggregate output frequency range from 2 to 4 GHz.	84
6.4	The addition of a second, identical ring oscillator compensates for DJ introduced by unequal pulse widths created by the edge detectors.	85
6.5	Spectre simulations show improvement in DJ performance without sacrificing lock range.	86
6.6	Latches at the output of each stage of the first ILO can be used to compare the phase relationship between the two ILOs.	87
6.7	Latch chains verify the multiplication factor by ensuring that there are exactly 2 rising and 2 falling clock edges within one half period of the reference clock.	88
6.8	A MILO can be powered down immediately if the TDC detects an out-of-lock condition or if the edge counter detects an incorrect frequency multiplication ratio.	89
6.9	If two MILOs are locked to the correct frequency, the power-down decision is made by determining which MILO is operating closest to its free-running frequency.	90

6.10	Power down of any unused MILOs is accomplished by blocking tail currents in all CML stages.	90
6.11	Converting the power down signal to CMOS logic levels ensures successful power-down and minimizes leakage power.	91
6.12	Timing diagram of the power-on sequence for a 1-GHz reference signal.	91
6.13	Power down of individual MILOs is enabled after 8 cycles and power down resulting from comparison of two correctly locked MILOs is enabled after 10 cycles.	92
6.14	Using four delay stages and an XOR gate as the building block for each component in the MILO ensures good matching between pulse widths and ILO free-running frequencies.	93
6.15	Die photo of the clock multiplier in 65-nm GP CMOS.	94
6.16	Measured free-running frequency of each MILO shows reasonable spacing between adjacent lanes and, if necessary, these free-running frequencies can be adjusted using the varactor control voltages.	95
6.17	Measured values of (a) ϕ_{ss} for various frequency offsets can be translated to (b) the PTC of the ILO. These measurements show good agreement with SPICE-level simulations.	96
6.18	Measurements show wide lock ranges for MILO1, MILO2 and the Breakout MILO, but problems with the reference clock distribution cause a reduction in the lock ranges of MILO3 and MILO4.	97
6.19	Two MILOs are able to increase the overall multiplier lock range to 55.7%. The point at which logic switches between MILOs is illustrated by measurements of the steady-state (SS) current drawn by each MILO, as well as the average (avg) current drawn when the multiplier is active for 50% of the time in 50-ns bursts.	99

6.20	Measured results show that the addition of ILO2 provides a reduction in DJ created by unequal pulse widths in the injected signal.	100
6.21	Decomposition of the measured jitter shows a DJ pattern that repeats in a 4-bit pattern due to the pulse widths of the injected signal.	101
6.22	A histogram of the total measured jitter shows four distinct peaks representing the DJ.	101
6.23	Including package parasitics at the simulated clock outputs identifies the cause of small measured clock amplitude.	102
6.24	Increasing driver power (a) increases output swing to approximately 100 mV _{pp} per side and (b) results in a decrease in DJ.	103
6.25	Histogram (a) and oscilloscope capture (b) showing that the RJ of the clock signal is approximately 1.4 ps-rms.	103
6.26	Measurements at the output of both ILOs shows that the RJ remains near 1.4 ps-rms across the entire lock range.	104
6.27	Test setup used to capture the transient response to an applied Start signal.	105
6.28	Repeated 50-ns bursts allow for power-on transient behaviour of the multiplier to be captured in (a). Zooming in on the start of these bursts (b) shows 3 ns delay between Start signal and output oscillations	106
6.29	Test setup used to capture the transient response of two output clocks.	106
6.30	Transient startup behaviour when (a) two and (b) three MILOs are enabled shows identification and power down of unlocked clock(s) in 10 reference clock cycles.	107
6.31	Simulated transient power consumption during one power-up cycle of the multiplier is compared to measured steady-state power consumption.	108
6.32	Energy consumed by the power-up sequence means that efficiency of the multiplier suffers when used in short bursts.	109

6.33	Average power consumption of the multiplier scales linearly with the percentage of active time.	110
6.34	Comparison to other clock multipliers shows the improvement in power-on time and frequency range achieved by this work.	110
B.1	Schematic of the variable threshold slicer used in the AFE adaptation scheme. Transistor sizes are given in μm with all gate lengths set to minimum size.	117
B.2	Schematic of the CML latch used in the TDC and power down logic of the MILO. Transistor sizes are given in μm with all gate lengths set to minimum size.	117
B.3	Schematic of the CML XOR gate used in the pulse generators and ring oscillators of the MILO. Transistor sizes are given in μm with all gate lengths set to minimum size.	118
B.4	Schematic of the two-stage, $50\ \Omega$ output driver used to send the MILO output off-chip. Transistor sizes are given in μm with all gate lengths set to minimum size.	118

Chapter 1

Introduction

1.1 Motivation

Wireline transceiver speeds have increased dramatically over the past decade. While this increase in speed is a driving force in the industry, other considerations such as transceiver flexibility are becoming increasingly important in next-generation devices. In many applications the need for transmitters and, to a greater extent, receivers that are adaptable and can function correctly in a variety of situations is becoming increasingly apparent. These applications include short range communications such as memory links for mobile devices, medium length transmissions such as ethernet links, and long distance lines such as cable TV and internet links. The following sections outline the motivation for adaptive wireline receivers in each of these applications.

1.1.1 Cable TV and internet

In cable TV and internet communication links, the downstream channel uses a point-to-multipoint scheme, which means that a carrier always exists, even when there is no packet being transmitted. As a result, the subscriber set-top box or modem has ample time to acquire synchronization to the downstream channel. The upstream channel,

however, which uses multipoint-to-point transmission, can be much more challenging since bursts from different subscribers arrive with different power levels, symbol timing, carrier phase/frequency offsets, and channel distortions [2]. The ability to receive and decode data in such a link requires a receiver that is flexible and able to adapt to a wide variety of operating conditions in real time.

1.1.2 Ethernet

As standards such as 100 Gigabit Ethernet continue to push the speeds of networks, the power consumption that such links require is encouraging the development of smarter link options. In this context, “smarter” means the ability to make the best use of available power and bandwidth constraints. One way of achieving this is through the use of real-time systems that monitor network congestion and adapt the video quality and bit-rate of streaming video. Such strategies to make best use of available bandwidth have recently been proposed [3].

Similarly, other protocols such as Energy-Efficient Ethernet focus on saving energy by powering down transceivers during periods of inactivity. In addition, although typical ethernet standards require that transmitters be able to provide enough power to send a signal up to 100 m, this new standard recognizes that significant power saving can be realized by determining channel length and reducing power for shorter cable lengths, which are often used in home networks. Each of these techniques require an intelligent receiver that is able to adapt to different data rates, received power levels, or amounts of channel attenuation.

1.1.3 Memory links

The speed of next-generation mobile memory interfaces for consumer electronics is continuing to increase with aggregate rates projected to reach 12.8 Gb/s in the near future [4]. One factor impeding this progress, however, is the fact that battery energy density

and thermal dissipation constraints are expected to remain essentially constant, which places a premium on both average and peak interface power consumption [4].

Further complicating the situation is the fact that bandwidth utilization in memory links varies by orders of magnitude over time, not just as different applications are executed, but also for short periods of time within an individual application. This, combined with the fact that average memory utilization is typically only a fraction of the peak bandwidth [4], is making it necessary to design an interface that is able to adapt to these dynamic workloads efficiently. In addition to this, the incorporation of a wide operating frequency range is becoming an attractive feature since this ensures interoperability with legacy devices [5], [6]. In order to make this practical, however, the power consumed by such a memory link must also scale with the operating frequency, which requires a great deal of transceiver flexibility in such a link.

1.2 Thesis Overview and Major Contributions

1.2.1 Analog Front-End Gain

In order to accommodate the wide range of requirements listed in the previous section, modern wireline receivers are, out of necessity, becoming more flexible. While variable data rates have become commonplace [7], [8], [9], recently reported receivers have gone a step further and now offer compatibility with a variety of standards such as PCIe, SATA and 1 to 10-Gb/s Ethernet [10]. A vital requirement of any such receiver is the ability to intelligently adapt the gain of a receiver's analog front-end (AFE) in order to function correctly in a variety of channel conditions.

The first stage in any receiver should be well matched to the characteristic impedance of the channel. It should provide as much gain as possible in order to optimize the receiver's noise performance and to open the received eye sufficiently for bit decisions to be made with a sufficiently low error rate. However, a gain that is too large can introduce

non-linear distortions, which are especially problematic in ADC-based receivers [11] or links employing multi-level signaling [7]. As a result, the addition of a variable gain amplifier (VGA) to the AFE can help improve the dynamic range of a receiver, thereby making it suitable for a wider range of channel conditions or transmitted powers.

Chapter 3, sections 2 and 3, introduce a high-bandwidth, low-noise, preamplifier that provides 10 dB of gain control while maintaining a good broadband match to the channel impedance across all gain settings. The topology is similar to the single-stage, fixed-gain preamplifier presented in [12], however, in addition to the gain control, it introduces a system that monitors and automatically adjusts the power consumption in order to provide a constant common mode level at the amplifier output. In doing so, this output common mode control also helps to maintain a good input impedance match across all gain settings and provides the added benefit of increasing power consumption in order to provide improved linearity only when it is required to accommodate a large received signal swing.

Experimental results of a prototype chip in 65-nm GP CMOS verify functionality of this preamplifier up to at least 15 Gb/s and show a good impedance match with S_{11} remaining below -8 dB from 1 to 24 GHz over all gain settings. These results were published at the International Symposium on Circuits and Systems (ISCAS) in 2010 [13].

1.2.2 AFE Adaptation

Another component of a receiver's AFE that must be flexible in order to accommodate varying channel conditions is the equalizer (EQ). Even in links using fixed channels such as PCB traces, channel losses can vary by several dB as changes in temperature or humidity are encountered. Furthermore, these variations have been shown to increase as the frequency of operation increases [14]. Therefore, whether EQ adaptation is implemented using a decision feedback equalizer or a linear equalizer, its coefficients must be intelligently adapted in order to compensate for these varying amounts of loss.

As a result, the method used to implement this adaptation is extremely important and has therefore received a great deal of attention in recent research publications. Currently, EQ settings are usually controlled by either minimizing the difference in frequency content between the EQ and slicer outputs [15], or monitoring the eye opening at the EQ output and adjusting either the decision threshold [16], the sampling time [17], or both [18] until a target bit-error rate (BER) is achieved. Unfortunately, the former method offers no guarantee that the resulting output eye is optimized and the use of several filters makes its implementation expensive in terms of chip area. Conversely, the latter method optimizes the received eye effectively but the eye monitoring required in adds a great deal of complexity and overhead to the receiver.

In addition to the preamplifier described in the previous section, Chapter 3 presents a method for simultaneous adaptation of both the EQ and the VGA by monitoring not only the vertical eye opening, but also the statistical distribution of the received data levels. This adaptation technique requires the addition of very little circuitry to the receiver and its ability to adapt both the EQ and VGA either while the circuit is active, or in an initial “set and forget” calibration sequence, makes its implementation practical, economical and effective. This research was published at the Custom Integrated Circuits Conference (CICC) in 2010 [19].

1.2.3 Adaptive Clocking

As discussed previously, the ability to switch between different power modes and operating frequencies can help to improve the energy efficiency of a wireline receiver. Unfortunately, the implementation of this functionality presents a significant challenge to the design of the receiver’s clocking systems. The long lock time required by traditional phase-locked loops (PLLs) when powering up or transitioning between operating frequencies often makes them unsuitable for systems that require these features. As a result, the receiver clock usually remains active at all times, with recent approaches relying on the

use of a variety of power-modes and bursting techniques to reduce average power [4].

While such techniques are effective, the need to operate the clock generating circuit at all times introduces a significant power penalty. In addition, the complexity involved in using many different power states with varying degrees of activity can often make their implementation impractical [20]. As a result, recent work has focused on implementing only two power states—either completely off or completely on—and on reducing the transition time between these two options [21].

Chapter 5, section 3 presents a 2.8-GHz clock multiplier based on a multiplying injection-locked oscillator (MILO) that is capable of transitioning from zero power to fully operational in less than 8 ns. This MILO is designed to have a lock range equal to 7%, which is sufficient to allow it to lock to its 2.8-GHz free-running frequency despite variations in voltage and temperature. This MILO was designed as part of a team while on internship at Rambus Inc. and was published as part of a fully bidirectional link at the Symposium on VLSI Circuits in 2011 [21].

Chapter 6 extends this work to show that the lock range of a MILO can be increased to over 40% of its free-running frequency through the use of frequency pre-conditioning. It then shows how the lock range can be further extended, in theory to any desired value, by using parallel MILOs with adjacent lock ranges. The resulting circuit minimizes the resulting power penalty through the use of several logic blocks that evaluate the output signals of adjacent MILOs and power down all but the one providing the clock at the desired frequency. These logic blocks operate quickly, allowing the circuit to transition from zero power to operation at the desired frequency in less than 10 cycles of the reference clock. The result is the first clock multiplier that is able to combine fast power-cycling with the ability to adapt to a wide range of operating frequencies. This frequency agility is achieved without the use of any external controls such as manual tuning of the oscillator's operating frequency. This work was developed at the University of Toronto, not during the Rambus internship, and has been submitted for publication at CICC in

September 2012.

1.2.4 ILO Modeling

Although MILOs are typically only capable of achieving narrow lock ranges on the order of a few percent of the free-running frequency [22], the lock range of the MILO described above was increased significantly with the help of a new, simulation-based technique for ILO modelling.

Traditionally, ILOs have been analyzed using the frequency domain model, first proposed in [23]. Although originally limited to LC oscillators using small, sinusoidal injected signals, variations of the frequency domain model have recently been introduced in order to handle large injection strengths, ring oscillators and other injected signal shapes [24], [25], [26]. Although attempts have been made to generalize the model to make it applicable in a wide range of situations [26], the results remain complex. In addition, all versions of the model use a variety of quasi-physical variables, which are not well defined and whose values are difficult to determine before measured results are available. As a result, this model is difficult to apply during the design phase of an ILO and is therefore often used only as a way of verifying or explaining measured results.

As an alternative, the phase transfer characteristic (PTC), which is an extension of the impulse sensitivity function [27], is proposed as a universal method to model an ILO's behaviour. This method uses SPICE-level simulations to determine the change in an oscillator's output phase in response to any injected signal. This phase response at the output is dependent on the phase at which the injected signal is applied, which gives rise to the term "PTC". Chapter 5 explains this concept in detail and shows how it can be used to predict the lock time, jitter tracking behaviour and lock range of an ILO. This technique is then demonstrated through the design of a MILO with wide lock range that is suitable for fast power-on applications, as described in the previous section. This work has been submitted for publication in *Transaction on Circuits and Systems - I (TCAS-I)*.

Chapter 2

Analog Front-Ends

Frequency-dependent channel impairments such as dielectric loss and skin-effect can severely limit a wireline receiver's ability to operate at high-speeds. As a result, many receivers incorporate an analog front-end (AFE) in order to compensate for these impairments and improve the quality of the received data before it can be converted to the digital domain. Although the elements contained in an AFE can vary depending on the intended application, variable gain amplifiers (VGAs) and equalizers (EQs) are two common circuits found in many front ends. This chapter presents a study of the state-of-the-art of each of these two components in order to serve as a foundation for the contributions presented in the following chapter.

In addition to the AFE circuit blocks themselves, this chapter examines the important issue of adaptation of their control signals. The ability to automatically adapt the control signals of the AFE components has a direct impact on their effectiveness and has therefore been studied extensively. A review of currently published adaptation schemes is presented in this chapter, which will provide a basis for the adaptation method proposed in the following chapter.

2.1 Variable Gain Amplifiers

In DSP-based receivers, or any link employing multilevel signaling, the analog front end (AFE) of the receiver must avoid introducing non-linear distortions to the signal before passing it to the multilevel slicer (ADC). In order to achieve the best possible dynamic range, the variable gain stage should be implemented as close to the receiver input as possible, so that overall gain can be reduced before nonlinearities are introduced by the input stages.

One method commonly used to implement a variable gain amplifier is to use a differential pair with variable source degeneration, as shown in Fig. 2.1. The addition of source degeneration resistor R_s helps to reduce the the signal swing applied between the gate and source of the input transistors, thereby making the input/output characteristic more linear [28] with a gain, A_v , given by

$$A_v = \frac{R_d}{R_s}. \quad (2.1)$$

By controlling the value of R_s the gain can therefore be controlled directly. If this stage is used as the first stage in a receiver's AFE then resistive components R_1 and R_2 can be added to the input to achieve a broadband impedance match. Control of these resistance values can also be implemented to achieve further gain control by attenuating the signal before it arrives at the VGA [29], resulting in a total overall gain of

$$A_v = \left(\frac{R_2}{R_1 + R_2} \right) \left(\frac{R_d}{R_s} \right) \quad (2.2)$$

A second, more recently proposed method for implementing a VGA is to vary the effective size of the input transistors by turning unit-sized differential pairs on or off, as shown in Fig. 2.2 [30]. In this topology since the bias current is kept constant regardless of the number of active differential pairs, N , the gain of the circuit is proportional to




Figure 2.1: Variable gain can be implemented by using source degeneration with adjustable resistance. Resistors R_1 and R_2 can be added to increase gain control and provide a broadband impedance match [29].




Figure 2.2: Variable gain can be achieved by varying activating a variable number of unit-sized differential pairs, effectively changing the input transistor width [30].

\sqrt{N} as

$$A_v = \sqrt{2\mu_n C_{ox} \frac{NW_u}{L} I_D R_d} \quad (2.3)$$

where W_u is the width of a transistor in a unit-sized differential pair. Therefore, for low-gain settings the current density in each active transistor is increased, helping to reduce distortion of the large received signal.

While this technique has been shown to achieve good dynamic range, with gain variation of over 6 dB per stage [30], it is difficult to implement as the first stage of an AFE since creating a broadband impedance match between the variable-sized input devices




Figure 2.3: The shunt-feedback configuration can provide a low-noise, broadband impedance match in electrical wireline applications [1].

and the channel is challenging. As a result a resistive matching network is often used, which can be detrimental to the bandwidth and noise performance of the preamplifier.

As an alternative, it is possible to eliminate the need for a purely resistive input matching network through the use of the shunt-feedback configuration commonly associated with optical-input transimpedance amplifiers. This topology, displayed in Fig. 2.3, is analyzed for electrical wireline applications in silicon technologies in [1] and is shown to outperform the preamplifiers in Fig. 2.1 and Fig. 2.2 in terms of bandwidth and power dissipation.

The linearity of this input stage can be improved by drawing current through feedback resistor R_f with the use of transistor M_1 as shown in Fig. 2.4 [12]. This raises the CM level at V_{out} and obviates the need for a common-source level-shifting stage following the preamplifier, which in turn increases the voltage swing that can be achieved at V_{out} without driving the common source transistor into triode operation. While the addition of M_1 to the receiver input can degrade noise performance, this can be mitigated by maximizing this transistor's size and minimizing its g_m [12].

The addition of gain control to shunt-feedback preamplifier stages can be achieved by adjusting the value of R_f [31] since this resistor determines the low-frequency tran-




Figure 2.4: Drawing bias current through R_f raises the CM level at the output and increases the voltage swing headroom for linear operation.

simpedance gain of the preamplifier, R_T , which can be approximated by the equation

$$R_T = R_f \frac{A}{1 - A} \quad (2.4)$$

where A is the open-loop voltage gain of the amplifier [32]. Changing R_f , however, also has the effect of changing the input impedance of the preamplifier, approximately given by

$$R_{in} = R_f \frac{1}{1 - A} \quad (2.5)$$

Since the preamplifier input must be matched to the impedance of the channel to avoid reflections, this change in R_{in} is undesirable. A common solution to this is to implement gain control in the second stage of the receiver chain, maintaining a fixed preamplifier gain and impedance match [33]. However, the fixed gain of the input stage using this solution limits the dynamic range of the receiver.

2.2 Equalization

The frequency-dependent losses introduced by the lossy channels used in most wireline communications links can severely degrade the quality of the received signal. This can in turn limit the operating speeds unless some method can be introduced to equalize the amount of loss encountered at all relevant frequencies. This can be accomplished through the use of either decision feedback equalizers (DFEs) [34], continuous time linear equalizers (CTLEs) [35], or some combination of the two schemes [10].

Of these two options, DFEs typically outperform CTLEs due to their ability to adapt a wide variety of channel conditions and to compensate for post-cursor intersymbol interference and reflections without enhancing noise or crosstalk. In addition, although typically limited to lower speeds than CTLEs, they have recently been shown to operate at speeds of up to 19 Gb/s [36].

This thesis, however, makes no effort to advance the state-of-the-art in equalizer design. Instead, the EQ used in the AFE presented in the following chapter is present to illustrate the operation of the AFE adaptation algorithm. As a result, a CTLE was chosen for its simplicity and ease of implementation. Therefore the discussion of EQ background presented in this section is restricted to CTLEs.

2.2.1 Continuous Time Linear Equalizers

The measured S_{21} of three lengths of a 75 Ω coaxial cable channel are illustrated in Fig. 2.5. These measurements clearly show that the loss of these channels increases as the frequency content of the transmitted signal increases. In order to equalize these losses CTLEs aim to introduce some combination of low-frequency gain (or loss) and high-frequency gain to the AFE. This concept is captured by the conceptual diagram shown in Fig. 2.6.

This concept can be implemented in a literal fashion by using a combination of dif-

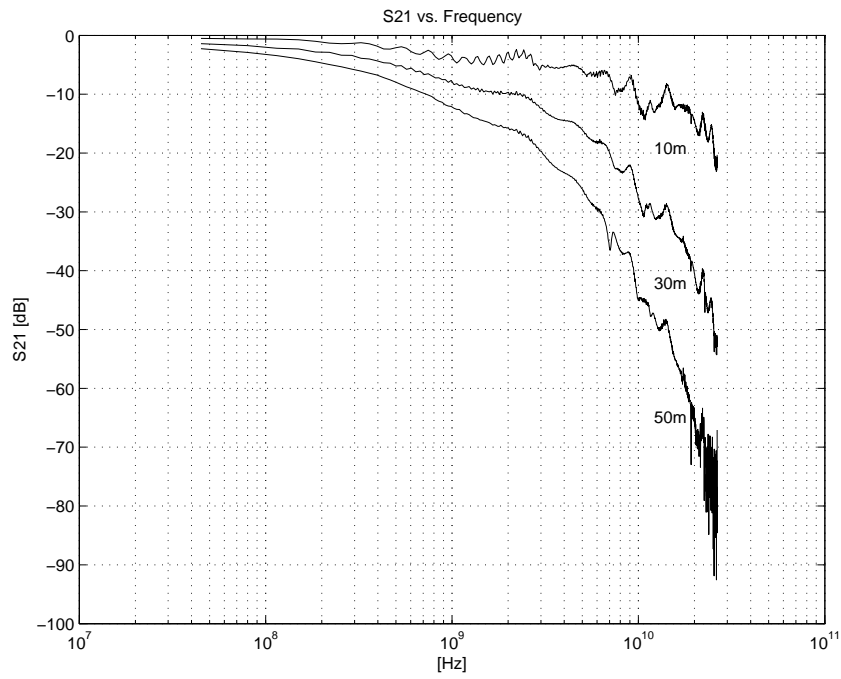


Figure 2.5: Measured losses of three coaxial cables illustrate frequency-dependent losses.




Figure 2.6: Channel losses can be equalized by introducing a combination of high frequency gain and low frequency gain/loss to the receiver.

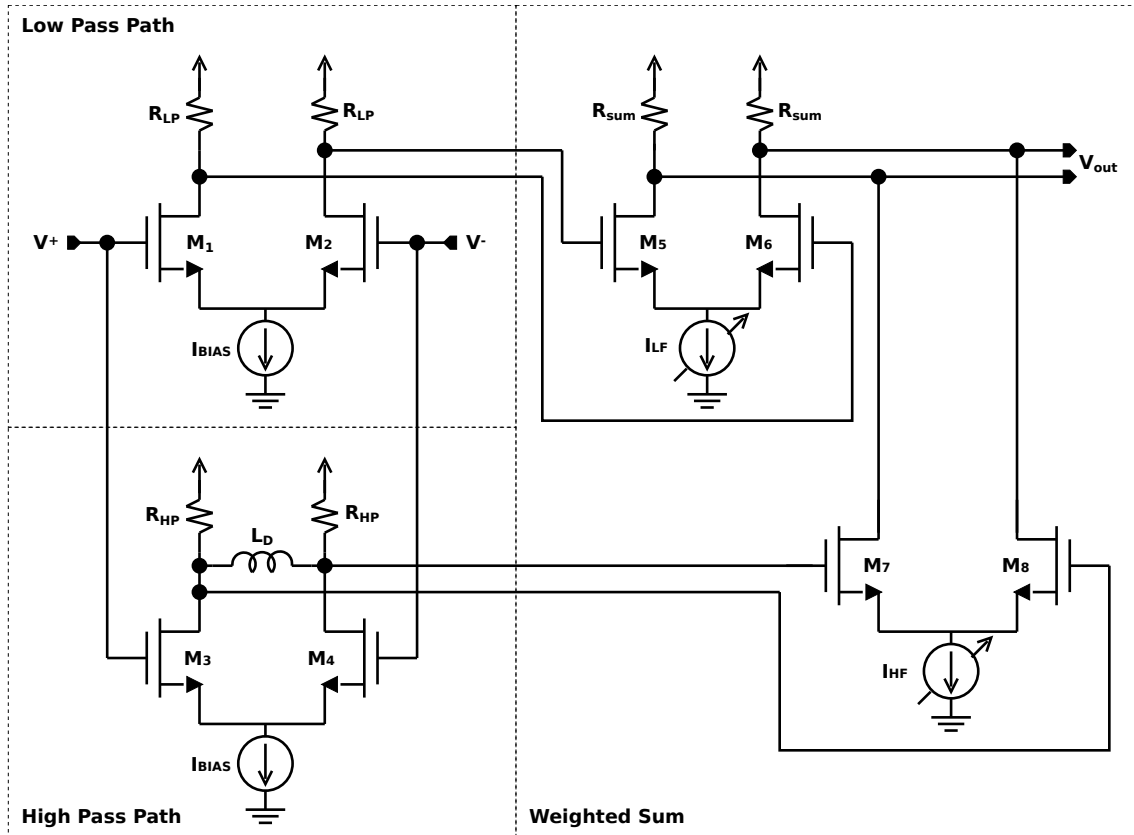


Figure 2.7: By combining the outputs of high-pass and low-pass amplification stages and adjusting the weights of their contributions, channel losses can be equalized across all frequencies of interest.

ferential pairs as shown in Fig. 2.7 [37]. By shorting the drains of transistors M_3 and M_4 at low frequencies, inductor L_D enables high frequency amplification of the input data while attenuating low frequencies. By controlling the current sources I_{LF} and I_{HF} the contributions from the high pass and low pass paths can be adjusting according to the losses of the channel.

While this technique offers an intuitive and straight forward way to implement equalization, a similar effect can also be achieved by introducing a low-frequency zero into a single amplification stage. One popular method of achieving this is through the addition of a capacitor to a source degenerated differential pair, as shown in Fig. 2.8. In this topology the source resistance R_s provides a reduction in gain, as described in the previous section. However, at high frequencies capacitor C_s becomes a short circuit and the




Figure 2.8: Source degeneration reduces the gain of the amplifier. At high frequencies this reduction is eliminated by the short circuit created by C_s resulting in higher gain at high frequencies and equalization of the channel losses.

effect of R_s is removed according to the gain equation

$$A_v = \frac{R_d}{R_s // \frac{1}{sC_s}} \quad (2.6)$$

$$= \frac{R_d(1 + sC_sR_s)}{R_s} \quad (2.7)$$

This allows for larger gain at high frequencies, thereby equalizing the high frequency channel losses. The frequency at which the gain begins to increase can be controlled by adjusting the value of C_s .

2.3 Analog Front-End Adaptation

The ability to intelligently control the AFE blocks in a way that optimizes receiver speed, sensitivity, dynamic range and noise performance, is non-trivial. As demonstrated in the previous section, channel losses can change dramatically when the length of the channel is changed. Even for fixed channels, variations in loss can occur due to factors such as temperature or humidity. As a result, control signals for the VGA and EQ, as shown in Fig. 2.9, should be generated automatically and should be able to adapt to a variety of

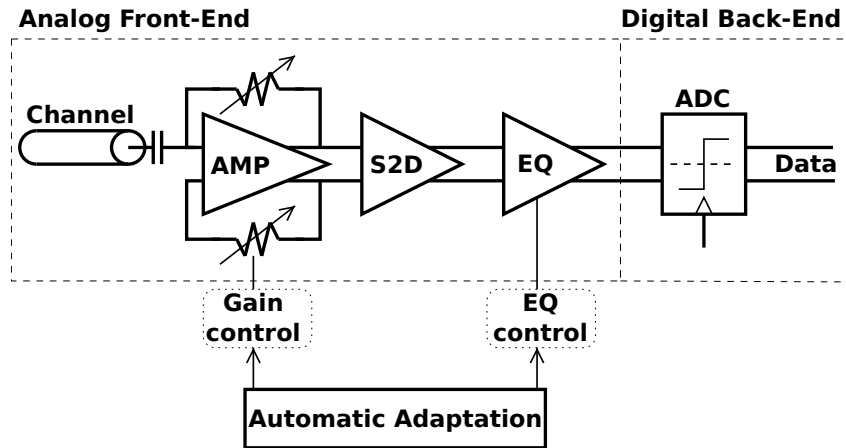


Figure 2.9: A typical wireline receiver with an AFE including gain and equalization control signals that should be generated automatically and intelligently.

channel conditions.

One way that the EQ control signal can be generated is by minimizing the difference between the high-frequency content of the EQ and slicer outputs [38] as shown in Fig. 2.10. While this technique can also include low-frequency gain control by adding a second, similar loop containing low-pass filters [15], the necessary analog filters can be difficult to design accurately and can consume a great deal of chip area. Furthermore, this technique offers no guarantee that the resulting EQ control achieves optimal bit error rate (BER) performance for the given channel conditions. Instead, this section focuses on digital adaptation schemes that monitor the opening of the received eye in the time domain in order to make decisions about AFE settings.

2.3.1 Eye Monitor-Based Adaptation

In order to optimize the BER, it is necessary to develop a picture of the eye opening at the slicer (or ADC) input. This requires moving the sampling point from its ideal location at the center of the eye until errors are detected by comparing this output to the actual data obtained by sampling at the ideal point. This movement away from the ideal point can be performed by using additional slicers with either a variable decision threshold for

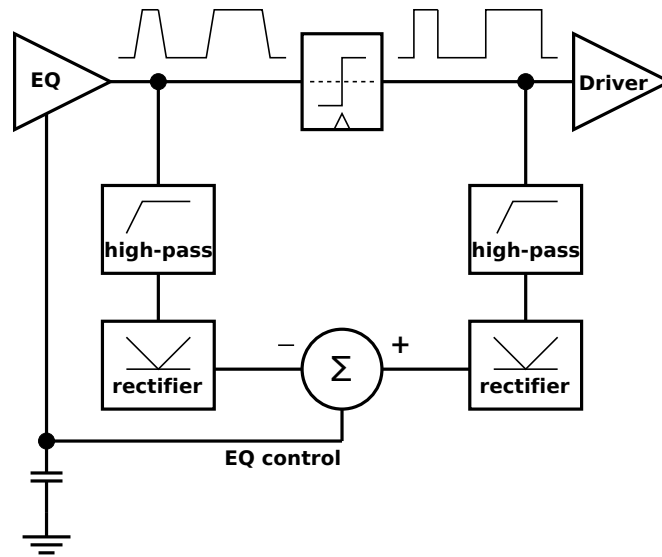


Figure 2.10: EQ adaptation can be performed by analyzing the high-frequency content of the EQ output but this does not necessarily maximize the resulting eye opening.

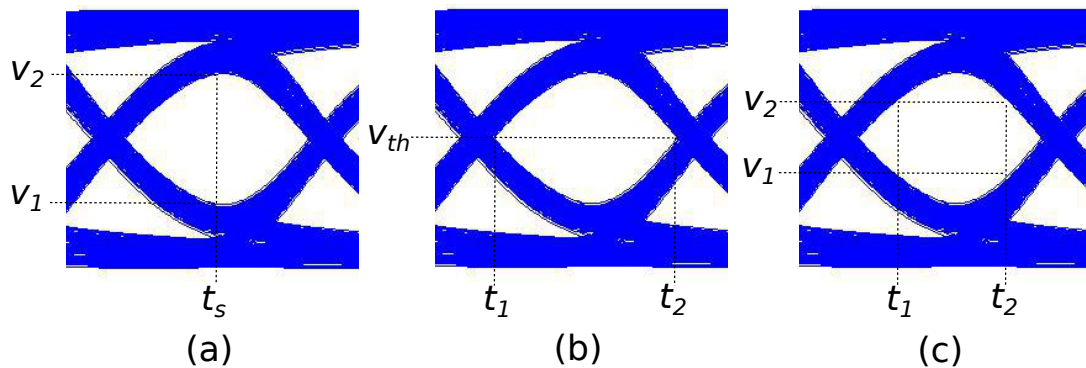


Figure 2.11: Adjusting the (a) sampling threshold, (b) sampling time or (c) both produces an indication of the eye opening that can be used to adapt the AFE settings.

vertical eye monitoring [16], variable sampling time for horizontal eye monitoring [17], or both variable threshold and sampling time for two-dimensional eye monitoring [18]. These three types of monitors are illustrated in Fig. 2.11 where the traditional goal for adaptation is to maximize the distance between V_1 and V_2 and/or t_1 and t_2 depending on the chosen scheme.

If the threshold levels and sampling phases can be varied independently and with a fine resolution then this technique can be extended further to obtain complete two-dimensional eye opening data and, by counting the frequency of errors at various sam-

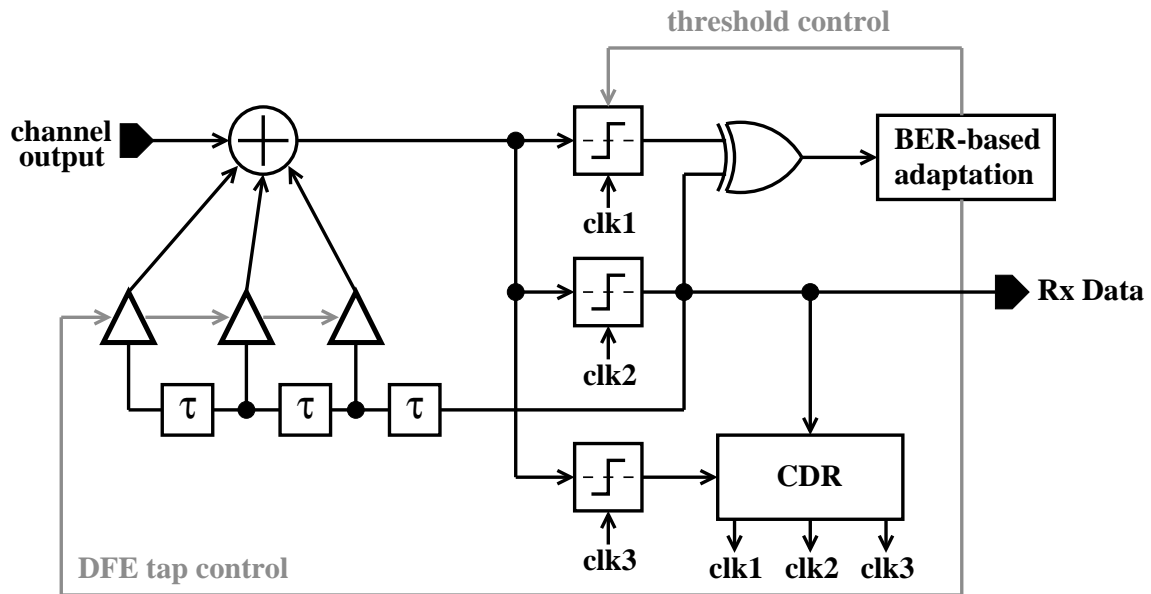


Figure 2.12: An example of a BER-based adaptation used to optimize DFE tap settings [39].

pling points, BER contours can therefore be developed. An example of a circuit used to implement this approach for adaptation of a DFE is shown in Fig. 2.12 [39]. By comparing the resulting BER to some target value, the adaptation algorithm is able to generate analog control signals for the taps of the DFE.

This technique can guarantee optimal receiver performance by minimizing the BER and represents the state-of-the art in adaptation techniques. It can, however, be very slow to converge and requires a significant amount of additional hardware including a high-speed XOR gate and error counter as well as multiple clocks with finely tunable phases. It also does not incorporate any gain control into the AFE, limiting the dynamic range of the receiver. As a result, the development of simpler adaptation schemes that can achieve similar performance present an attractive area for further work on this topic.

2.4 Summary

Variable gain and equalization are two important functions performed by many AFEs used in wireline communication links. Adding equalization at some stage of the AFE can help to mitigate frequency-dependent channel losses and help to improve the eye opening. Adding variable gain to the input stage of the AFE can help to maximize the dynamic range of the receiver. The use of shunt-feedback in the input stage can provide low-noise, broadband functionality and presents the opportunity to implement gain control by varying the resistance of the feedback path. This implementation, however, can degrade the impedance match with the channel, creating undesirable reflections.

Both gain and equalization controls must be generated automatically and should be intelligently adapted in order for the receiver to function correctly over a range of channel conditions. While two dimensional monitoring of the equalized eye diagram can adjust an equalizer's settings to guarantee optimal BER operation, the complexity involved in such an adaptation scheme can be prohibitive.

Chapter 3

Analog Front-End Adaptation

This chapter introduces a technique for the automatic adaptation of both the gain and equalizer control signals in an AFE. Rather than targeting a specific minimum BER as in [39], this algorithm tightens the distribution of the received signal amplitude and centers it at a specific, pre-determined optimal level. This ensures that the vertical dimension of the received eye is as open as possible for the given channel conditions, while avoiding excessive gain, which can compromise the receiver's dynamic range. This technique requires only a single sampling phase at the center of the received eye and is therefore able to operate independently of the CDR.

It also requires only minimal hardware overhead as it does not require a variety of clock phases or a high-speed XOR gate to count errors but instead uses only the DC output voltage of one additional comparator to determine the appropriate EQ and VGA control signals. These factors make this adaptation strategy attractive not only as an efficient means of optimizing AFE settings, but also a way of quickly and easily detecting faults and predicting circuit performance in a production testing environment, which can otherwise be costly and time-consuming [40].

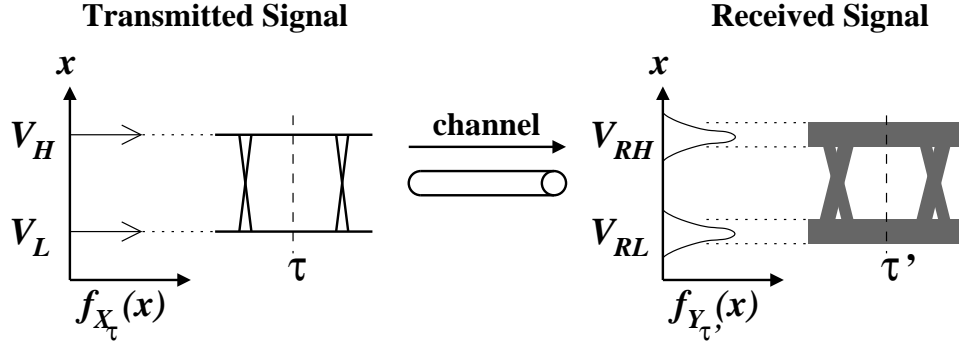


Figure 3.1: A noiseless transmitted signal encounters noise and ISI in a frequency-dependent channel, which results in spreading of the PDF of the received signal.

3.1 PDF-Based Adaptation

A sequence of random binary data can be represented by a random process X . If this data is generated by an ideal, noise-free binary transmitter and sampled at the baud rate, T_b , and at phase τ , corresponding to the midpoint of each bit then each sample is a random variable represented by $X_\tau(k)$ as

$$X_\tau(k) = X(kT_b + \tau) \quad (3.1)$$

where k is an integer. In the absence of noise or other non-idealities in the transmitter, $X_\tau(k)$ will be constrained to two possible states: logic high, V_H , and logic low, V_L , which are assumed to be equally probable

$$P(X_\tau(k) = V_L) = P(X_\tau(k) = V_H) = 0.5 \quad (3.2)$$

and whose values are defined by the voltage swing constraints inherent to the transmitter.

The result is the transmitted eye diagram and the probability density function (PDF) of X_τ , denoted by $f_{X_\tau}(x)$, is shown on left side of Fig. 3.1. Here the PDF is confined to two impulse functions corresponding to V_H and V_L . After passing through a channel and part of the AFE, the eye at the EQ input will show intersymbol interference (ISI) in

the form of jitter and a spreading of the received signal amplitudes due to the frequency-dependent losses of the channel, reflections at connection points, and bandlimited stages at the input of the receiver's AFE. A new random variable, $Y_{\tau'}(k)$, is then obtained by sampling the received data at a phase corresponding to the midpoint of each received bit, given by τ' . The PDF of this variable is given by $f_{Y_{\tau'}}(x)$ and is illustrated on the right side of Fig. 3.1.

The received data is no longer exclusively constrained to two discrete voltage levels but is instead spread over a range of values by the introduced ISI and noise. Hence, the PDF of the received data is not represented by impulse functions, but is instead distributed around new high and low levels, V_{RH} and V_{RL} . A detailed derivation of $f_{Y_{\tau'}}(x)$ in terms of the transmitter and channel characteristics can be found in [41]. Although the PDF has been used previously as an efficient way to detect the BER of a received signal [32], its use as a metric for the adaptation of AFE components has not. The premise of this work is to monitor the PDF of the received data and adjust gain and high-frequency peaking controls in the AFE in an effort to return to the ideal PDF of the transmitted signal.

3.1.1 Measuring the Received PDF

If the threshold voltage of the slicer at the AFE output is moved far enough away from the center of the received eye, the slicer will begin to generate errors. If this output is available at the same time as the error-free output from a slicer with the correct threshold level, then these errors can easily be detected by passing both slicer outputs to an XOR gate. Counting the number of errors generated by each threshold step can be used to generate BER contours, which can in turn be used to select an optimum slicer threshold [42] or optimize equalizer tap coefficients [39].

In this work this process is simplified by observing that as the threshold voltage is swept, the DC output of the slicer, V_{out} , is proportional to the cumulative distribution

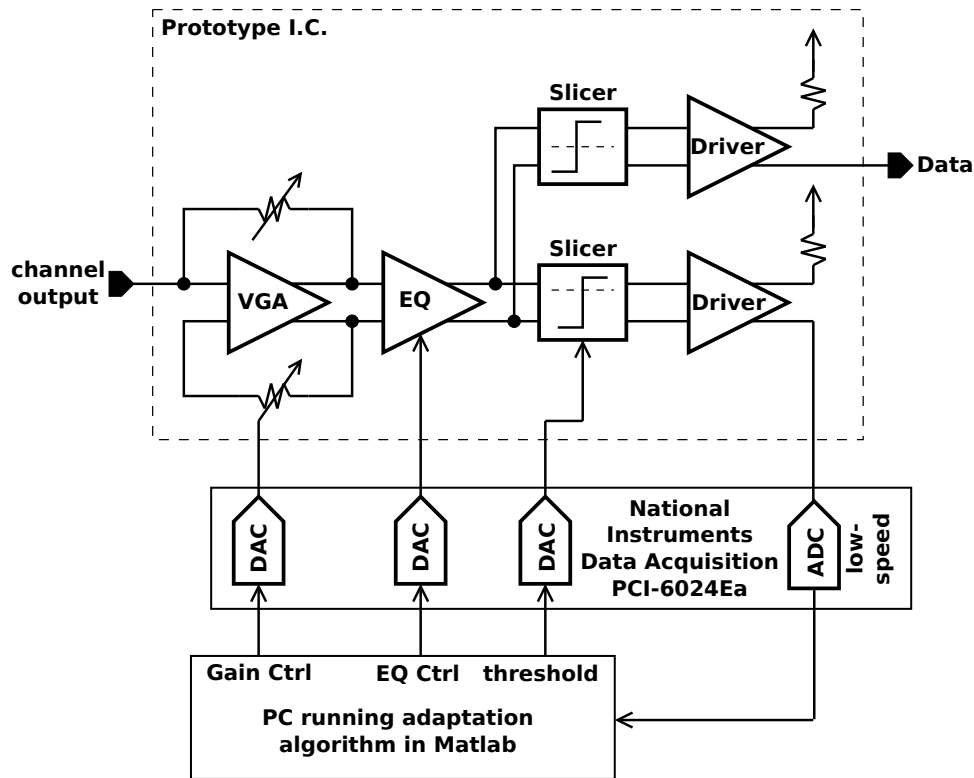


Figure 3.2: A block diagram of the VGA and EQ adaptation performed in this work. The low-speed ADC digitizes the DC (average) output of the offset slicer. This information is used by the adaptation algorithm to minimize the spreading of the received PDF.

function (CDF) of the input signal. Since the CDF of the received signal yields information equivalent to a BER contour, it is possible to replace the high-speed XOR gate and error counter with a simple low-pass filter. The prototype chip along with the off-chip hardware used to demonstrate this operation is shown in Fig. 3.2. The fabricated chip is a 10 Gb/s binary wireline receiver in 65-nm CMOS. A schematic of the variable-threshold slicer can be found in Appendix B, Fig. B.1. The output of this variable-threshold slicer is taken off-chip and passed to a low-speed ADC (i.e. an ADC preceded by a low-pass filter), which produces the DC (average) voltage of the receiver output. The ADC and DACs are provided by a National Instruments PCI-6024E data acquisition card, which runs at 200 kS/s and provides 12 bits of resolution. This card allows for the adaptation algorithm to be run using a PC where simple Matlab control logic is used to generate the gain, equalizer and threshold controls.




Figure 3.3: The DC output of the slicer drops as V_{th} is increased. The magnitude of the slope of this decline yields the PDF of the received data.

The PDF of the received signal is equal to the slope of the CDF obtained by sweeping the threshold of the auxiliary slicer, as illustrated in Fig. 3.3. The equalizer peaking is then chosen to maximize the slope of the CDF (i.e. the peak value of the PDF) since this corresponds to the narrowest possible spread of the PDF and therefore the lowest amount of amplitude noise in the received eye.

3.1.2 Equalizer's Effect on the PDF

The EQ implemented in this receiver is a continuous time linear equalizer, which creates high-frequency peaking by decreasing low-frequency gain. The chosen topology is the same as that introduced in [37], as shown previously in Fig. 2.7. This equalizer was implemented in 65-nm CMOS and Spectre simulation results using post-layout extraction, shown in Fig. 3.4, illustrate that the equalizer is capable of producing high-frequency peaking of up to 8 dB. It should be noted that this peaking occurs at a frequency close to 20 GHz, which was accidentally overdesigned for the target receiver speed. All transistors were implemented with minimum gate lengths and with widths and bias currents




Figure 3.4: Simulated gain of the EQ using Cadence Spectre with full RC-extraction of the circuit shows up to 8 dB of high-frequency peaking.

	Identifier from Fig. 2.7	Size
Transistors	$M_1 - M_8$	$20 \mu\text{m}$
Bias Currents	I_{BIAS}	6 mA
	I_{LF}	1 mA - 6.4 mA
	I_{HF}	6.4 mA - 1 mA
Resistors	R_{LP}, R_{HP}	100Ω
	R_{sum}	66Ω

Table 3.1: Component values used to implement the EQ design shown in Fig. 2.7.

as listed in Table 3.1.

To illustrate the adaptation of this EQ, Fig. 3.5 shows measured results taken from a prototype binary receiver fabricated in 65-nm CMOS. In Fig. 3.5(a) the DC slicer output voltage is used to infer the percentage of time that the slicer spends at logic high (V_{RH}), which is equivalent to the CDF of the received random variable. As the threshold level is increased, this percentage increases from 50% when the threshold is in the center of the received eye, to 100% when the threshold is outside the eye. When alternating data is transmitted, a sharp rise in the CDF is observed near a threshold level of 200 mV because there is no ISI present. Conversely, the presence of ISI when 2^7-1 PRBS data is transmitted means that a more gradual change is observed in the CDF. This effect is

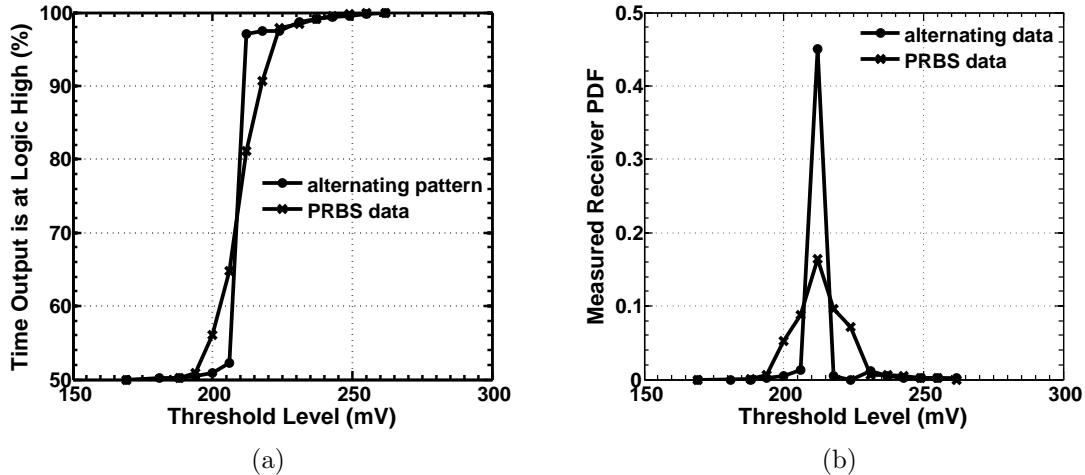


Figure 3.5: Measured DC output voltage is proportional to (a) the percentage of time that the output spends at logic high. The slope of this curve is then used to create (b) the PDF of the the received signal.

apparent in Fig. 3.5(b) where the PDFs of the two received patterns are compared.

The adaptation algorithm was tested when receiving a 2^7-1 , 2 Gb/s PRBS signal, transmitted across a 10 m BNC coaxial cable. Fig. 3.6(a) shows the PDF of the logic high level of the received data for five different equalizer peaking settings. As the EQ control voltage V_{eq} is increased, the low-frequency content of the received signal is attenuated, meaning that the peak slope of V_{out} occurs at incrementally lower threshold levels. At the same time, the resulting emphasis of the high-frequency content helps to reduce ISI, narrowing the PDF of the received data and increasing its peak value. In this case, the adaptation algorithm selects $V_{eq} = 0.65$ V to be the best equalizer setting because it produces the highest peak value of the PDF (i.e. the largest slope of the CDF). Since adjusting the settings of a peaking EQ effects the signal amplitude, this adaptation is performed before the VGA adaptation, which is described in the following section.

3.1.3 VGA's Effect on the PDF

The VGA can help keep the signal amplitude within some specified range throughout the receiver signal path in order to ensure the best possible dynamic range in the AFE.




Figure 3.6: Measurement results show that (a) increasing EQ peaking narrows the PDF and increases its peak value and that (b) increasing VGA gain moves this peak towards the target maximum threshold level, in this case set to 100 mV.

Fortunately, the amplitude of the signal at the equalizer output can be readily observed using the additional slicer. As the gain of the AFE is increased or decreased, the received signal PDF is observed as described in the previous section and its peak can be used to indicate the signal amplitude. For the 65-nm CMOS receiver in this work, it was determined that the voltage swing of the equalizer output eye should not exceed 100 mV (peak-to-peak per side) in order to avoid distortion.

As will be explained in the following section, the gain of the VGA used in this work decreases as control signal V_{gain} is increased. From the measured results in Fig. 3.6(b) it is apparent that the peak of the PDF shifts to higher voltages for higher preamplifier gain settings, as expected. In this case the adaptation algorithm chooses a V_{gain} setting of 1.2 V since this places the peak of the PDF as close as possible to the target value of 100 mV.

A single round of optimization of the EQ and VGA may be enough to ensure an open eye in a “set and forget” initial calibration. The adaptation algorithm used in such an implementation is illustrated in Fig. 3.7. If a set and forget adaptation is not appropriate for the intended application then it is possible to maintain optimal settings

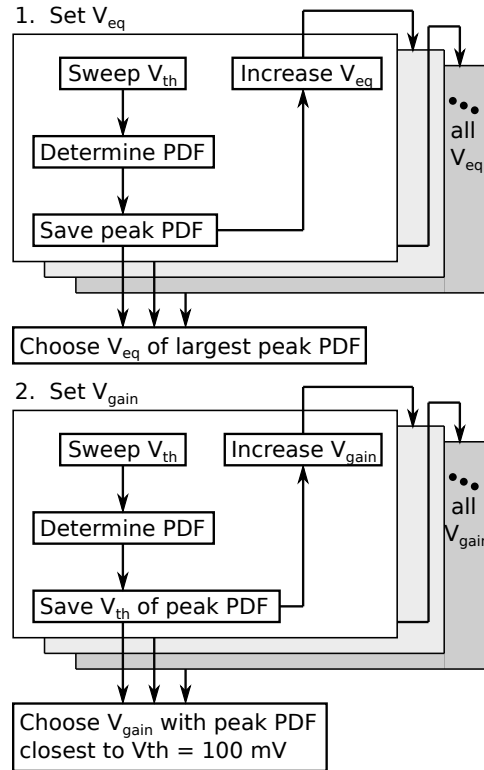


Figure 3.7: Flow diagram showing how the equalizer and gain settings are determined.

for the AFE as channel conditions vary over time by continuously repeating the EQ and VGA adaptation algorithms in the background. In this case, once the initial calibration algorithm is run, small changes to the gain and peaking settings can be tested to see if they improve the location and size of the peak PDF. Since both the equalization and gain control settings are determined from the same set of measured DC output voltages, requiring that only limited overhead be added to the receiver, the power penalty incurred by running these algorithms continuously is minimized. This adaptation technique was published at CICC in 2010 [19].

3.2 Variable-Gain Preamplifier Design

Variable gain is achieved in the AFE presented in this work by incorporating gain control into the input preamplifier stage, as illustrated in Fig. 3.8. This obviates the need for

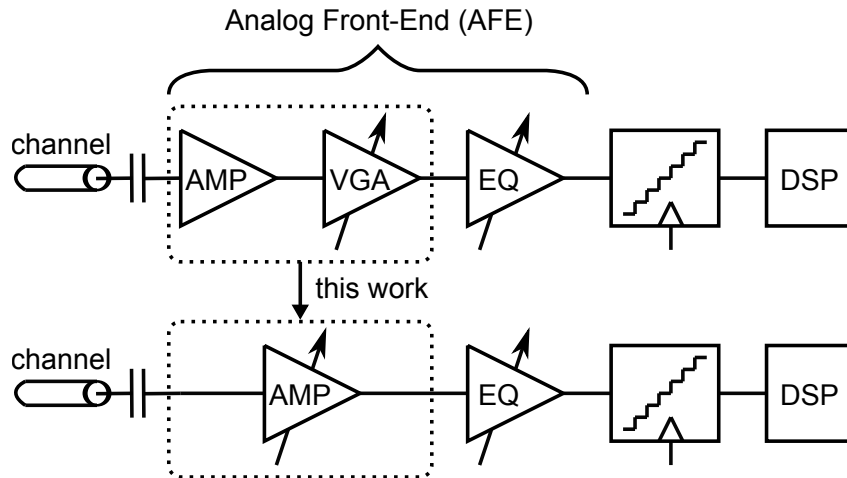


Figure 3.8: A preamplifier with gain control replaces a fixed-gain preamplifier and VGA in applications requiring high linearity and wide dynamic range.

a separate VGA stage, which can help to minimize power consumption and area while maximizing the dynamic range.

This section describes a high-bandwidth, low-noise, adjustable-gain preamplifier based on the fixed-gain topology shown in Fig. 2.4 [12]. In addition to implementing gain control to maximize linearity and dynamic range, the circuit presented in this section also introduces a control loop to automatically maintain a constant output common mode (CM) level across all gain settings in order to simplify the design of the following differential stages in the receiver. This output CM control also automatically adjusts the power consumption of the preamplifier in response to linearity demands and to help maintain a broadband input match across all gain settings. Fig. 3.9 shows a detailed schematic of the circuit and illustrates how the CM control is able to perform these tasks without loading the signal path of the TIA by operating on a copy of the preamplifier. This circuit, and the measured results presented in the remainder of this chapter, were published at ISCAS in 2010 [13].

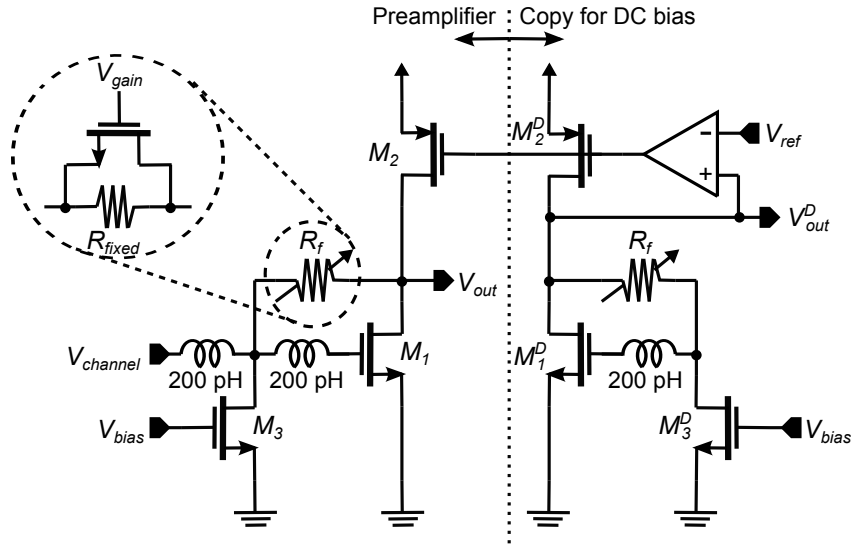


Figure 3.9: Schematic of the preamplifier introduced in this work.

3.2.1 Gain Variation

A single-ended signal arrives at the preamplifier at $V_{channel}$ where two inductors are used to resonate with the input capacitances to extend the bandwidth of the input impedance match. As discussed in the previous chapter, both the transimpedance gain and the input impedance of this amplifier are set by the value of the feedback resistance, R_f , and the open loop gain, $|A|$. Although a common solution to this problem is to separate the variable gain stage from the input stage, however, in this approach non-linear distortions may be introduced before the variable gain stages are reached.

The approach introduced in this work is to mitigate the impact of input impedance variations by decreasing $|A|$ along with R_f and by designing the input stage to be well matched to the channel when the received signal is smallest and therefore most sensitive to unwanted reflections. The gain of the RC-extracted preamplifier was simulated using Spectre and is shown for four values of V_{gain} in Fig. 3.10. These simulations were conducted using post-layout, RC extraction of the preamplifier and typical corners for transistor models.

Another concern with the variable gain implementation in this topology is that bias

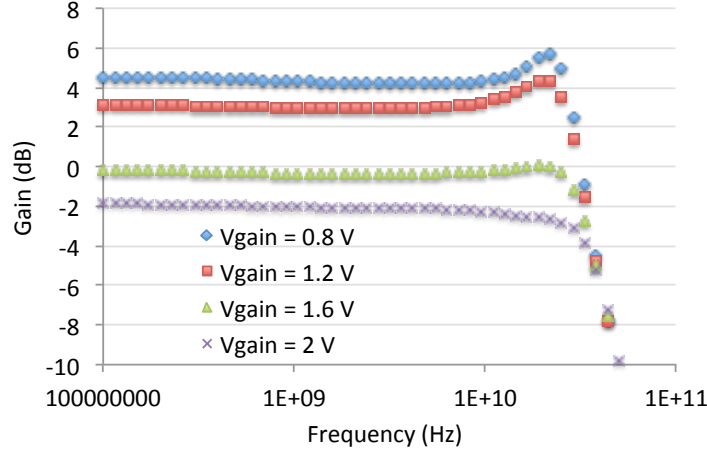


Figure 3.10: Simulated gain of the preamplifier as V_{gain} is varied. Simulations were performed using Spectre with RC-extraction of the entire preamplifier.

levels at the input and output of the amplifying transistor are set by transistor M_3 , which pulls a DC current through R_f . This means that varying R_f will also change the CM levels of the preamplifier, which necessitates the need for the CM regulation presented in the following section.

3.2.2 Output CM Control

A copy of the preamplifier ($M_1^D - M_3^D$) is used to replicate the DC biasing. The resulting pseudo-differential signal ($V_{out} - V_{out}^D$) is passed to subsequent differential stages providing power supply and common-mode noise rejection. However, in this topology changing R_f will impact the CM levels of the preamplifier. To control this, the op amp shown in Fig. 3.9 monitors the DC output level at V_{out}^D and fixes both V_{out} and V_{out}^D to V_{ref} by manipulating the current sourced by M_2 and M_2^D .

By reducing the effective impedance of M_2 and M_2^D at low gain settings the control loop also has the effect of reducing the magnitude of the open-loop gain, $|A|$, of the preamplifier. This helps to mitigate the change in input impedance according to Equation (2.5) and also further reduces the transimpedance gain according to Equation (2.4). Since this reduction in gain occurs without impacting the dominant pole of the amplifier, it

also results in an increase in bandwidth for these low-gain settings. Simulated results in a standard 65-nm GP CMOS process show that this effect leads to an increase in bandwidth by a factor of 1.6x when compared to a preamplifier with no CM level control.

The CM stabilization also adjusts the current density of M_1 . For low gain settings (used when the received signal swing is large) the DC bias voltage at its gate is increased, providing a large overdrive voltage and, hence, high linearity. Simulated results show that this setting results in a total harmonic distortion (THD) of less than -34 dB. For high gain settings (used when the received signal swing is small) the gate voltage of M_1 is decreased, sacrificing linearity for an improvement in its noise performance. This setting helps the preamplifier achieve good measured sensitivity of 8 mV. In total, the measured current drawn by 60 μm wide transistor M_1 ranges from 4.9 mA at the high-gain setting, to 24.3 mA at the low-gain setting.

The op amp used in this control loop is a simple single stage amplifier with NMOS inputs and an active, current mirror load. Simulated results show that the compensation provided by the gate capacitances of transistors M_2 and M_2^D is enough to ensure stability with a phase margin of at least 70° .

3.3 Preamplifier Measurement Results

The fabricated receiver occupies approximately 0.23 mm² (excluding pad frame) and has a measured maximum power consumption of 252 mA from a 1.2 V supply. The AFE accounts for as much as 67.1 mA of this total (with the preamplifier at the low-gain setting), with the rest being used in the output drivers and current-mode logic of the digital back end. A die photo of the fabricated receiver is shown in Fig. 3.11. All measurements were made on-wafer and alignment of the clock and data signals was performed manually off-chip.




Figure 3.11: Die photo of the receiver fabricated in 65-nm CMOS.

3.3.1 S-parameters

S-parameter measurements were taken to evaluate the input match of the receiver. The measured S_{11} results are shown in Fig. 3.12 for various gain settings. For the highest preamplifier gain settings ($V_{gain} = 0.8$ to 1.2 V), which are used when the received signal is smallest and a good input match is critical, S_{11} remains below -12 dB to well beyond 20 GHz. For the lowest preamplifier gain settings ($V_{gain} = 1.6$ to 2 V) the input match remains below -8 dB across the measurable range. For prototype testing, a dedicated, off-chip source was used to generate the preamplifier V_{gain} settings. High voltage devices are unnecessary as the gate-source and gate-drain voltages of the feedback transistor do not exceed 1.2 V at any gain setting. In later iterations a high-voltage generator might need to be used, or it might be possible to replace the NMOS feedback transistor with an equivalent resistance PMOS device and use low control voltages instead.

The gain variation of the AFE was measured using a network analyzer with very small input signals to keep the digital CML logic operating linearly. These measurements show that the preamplifier gain can be adjusted by 10 dB, while the low-frequency gain of the

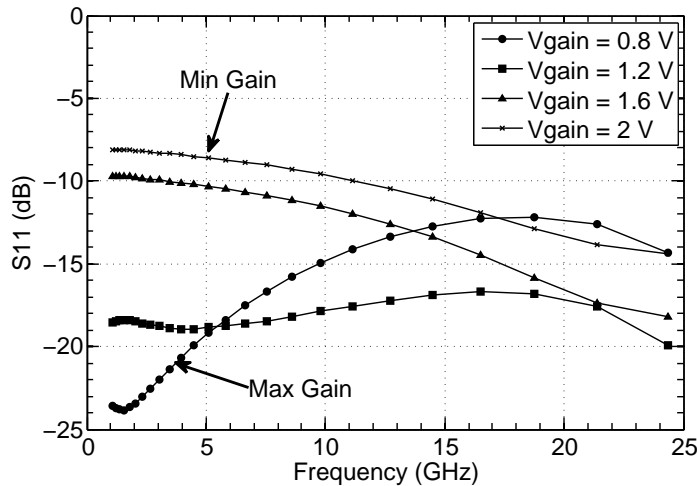


Figure 3.12: Measured S_{11} shows broadband input matching for various preamplifier gain settings.

equalizer can be adjusted by 9 dB for a total low-frequency gain variation of more than 19 dB.

3.3.2 Time Domain

In order to verify that the AFE can avoid distorting large signals, multilevel signaling tests were performed. Using the 4-PAM transmitter reported in [43], a 1.1 GS/s, 146 mV_{pp}, single-ended, 2^7-1 length 4-PAM signal, shown in Fig. 3.13(a), was applied directly to the receiver input. Note that although the receiver is designed to operate at higher speeds than 1.1 GS/s, test setup limitations prevented the generation of 4-PAM signals faster than this. With no clock applied to the slicers, and with the preamplifier set to minimum gain, these signals should pass through the receiver to arrive at the chip output showing a small amount of signal distortion as can be seen in Fig. 3.13(b). Since these results include the digital back end logic and output drivers, it is likely that the distortion seen in the 4-PAM output is caused by the number of gain stages present after the AFE. As a result, the AFE alone is likely able to accept even larger input signals without introducing signal distortion. There was, however, no test point to permit measured

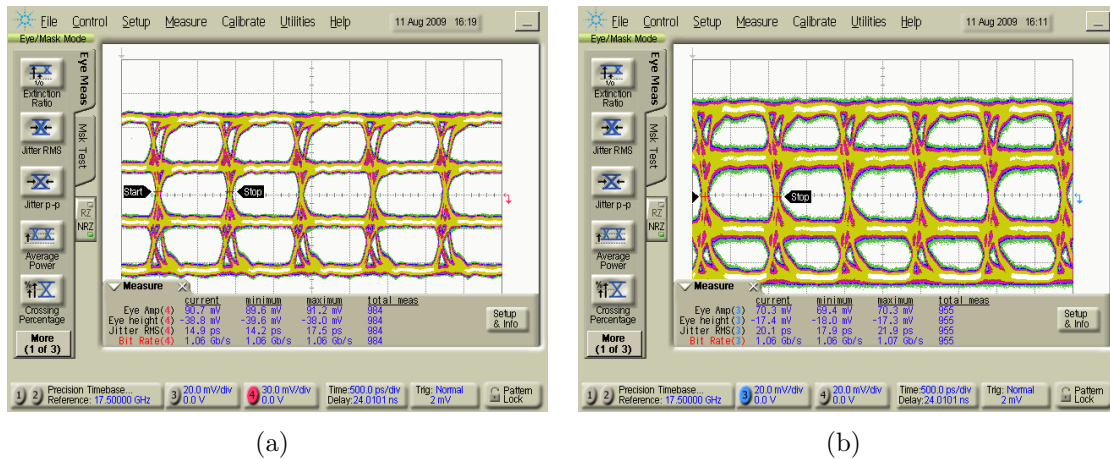


Figure 3.13: (a) 4-PAM, 146 mV_{pp}, 1.1 GS/s test signal applied to the receiver input and (b) corresponding receiver output shows little distortion of the 4-PAM eyes.

confirmation of this on the prototype chip.

To test the receiver's sensitivity a 1-Gb, 2^7-1 length PRBS signal with an amplitude of 8 mV peak-to-peak was applied to the receiver input. The corresponding receiver output was found to be error-free and is displayed in Fig. 3.14(a), indicating that the preamplifier has a sensitivity of at least 8 mV and therefore has a dynamic range of 25 dB. It should be noted however, that bit error rate testing was not performed and the absence of errors was determined by manually examining the input and output bit patterns using Matlab code. To test the overall receiver speed, a 15 Gb/s, 2-PAM PRBS signal of pattern length 2^7-1 was sent to the receiver across a 10-m coaxial cable channel with a loss of 9 dB at 7.5 GHz, as shown previously in Fig. 2.5. This loss was compensated for by the EQ and error-free retimed data at the prototype output is shown in Fig. 3.14(b).

The measured and simulated results of the preamplifier are summarized in Table 6.2, along with a comparison to other recently reported variable-gain preamplifier designs. This helps to illustrate the preamplifier's ability to provide a moderate amount of gain control without sacrificing bandwidth, linearity or input sensitivity. It should be noted that the bandwidth of the preamplifier itself could not be measured directly but simulation results indicate that it should be capable of operating at speeds well beyond 15

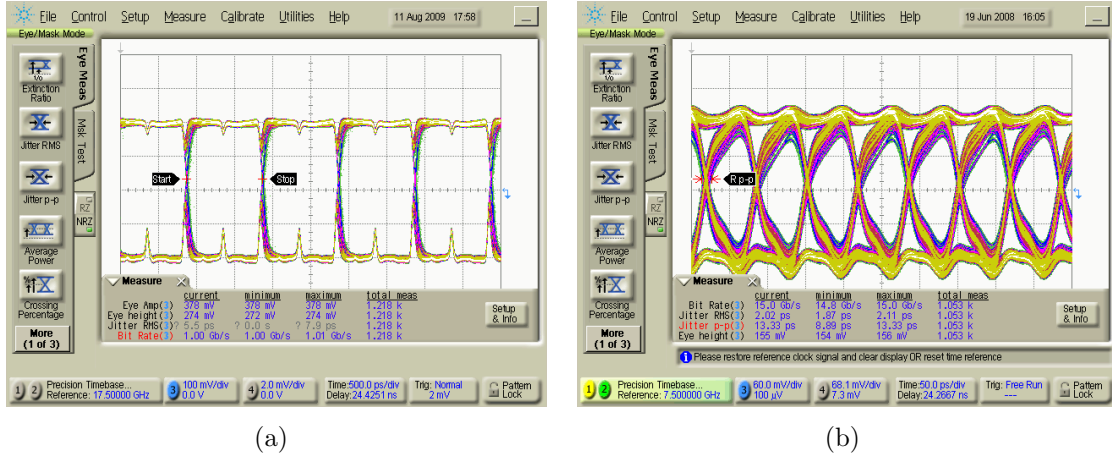


Figure 3.14: Receiver output for PRBS inputs at (a) 1 Gb/s with an 8 mV_{pp} eye opening and (b) 15 Gb/s with a 50 mV eye opening demonstrate receiver sensitivity and speed.

Ref.	Technology	3-dB Bandwidth	THD	Gain Control	Elec. Sensitivity	Preamp Power
[29]	90 nm	7 GHz*	-45 dB*	31 dB*		
[30]	65 nm	5 GHz	-38 dB*	23 dB	-7 dBm	
[33]	90 nm	22 GHz		2 kΩ**	-20 dBm	75 mW
[44]	0.18 μm	3.9-7.6 GHz		52 dBΩ**	-19 dBm	34 mW
This Work	65 nm	30-36 GHz*	-34 dB*	10 dB	-29 dBm	20-68.1 mW

* simulated result

** fixed gain input stage

Table 3.2: Variable-gain broadband preamplifier performance summary and comparison.

Gb/s. Instead it is likely that this limitation is due to the cascade of additional stages in the receiver.

3.4 Adaptation Results

With the proposed preamplifier in place at the input of the AFE, the robustness of the PDF-based adaptation algorithm was tested both through simulation and measurement. Fig. 3.15 shows the eye diagrams of at the output of the equalizer, as simulated in Spectre using schematic level transistors, for a variety of V_{eq} settings. Fig. 3.15 also shows the corresponding PDFs determined by taking the slope of the average value or the output signal as V_{eq} is varied. By selecting the PDF with the largest peak value, the algorithm settles to the correctly equalized eye. These simulations also show that the peak of the

PDF corresponds to the amplitude of the eye opening, allowing for correct optimization of V_{gain} as well.

Fig. 3.16 compares the measured loss of each channel to the equalizer peaking and preamplifier gain settings chosen by the adaptation algorithm. All tests were performed using length 2^7-1 PRBS data at a speed of 4 Gb/s for the PCB tests and 10 Gb/s for the coaxial cable tests (with the exception of the 30 m cable, where the speed was reduced to 5 Gb/s to achieve an open received eye).

In both the PCB and coax cases, the adaptation algorithm responds to the increase in channel losses by increasing equalizer peaking and preamplifier gain. This intuitive result was further verified by examining the receiver's output eye diagram after adaptation had taken place. In all cases the resulting eyes showed error free operation. One example of such an eye is shown in Fig. 3.17 for a 10 Gb/s signal sent across a 10 m coaxial cable.

3.5 Conclusion

In order to avoid non-linear distortions receiver front-ends require gain control, which should begin in the first stage in the receiver chain but must avoid adversely affecting the impedance match with the channel. This chapter has introduced a preamplifier that is suitable for this task with 10 dB of gain control and automated regulation of its output common-mode level. It is implemented as part of an AFE that is able to produce a total of up to 19 dB of gain control.

The preamplifier provides a broadband match with a measured S_{11} of less than -8 dB up to 25 GHz and across all gain settings. Simulation results show a bandwidth of at least 30 GHz and high linearity with a THD of -34 dB. Its fabrication in 65-nm CMOS as part of a complete receiver design was used to verify its ability to avoid non-linear distortions and to operate at speeds of at least 15 Gb/s.

In addition, the ability to generate and automatically adapt the control signals of

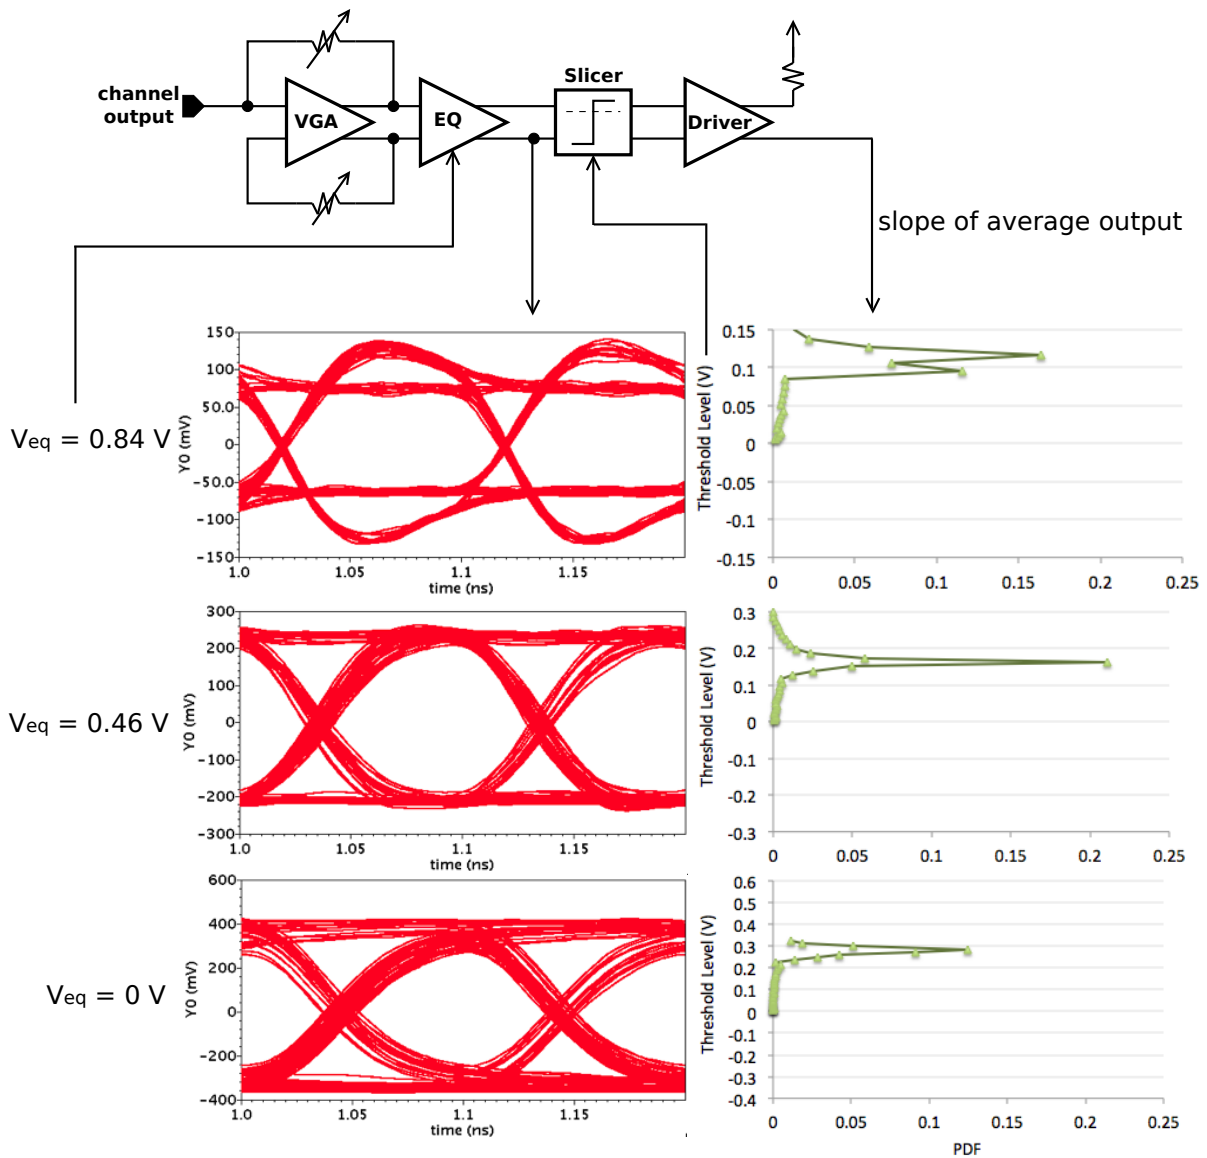


Figure 3.15: Schematic-level simulations of the adaptation scheme in Spectre compare equalizer output eyes to the corresponding PDF and show that by choosing the peak PDF, the algorithm settles to the correctly equalized eye.

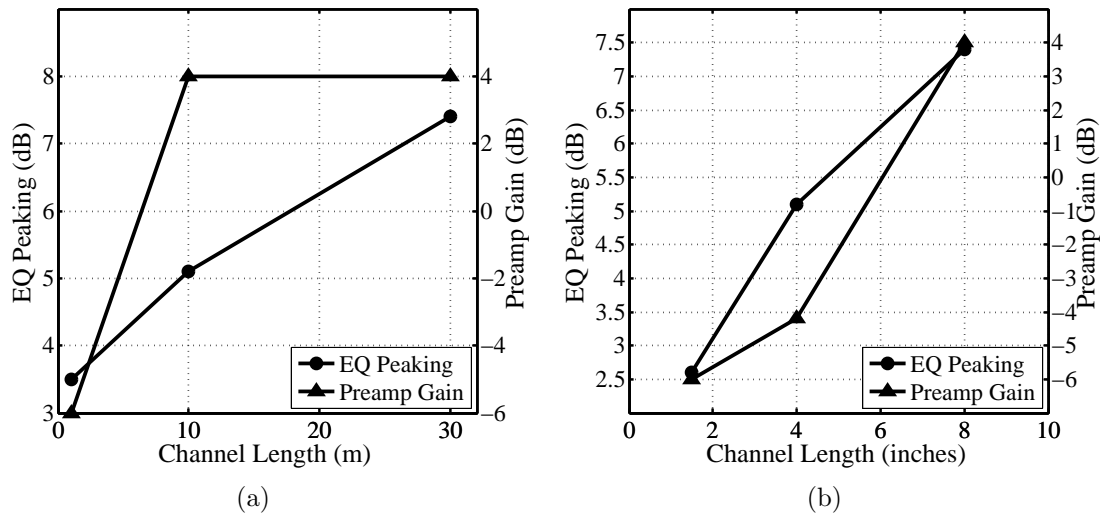


Figure 3.16: Measured results show that the adaptation algorithm automatically increases equalizer peaking and preamplifier gain to compensate for increasing channel losses across (a) coaxial cables and (b) PCB traces.

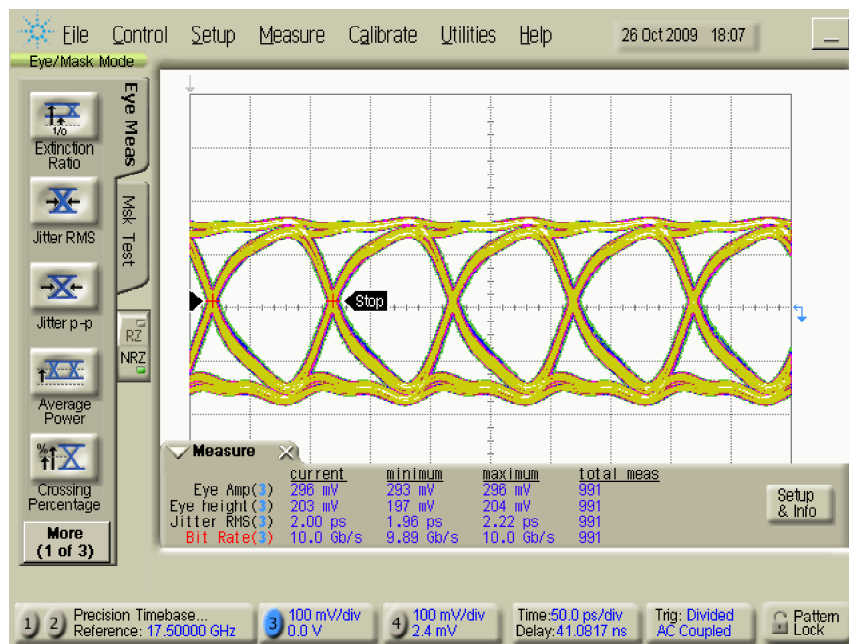


Figure 3.17: Eye diagram of the receiver output after automatic adaptation when receiving 10 Gb/s data sent across a 10 m coaxial cable.

both this preamplifier and an equalizer in a receiver's front-end is essential in order to maintain optimal receiver operation. The adaptation method presented in this work is able to generate these signals by adding only minimal hardware overhead to the receiver. By observing the DC output of a single additional slicer with variable decision threshold, a PDF of the received data is obtained, which contains the information necessary to intelligently adapt the control signals to a variety of channel conditions. By minimizing the spreading of the PDF caused by ISI and maximizing the amplitude of the received signal within the limits of linear operation, the adaptation scheme ensures that the vertical opening of the received eye is optimized.

To demonstrate its effectiveness, the technique was used to adapt the control settings of a binary receiver fabricated in 65-nm CMOS technology. Measured results show that the adaptation scheme operates correctly when used with a variety of channel types and lengths, and at speeds ranging from 2 to 10 Gb/s.

Chapter 4

Injection Locked Oscillators

In order for a communications link to achieve flexibility by adapting its data rate in response to varying demands or operating conditions, it must contain a clock source with the ability to function over a wide range of frequencies. In addition, the ability to transition between not only different operating frequencies but also levels of power consumption can further enhance the usefulness of such a feature.

Although traditionally not well suited for frequency agile applications due to their narrow lock ranges, injection-locked oscillators (ILOs) are becoming increasingly common as frequency dividers [45], multipliers [46], or as alternatives to phase-locked loops (PLLs) [24]. This is due in large part to their small power and area requirements as well as their ability to operate at high speeds [45] and to quickly transition between operating states [21]. Due to their potential for ubiquitous use, a great deal of attention has recently been paid to developing a comprehensive model for the injection-locking behaviour of oscillators. Despite this, an ILO model that is accurate, intuitive and applicable for all types of oscillators under any strength of injection signal has yet to be developed.

This chapter first presents an introduction to injection locked oscillators and then examines the phenomenon of injection locking, through which an oscillator will synchronize its output to an external signal. This phenomenon is best understood through analysis of




Figure 4.1: Oscillation amplitude of an LC tank will decay due to resistive losses unless compensated for by the addition of a negative resistance.

existing ILO models, including the strengths and weaknesses in their abilities to predict measured ILO behaviour.

4.1 Injection Locked Oscillators

If a charged capacitor is connected across an inductor, this charge will flow back and forth between the inductor and capacitor causing the voltage across the capacitor to oscillate at a frequency of [28]

$$\omega_0 = \frac{1}{\sqrt{LC}} \quad (4.1)$$

Due to the parasitic resistances associated with any real inductance and capacitance used to create this “LC tank” circuit, some of this charge will be lost in the form of heat in these resistances in every cycle. As a result, the oscillation amplitude will decay with time unless this charge can be replaced. This charge replacement can be accomplished using the transconductance from an active device, denoted as $-G_m$ in Fig. 4.1.

If charge from a second signal, I_{inj} , which is oscillating at a frequency of ω_{inj} , is injected into this LC tank, as shown in Fig. 4.2, the tank output will be perturbed. If the frequency of the injected signal is close enough to ω_0 and the strength of the injected signal is large enough then the frequency of the tank output will lock to ω_{inj} . Similarly, if the N^{th} harmonic of ω_{inj} is close enough to ω_0 then the tank output will lock to $N(\omega_{inj})$. As a result, any oscillator that is subjected to a signal that satisfies these criteria can be called an injection locked oscillator (ILO).

In electrical oscillators, this phenomenon was observed and exploited in applications




Figure 4.2: Injection of an oscillating signal, I_{inj} , can cause the LC tank output to lock to the frequency, ω_{inj} , of this injected signal.

such as frequency-modulation (FM) receivers as early as 1944 [47]. In this application, it was recognized that an ILO was suitable for this application since it is capable of

- downconverting the received FM signal frequency.
- providing a constant output amplitude regardless of the strength of the received signal.
- rejecting received signals in adjacent frequency bands that are far from ω_0 .

However, no model had yet been presented to account for these observed behaviours or to address ambiguities such as how “strong” an injected signal must be, or how “close” to ω_0 it must be, in order to achieve injection locking. Such models were proposed in subsequent years and are reviewed in the following sections.

4.2 The Frequency Domain Model

In 1946, Robert Adler developed a frequency domain-based model to describe injection locking phenomena observed in LC oscillators [23]. In this model the instantaneous phase difference between the injected and free running oscillator signals, $\Delta\phi(t)$, was defined as

$$\frac{d\Delta\phi(t)}{dt} = \Delta\omega - \frac{I_{inj}}{I_{osc}} \frac{\omega_0}{2Q} \sin(2\pi\Delta\phi(t)) \quad (4.2)$$

where ω_0 is the free running oscillator frequency, $\Delta\omega$ is the difference between ω_0 and the injected signal frequency, Q is the quality factor of the tank, I_{inj} is the injected signal strength, and the strength of the free running oscillations is I_{osc} .

When the injection locked oscillator has settled to a steady state then

$$\frac{d\Delta\phi(t)}{dt} = 0 \quad (4.3)$$

and Equation (4.2) can be simplified to

$$\Delta\omega = \frac{I_{inj}}{I_{osc}} \frac{\omega_0}{2Q} \sin(2\pi\Delta\phi_0) \quad (4.4)$$

From this, the lock range, $\Delta\omega_{max}$, of the injection locked oscillator can be found to be

$$\Delta\omega_{max} = \frac{I_{inj}}{I_{osc}} \frac{\omega_0}{2Q} \quad (4.5)$$

Although this original analysis proved accurate for the case studied in [23], it relied on the following assumptions:

- The injected input and oscillator output are both sinusoidal.
- The oscillator uses an LC tank (since Q is required).
- The strength of the injected signal is much smaller than that of the free running oscillator.

Since these assumptions are not always valid for ILOs, many later publications have since expanded upon this frequency domain model.

As one example of a situation in which these assumptions are invalid, injection locking can be achieved through the use of narrow pulses, in place of sinusoidal injection. This type of injection is often preferable in ILOs used as frequency multipliers where narrow pulses can reduce jitter and duty cycle distortion in the output of a ring oscillator [21] or an LC oscillator [48]. Not only does this type of injection violate the sinusoidal assumption, but it also makes determining the relative strength of the injected signal difficult since it is not clear if I_{inj} in this case should be taken as an average over time or




Figure 4.3: Shorting the tank during narrow injection pulses makes application of the frequency domain model difficult.




Figure 4.4: Equivalence between (a) the physical model of the series resistive losses in the inductance of an LC tank and (b) the parallel resistance used to develop the frequency domain model breaks down for low Q values.

simply as the pulse amplitude. Furthermore, if the injection signal is applied by shorting the tank during injected pulses [48], as shown in Fig. 4.3, then the resulting effect on I_{inj} is unclear. As a result, a modified version of the frequency domain model must be developed and applied in such cases [49].

As for the assumption that the injected signal strength is much smaller than that of the free running oscillator, this condition is often intentionally violated depending on the injection scheme used and the intended application of the ILO. For example, in applications using an ILO to provide deskew in a clock-forwarded link, it was found that stronger injection strengths provide a wider range of achievable phase shifts as well as improved linearity and resolution of the phase steps [50]. Further complicating this situation is the fact that when strong injection is applied to an oscillator with a low Q factor, the accuracy of the frequency domain model deteriorates as the equivalence between the two tank models shown in Fig. 4.4 is lost [25]. This leads to asymmetry in an ILO's lock range, which is a phenomenon that cannot be predicted without significant modifications to the frequency domain model [25].

As a result of these and similar issues, there now exist a wide variety of different




Figure 4.5: Impulses applied to an oscillator have varying impacts on the oscillator output depending on the relative phase at which they are applied.

versions of the frequency domain model, each of which are suitable for some injection locking cases, but not others. Therefore despite the ability to accurately reflect measured results of the lock range, transient phase step response, jitter tracking bandwidth and phase noise [9] for certain cases, this accuracy comes at the cost of complexity or a loss of generality. While attempts have been made to generalize the frequency domain model [51] to make it applicable under all conditions, the results remain complex and difficult to apply during oscillator design. As a result, the frequency domain model is often relegated to the task of explaining measured ILO behaviour instead of predicting it.

4.3 The Impulse Sensitivity Function

The impulse sensitivity function (ISF) was developed in [27] to describe phase noise in oscillators by observing that when small noise current impulses are applied to an oscillator, their impact on the output phase of the oscillator depends on the relative phase at which they are applied. Fig. 4.5 illustrates this concept, showing that for current impulses applied at phases ϕ_1 , ϕ_2 and ϕ_3 the resulting output phase change $\Delta\phi$ is negative, zero or positive, respectively.

It is therefore possible to determine the sensitivity of an oscillator's output phase for all possible applied impulse phases. Typical examples of the resulting impulse sensitivity function, denoted by Γ , are shown in Fig. 4.6. Note that different injection techniques can produce different Γ functions for a single ILO. Hence, in common practice, simulations of the oscillator in the presence of very small impulsive injections are used to obtain Γ .




Figure 4.6: Example impulse sensitivity functions for oscillators with (a) sinusoidal and (b) square wave outputs.

The ISF has traditionally been used to analyze oscillator phase noise and is shown in [27] to have advantages over the Leeson model [52] in its ability to predict $1/f^2$ and $1/f^3$ noise as well as the influence of cyclostationary noise sources. It also offers circuit designers insight into how the shape of the oscillator's output waveform, as well as the method of applying the negative resistance to restore energy to the oscillator, can affect the phase noise performance.

4.3.1 ISF-Based ILO Modeling

Although well suited to modeling phase noise, the ISF model cannot be directly applied to model the injection locking behaviour of an oscillator [53] without some modification. Unlike the noise sources for which the ISF model was developed, the injection waveforms in ILOs are deterministic. They therefore cause the ISF to change significantly, especially under strong injection. For example, straightforward application of the ISF model cannot account for locking an oscillator to a frequency other than its free-running frequency [53]. This is because, according to the ISF model, the phase at the output of an oscillator can be calculated as

$$\phi(t) = \int_0^t \Gamma(\tau)b(\tau)d\tau \quad (4.6)$$

where τ is the time of injection and $b(\tau)$ is an injected signal with a period close, but not equal to that of the free-running oscillator, ω_0 .

Since $\Gamma(\tau)$ has the same frequency as ω_0 , this means that the frequencies of $b(\tau)$ and $\Gamma(\tau)$ will not be equal and that the integral of their product in Equation (4.6) will contain no DC component. This contradicts the known result for an oscillator that is injection locked to $\omega_{lock} \neq \omega_0$. In this case the output phase of the ILO should increase linearly with time relative to the phase of the free-running oscillator, which should be represented by a DC component in the solution to the integral in Equation (4.6).

This idea can be represented graphically by examining the I and Q state variables of a four-stage ring oscillator. These state variables can be observed most easily by examining the output voltages of two unadjacent stages of such a ring oscillator. Simulation results of this oscillator, performed using Spectre and reported in Fig. 4.7(a), show that the injection of an impulse causes a temporary deviation (dotted line) from the steady-state oscillator's trajectory through state-space (solid line), where the magnitude of this deviation is related to the ISF of the oscillator and the strength of the injected impulse. Conventional wisdom dictates that, in order for the ISF to be successfully applied to any future impulses, the transient response of the oscillator must first settle back to its steady-state trajectory. This implies that the oscillator's frequency must remain unchanged and that the ISF is therefore unsuitable for use in the presence of a series of injected impulses designed to lock the ILO to a frequency other than ω_0 , since this would result in shifts in the I and Q states as shown in Fig. 4.7(b) and thereby continuously require new ISFs.

There have been attempts to extend ISF analysis to accommodate injection locking. For example, in [54] it is assumed that with each injected impulse, the ILO's Γ function undergoes a change in phase equivalent to that of the ILO output and any future injected impulses will be applied to the new, phase-shifted ISF as shown in Fig. 4.8. However, the model still fails to account for changes in the amplitude or shape of the ISF that inevitably arise when the oscillator's trajectory through state space deviates significantly from its




Figure 4.7: Spectre simulations of the I and Q states of a four-stage VCO show that (a) an injected impulse causes a perturbation (dotted line) from the steady state (solid line). Repeated impulses (b) can lock the VCO to a different frequency but this requires a new ISF to model the oscillator's new trajectory through state-space.

free-running trajectory, as can result from strong injection. Moreover, the analysis is complex and difficult to generalize.

4.3.2 ISF-Based Modeling for Strong Injection

The challenge of modeling oscillators under strong injection is best illustrated by way of example. In Fig. 4.9 an injected signal is divided into impulses of area $b(\tau)h$. The first impulse produces a shift in oscillator output phase, ϕ , obtained by multiplying the

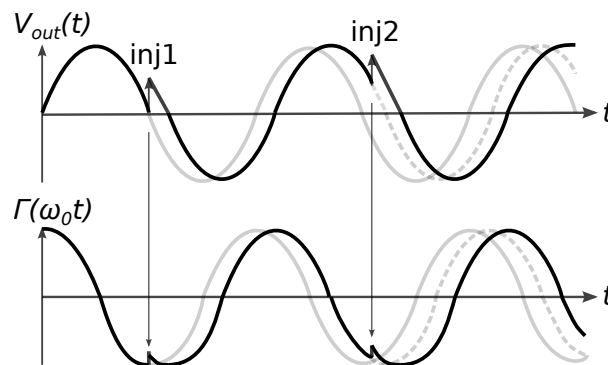


Figure 4.8: Injected impulses result in step changes in the oscillator output, which must be accounted for by introducing step changes to the ISF in order to model injection locking.




Figure 4.9: Dividing an injected signal into impulses that act immediately on the ILO output phase also shift the corresponding $\Gamma(t)$ function, allowing for injection locking to be accurately modeled by the ISF.

pulse area $b(\tau_1)h$ by $\Gamma(\tau_1)$. In [54], this same phase shift is applied to the ISF so that the second impulse is multiplied by $\Gamma(\tau_2 + \phi(\tau_2))$, and so on.

The application of this technique can be not only cumbersome and time consuming, but also inaccurate if the injected signal is large. For example, Fig. 4.10(a) shows how the ISF of a 4-GHz ILO can be determined through transistor-level simulation by injecting an impulse (in this case a 5-mV, 10-ps pulse) into an oscillator at various phases, ϕ , in relation to the oscillator output. Once this ISF has been determined, it is possible to use the method described by Fig. 4.9 to predict the ILO's sensitivity to other injected signals. In Fig. 4.10(b) the amplitude of the applied signal has been increased by a factor of 10. Using the ISF, one would expect the resulting oscillator phase shift to also increase by a factor of 10. Similar predictions can be made for an increase in pulse width, as shown in Fig. 4.10(c). Simulations of the ILO show that these predictions are relatively accurate, so long as the resulting phase shifts remain small. Unfortunately, larger phase shifts are often required in order to implement an ILO with a wide lock range or a fast lock time. Fig. 4.10(d) shows an injected pulse with both large amplitude and width such as would be required in a fast-locking ILO. The ISF prediction method in this case greatly

overestimates the actual phase shifts.

4.4 Summary

The frequency domain model has been shown to accurately model measured ILO behaviour for both LC and ring oscillators, using sinusoidal or pulse train injection, with either weak or strong injection strengths. However, each of these situations requires its own variation of this model. Such variations add significant complexity to the model and the resulting loss of generality limits its usefulness. In addition, in many situations it is difficult to determine the values to be used for the model parameters until measured results are available, thereby limiting its applicability during the design and simulation stages of an ILO.

In contrast, the ISF-based ILO model can predict the phase behaviour of any oscillator so long as the injected signal consists of small impulses. Its inability to be easily used with large injected signals and to be extended to predict other important ILO behaviour such as lock range, lock time and jitter tracking bandwidth have thus far limited its usefulness and it has yet to find widespread application.

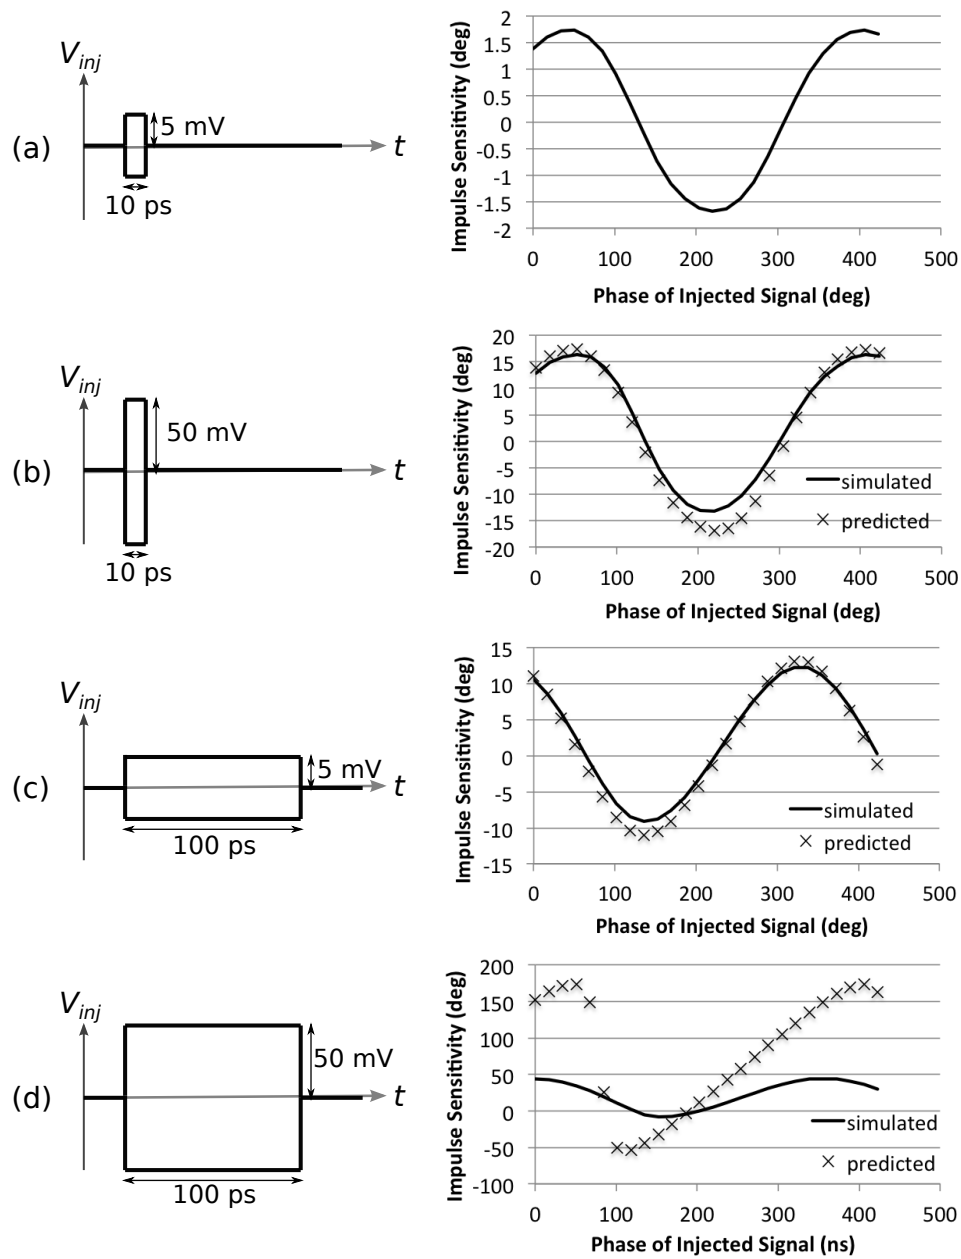


Figure 4.10: Transistor-level simulation of an ILO's response to an injected impulse (a) determines the ISF of the oscillator. This ISF can then be used to predict the ILO's sensitivity to pulses with (b) larger amplitudes or (c) wider pulse widths but fails to accurately predict large, wide pulses (d).

Chapter 5

Phase Transfer ILO Model

As described in the previous chapter, a behavioural ILO model that can be universally applied to any type of oscillator and under any type or strength of injected signal has yet to be presented. This chapter introduces such a model by using the proposed phase transfer characteristic (PTC) of an ILO. The parameters of this model can be extracted from a relatively short set of transient simulations of the oscillator in question. Once extracted, the model can be used to infer a great deal of information about the oscillator, such as lock range, lock time, input and output phase relationships and jitter tracking bandwidth, using only behavioural simulations.

This chapter will show that the use of this model presents a significant saving in time and computing resources when compared with determining this information directly through traditional simulations. Moreover, the ILO model can be incorporated into larger system-level behavioural models, such as phase-locked loops and clock distribution networks. The utility of this new model is then demonstrated by using it to develop an ILO with a wide lock range.

5.1 The Phase Transfer Characteristic

Instead of simulating the ILO under impulsive injection and feeding the result into complex, and in some cases inaccurate, expressions, we instead simulate the ILO's phase transient with the actual injected pulse shape being studied. Such simulations more directly provide an intuitive understanding of how ILO design goals can be translated to circuit topology. Moreover, the resulting model is accurate even under strong injection, and is readily incorporated into behavioural simulations.

To define the model we begin by defining the relationship between the injected signal and the ILO output. First, we assume that when the ILO is locked by some injected signal to an angular frequency ω_{lock} , its output is

$$V_{lock}(t) = f(\omega_{lock}t + \theta(t)) \quad (5.1)$$

where f is some periodic function with period 2π describing the oscillating waveshape (e.g. square, sinusoidal, etc.) and $\theta(t)$ represents the phase of the signal.¹ Similarly, the injected signal can be represented by

$$V_{inj}(t) = b\left(\frac{\omega_{lock}}{N}t + \theta_{inj}(t)\right) \quad (5.2)$$

where b is some periodic function (i.e. a sinusoid or a pulse train) with frequency 2π , $\theta_{inj}(t)$ is the phase of the injected signal and N is an integer that represents some multiplication factor between the injected and output signal frequencies.

In the case where the injected signal is no longer present, the ILO will free-run at a frequency $\omega_0 = \omega_{lock} - \Delta\omega$, and its output therefore becomes

$$V_{lock}(t) = f((\omega_{lock} - \Delta\omega)t + \theta(t)). \quad (5.3)$$

¹Note that voltage state variables $V(t)$ are used in this work, any of the relevant signals may be branch currents instead of voltages.




Figure 5.1: Definitions of the ILO input and output signals that will be used to develop the ILO model in the remainder of this chapter.

Here $\Delta\omega$ is the difference between the frequency of the locked ILO and its free-running frequency,

$$\Delta\omega = \omega_{lock} - \omega_0. \quad (5.4)$$

These definitions are illustrated in Fig. 5.1.

In this work, the ILO's response to one full period of the injected signal, $V_{inj}(t)$, is simulated. If the ILO includes any peripheral circuitry such as narrow pulse generators to condition the injected signal [48] then this can be included in the simulation to ensure that the effects of this circuitry are accurately captured. Each period of this injected signal changes the phase of the ILO's output, $\theta(t)$, by an angle, P . The procedure used to determine P through transistor-level simulation in Spectre is shown in Fig. 5.2 where two copies of the ILO (including any peripheral circuitry) are simulated over a small number of cycles of the output clock. By injecting one period of the intended signal into the ILO and comparing the times of the resulting zero crossings to those of the unperturbed ILO, the output phase change can be determined.

This phase change depends upon $\phi(t)$, defined as the phase of the ILO output signal $\theta(t)$ subtracted from the phase of the injected signal, $\theta_{inj}(t)$, such that

$$\phi(t) = \theta_{inj}(t) - \theta(t). \quad (5.5)$$

Hence, we define the phase transfer characteristic (PTC), $P(\phi)$, as the ILO's phase change for each injection of one period of $V_{inj}(t)$ as a function of the relative phase of

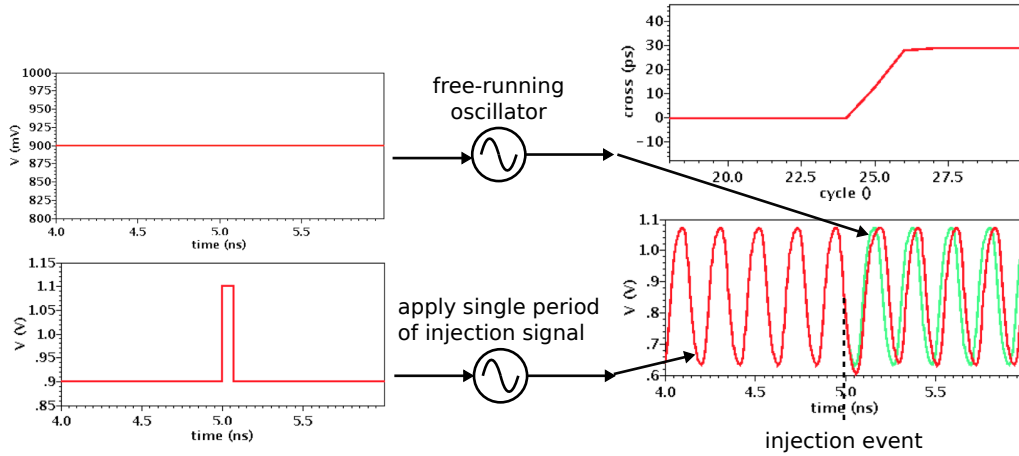


Figure 5.2: By comparing the zero crossing times of the ILO output to that of an unperturbed copy, the phase change created by one period of the injected signal can be determined through transistor-level simulation in Spectre.

this injection, $\phi(t)$. The PTC is readily extracted from a series of transient simulations, as demonstrated by Fig. 5.3 where two samples of $P(\phi)$ are determined by applying one period of $V_{inj}(t)$ at phases ϕ_1 and ϕ_2 .

Since the PTC is specific to the injected signal, $V_{inj}(t)$, it can have a wide variety of possible shapes depending on the amplitude, shape and frequency of the injected signal and the injection scheme used. While this technique means that PTC simulations must be redone if the injected pulse shape is changed, these simulations can be run in a short time and the results are accurate and provide insight. The time required to simulate the PTC will be analyzed and compared to direct spice-level simulation of the ILO characteristics later in Section 5.2.5.

5.2 PTC-Based ILO Modeling

With the PTC in hand, a behavioural model for the ILO is formed under the assumption that injection of the waveform b at a relative phase ϕ causes an immediate change in

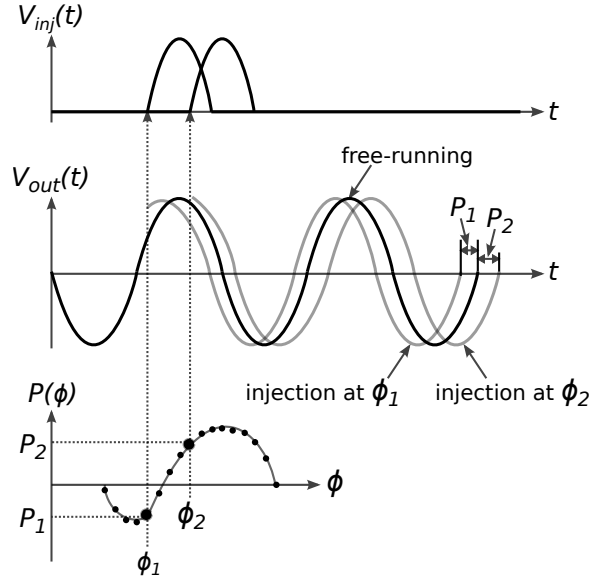


Figure 5.3: The PTC, $P(\phi)$, is determined through simulation by applying one period of the injected signal, $V_{inj}(t)$, at different phases, ϕ_n , relative to the oscillator's output signal, V_{out} , and observing the resulting change in output phase, P .

the ILO's output phase equal to $P(\phi)$.² Moreover, it is assumed that between injection events, the ILO operates at its free-running frequency, ω_0 , causing its phase relative to the lock frequency to drift by $2\pi\Delta\omega/\omega_0$ radians each period.

If we treat each period of the injected signal as a discrete event then we are interested in the phase difference at the start of the k^{th} injection and Equation (5.5) becomes

$$\phi_k = \theta_{inj,k} - \theta_k. \quad (5.6)$$

While an ILO is locking, the difference between the injected signal phase and ILO output phase evolves along the sequence $\phi_1, \phi_2, \dots, \phi_k$. This means that the phase shift introduced by injection event k is $P(\phi_k)$ such that

$$\phi_{k+1} = \phi_k - P(\phi_k). \quad (5.7)$$

²This is an approximation since, in fact, it may generally take some time for the ILO's output phase to react to the injected input. However, the accuracy of this approximation is borne out by later comparison of the model with transistor-level simulations and measurements.

The negative sign is included because an increase in ILO output phase results in future injection events being applied earlier relative to the ILO output.

In the event that $\omega_{lock} \neq \omega_0$ (i.e. $\Delta\omega \neq 0$) an additional phase shift of $-2\pi N\Delta\omega/\omega_0$ is added to Equation (5.7). The resulting expression for the phase difference between the injected signal and the ILO output becomes

$$\phi_{k+1} = \phi_k - P(\phi_k) - \frac{2\pi N\Delta\omega}{\omega_0} \quad (5.8)$$

Finally, since any perturbations of the phase of the injected signal phase (i.e. cycle-to-cycle jitter) can be represented as

$$\Delta\theta_{inj,k} = \theta_{inj,k+1} - \theta_{inj,k} \quad (5.9)$$

this should be included in the model and Equation (5.8) becomes

$$\phi_{k+1} = \phi_k - P(\phi_k) - \frac{2\pi N\Delta\omega}{\omega_0} + \Delta\theta_{inj,k} \quad (5.10)$$

One may also wish to consider the ILO's behaviour in terms of an absolute phase reference in order to model external phase perturbations and to make this model applicable in other, larger systems. This can be done by substituting Equation (5.6) into Equation (5.10), resulting in

$$\theta_{k+1} = \theta_k + P(\phi_k) + \frac{2\pi N\Delta\omega}{\omega_0} \quad (5.11)$$

This relationship incorporates the nonlinear PTC, $P(\phi_k)$, and can be represented by the system drawn in Fig. 5.4. The absolute phase reference of this model means that it can easily be implemented to model ILO behaviour in system level simulations.

Equations (5.11) and (5.6) and Fig. 5.4, comprise a general nonlinear behavioural model of an ILO. The nonlinear PTC function, $P()$, may be extracted from a relatively quick series of transient simulations of the ILO to be modeled, as described above. In




Figure 5.4: Model representing nonlinear the phase relationship between the injected signal and the ILO output.

the next section, it will be shown how $P()$ may also be extracted from measurements of an ILO. The following sections will show how the model may be used to very quickly and accurately find the phase relationship, lock range, lock time and tracking bandwidth of an ILO. Each of these ILO performance metrics would otherwise require extensive transistor-level simulations; hence, the model greatly accelerates design iterations, affording the designer insight. The model may also be integrated into larger behavioural system-level models such as phase-locked loops and clock distribution networks.

5.2.1 Steady-State Phase Shift

When locked, the oscillator and injected pulses will settle to some steady-state phase relationship, ϕ_{ss} , where each injected period causes a phase shift $P(\phi_{ss})$ that is just sufficient to cancel the phase drift resulting from $\Delta\omega$. This concept is illustrated by the Spectre simulation results shown in Fig. 5.5 for a ring oscillator with $\omega_0 = 4.7$ GHz, which is being injection locked to $\omega_{lock} = 4.72$ GHz by pulse injection that occurs at a frequency of 1.18 GHz (i.e. $N = 4$). In part (a) the ILO output signal is compared to an ideal reference signal oscillating at 4.7 GHz. Due to the 20 MHz difference between ω_{lock} and ω_0 the phase difference between these signals, as measured by the difference between their zero crossing times as shown in part (b), shrinks during cycles 74, 75 and 76. When an injection event then occurs during cycle 77 the resulting phase change, given by $P(\phi_{ss})$, is sufficient to offset the phase difference that accumulated during the previous four cycles. This allows the oscillator to settle to the steady-state phase relationship shown in part

(c).

In order for an ILO, with free-running frequency of ω_0 to lock to ω_{lock} , the phase change produced by each period of the injected signal at steady-state, $P(\phi_{ss})$, must be sufficient to eliminate the phase drift accumulated over N cycles of the oscillator output such that

$$P(\phi_{ss}) = -\frac{2\pi N \Delta\omega}{\omega_0}. \quad (5.12)$$

Since N , $\Delta\omega$, ω_0 and ϕ_{ss} of a physical oscillator can all be directly observed, Equation 5.12 therefore provides a way to determine the PTC of a fabricated ILO and compare it to simulated results. This will be illustrated in the following chapter.

Equation (5.12) shows that the steady-state phase relationship, ϕ_{ss} , between the injected signal and the ILO output is determined by the frequency difference, $\Delta\omega$. This observation is intuitive since the steady-state phase relationship between the injected and output signals, ϕ_{ss} , of an ILO has previously been exploited in applications such as clock deskew, where ILO output phase can be adjusted by tuning the free-running frequency of a VCO [24]. Fig. 5.6 illustrates this concept by showing ϕ_{ss} for a variety of injection frequencies. When $\omega_{lock} = \omega_0$ (i.e. $\Delta\omega = 0$) each period of the injected signal has no need to influence the oscillator's output and therefore settle to a steady-state relationship where, according to the PTC, they will have no effect on the output phase (i.e. $P(\phi) = 0$). When $\omega_{lock} < \omega_0$ each injected pulse must decrease the oscillation frequency, meaning that the pulses settle to a steady-state relationship where they will each create a positive change in oscillator phase, given in this example by P_+ . These steady-state relationships are reached, after some settling time, regardless of the phase at which the injected pulses begin.

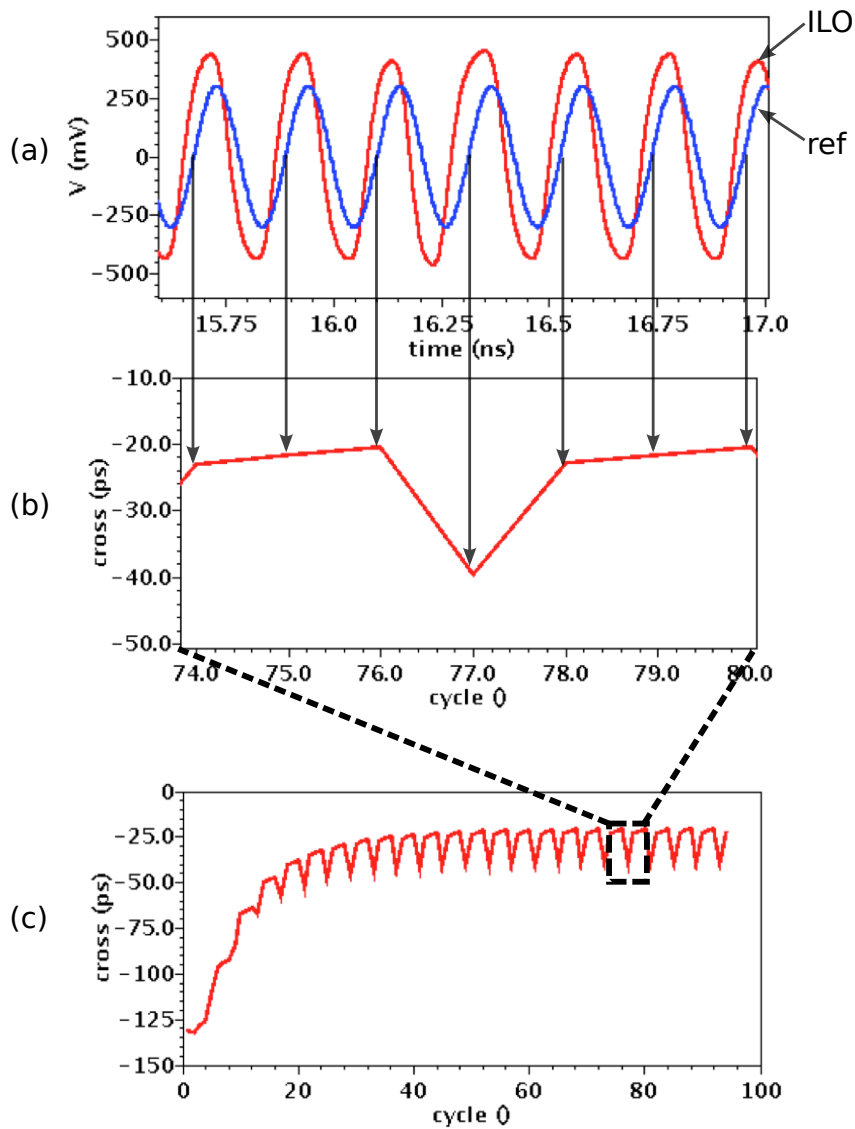


Figure 5.5: Transistor-level simulation results (a) comparing ILO output to an ideal reference signal at ω_{lock} shows that (b) the phase change produced by an injection event is sufficient to cancel the phase drift resulting from the difference between ω_{lock} and ω_0 . This allows the ILO to settle to the steady-state phase behaviour plotted in (c).




Figure 5.6: Steady-state phase relationships are determined by the difference between ω_{lock} and ω_0 .

5.2.2 Lock Range

The analysis of an ILO's lock range under subharmonic injection is of particular interest [24], [21]. Fig. 5.6 shows that a natural extension of the steady-state phase shift modeling is that the lock range of an ILO can be determined directly from its PTC since the ILO can only successfully lock to an injected signal that produces a large enough phase change in the oscillator output to compensate for the difference in their frequencies. In other words, the maximum value of $\Delta\omega$, which we define to be $\Delta\omega_{high}$, can be found using Equation (5.12) to be

$$\Delta\omega_{high} = -\frac{\omega_0 P_{min}}{2\pi N}. \quad (5.13)$$

Similarly, the minimum value of $\Delta\omega$, which we define to be $\Delta\omega_{low}$, is

$$\Delta\omega_{low} = -\frac{\omega_0 P_{max}}{2\pi N}. \quad (5.14)$$

The total output referred lock range, $\Delta\omega_{max}$, is therefore

$$\Delta\omega_{max} = \Delta\omega_{high} - \Delta\omega_{low} \quad (5.15)$$

$$= \frac{\omega_0 (P_{max} - P_{min})}{2\pi N} \quad (5.16)$$

$$= \frac{\omega_0 P_{pp}}{2\pi N} \quad (5.17)$$

where $P_{pp} = P_{max} - P_{min}$ has been defined as the peak-to-peak value of the PTC. Simulation results of an example oscillator topology comparing these lock range equations to SPICE-level simulations are presented in Section 5.3.1.

Treating the maximum and minimum PTC values separately, as in Equations (5.13) and (5.14), identifies cases where the lock range is not centered equally about the free-running frequency. This effect can be present in ILOs for a variety of reasons, especially during strong injection and, although it has previously been reported in [25], it is often ignored in ILO models. Furthermore, it should be noted that the lock range calculation

given by Equation (5.17) does not require that the circuit designer determine an effective Q , injection strength, or any other oscillator parameter. Instead it relies only on the PTC, which can be efficiently extracted from simulation.

5.2.3 Lock Time

When an injected signal is applied, the time that it takes for an ILO's output to settle to a steady-state phase relationship with this injection is known as the lock time. This transient relationship can be useful in determining the ILO's jitter tracking capabilities [9] and can also be important in systems that require fast locking, such as frequency hopping [55], burst mode [56] or fast power-on applications [21].

The PTC allows us to predict lock time variations that are not obvious in the frequency-domain ILO model. Although it has been shown that the frequency domain model can be manipulated to predict these lock time variations [57], the complexity introduced by such manipulations can make this approach unattractive. Whereas it is usually suggested that lock time depends only upon $\Delta\omega$ and injection strength, the PTC model indicates that there is also a strong dependence on the initial phase relationship, ϕ_0 , between the ILO output and the injected signal.

For the analysis of lock time we assume that no phase perturbations are introduced by the injected signal ($\Delta\theta_{inj,k} = 0$). In this case, the model given by Equation (5.10) shows that the ILO phase will settle to its steady-state condition, ϕ_{ss} , when

$$\phi_{k+1} = \phi_k = \phi_{ss} \quad (5.18)$$

and therefore

$$P(\phi_k) = P(\phi_{ss}) = -\frac{2\pi N \Delta\omega}{\omega_0} \quad (5.19)$$

which corresponds to Equation (5.12) as determined previously.

This means that an injected signal that begins at a phase that is far from the desired




Figure 5.7: An ILO's lock time depends on the initial phase difference, $\Delta\phi$, between the injected signal ϕ_0 and the steady-state phase difference, ϕ_{ss} , required by the frequency of the injected signal.

steady-state relationship, ϕ_{ss} , will require more injection events, and therefore a longer time, to reach ϕ_{ss} . Fig. 5.7 demonstrates this relationship for an example case where two identical injected signals are applied individually to an ILO at initial phases ϕ_{01} and ϕ_{02} . Since ϕ_{01} is much closer to ϕ_{ss} than ϕ_{02} , therefore $\Delta\phi_1 < \Delta\phi_2$, resulting in a shorter lock time for the injected signal that begins at ϕ_{01} . Note that although $P(\phi_u) = P(\phi_{ss})$, the steady-state ILO output phase cannot settle to this point. A small deviation to the left of ϕ_u , resulting from noise or a slight frequency difference between ω_{lock} and ω_0 , will produce a small positive phase shift, which will then shift the phase difference further to the left of ϕ_u , in turn producing a larger positive phase shift, and so on until ϕ_{ss} is reached. A similar effect occurs in the opposite direction if the shift occurs to the right of ϕ_u . Due to the small steps that begin this settling, the increase in lock time that occurs when $\phi_0 \approx \phi_u$ is significant.

To illustrate this effect, Fig. 5.8(a) shows the PTC of a 4-stage ring oscillator obtained using Spectre (SPICE) simulations of a 400-mV_{pp} (differential) injected pulse with a width of 70 ps. When this injected signal is at a frequency close to ω_0 then ϕ_{ss} is where $P(\phi_{ss}) = 0$ on the rising edge of the PTC, as indicated. If the injected signal begins its injection at a phase that is close to ϕ_{ss} it will therefore settle quickly, following a simple, first order exponential settling step response. Indeed, the time constant of this exponential settling is expressed in the following section as a function of the jitter tracking bandwidth, f_{TB} , in Equation (5.25), in accordance with other linear modeling




Figure 5.8: Spectre simulations of the (a) PTC and (b) transient phase response of a 4-stage ring ILO. When injection begins far from ϕ_{ss} at ϕ_{01} the lock time is significantly longer than when it begins at ϕ_{02} .

approaches. If the injected signal begins farther from ϕ_{ss} , especially if it begins near the unstable operating point ϕ_u , it will require a much longer lock time. This effect is captured by the PTC model as shown in Fig. 5.8(b) where 2 identical signals are injected into the 4-stage ring oscillator but beginning at initial phases ϕ_{01} and ϕ_{02} , respectively. This can lead to large variations in the lock time of an ILO for a given injection frequency, but is not captured by traditional analyses based on frequency domain modeling.

A more complete picture of the lock time of an ILO as a function of its initial phase is shown in Fig. 5.9 for both SPICE simulations and as predicted using the PTC model. In these simulations the lock time is defined as the time taken for the oscillator's output phase to settle to within 1° of its steady-state phase. Although the lock time reaches a maximum value near 12 ns, it should be noted that there is no fundamental limit to this and it is possible to observe very long settling times in an ideal, noise-free simulation environment. In practice noise will push the oscillator phase away from ϕ_u , thereby causing it to lock more quickly.




Figure 5.9: Lock time varies greatly depending on the phase at which the injected signal begins. This effect is seen in both SPICE-level simulation simulation and the PTC-based lock time model.

5.2.4 Tracking Bandwidth

Although a strength of the PTC-based model is that it captures the nonlinear phase response of the ILO during large phase transients, it can also be used to find linear performance metrics in the presence of small phase deviations such as phase tracking bandwidth. Specifically, consider an ILO that has reached steady-state at a lock point with a relative phase shift ϕ_{ss} defined by Equation 5.12. Small perturbations around this lock point due to phase changes (i.e. jitter) in the injected signal, $\theta_{inj,k}$, result in restoring phase shifts that are proportional to the phase error. The constant of proportionality is the slope, m , of the PTC around ϕ_{ss} . Hence, under small phase perturbations, a first-order phase tracking model, JTF , may be applied,

$$JTF = \frac{1}{1 + \frac{\omega_j}{\omega_{TB}}} \quad (5.20)$$

where ω_j is the frequency of the jitter and the 3-dB tracking bandwidth is given by ω_{TB} .

When the phase at the input of the ILO is perturbed by an amount, $\theta_{inj,k}$, then the

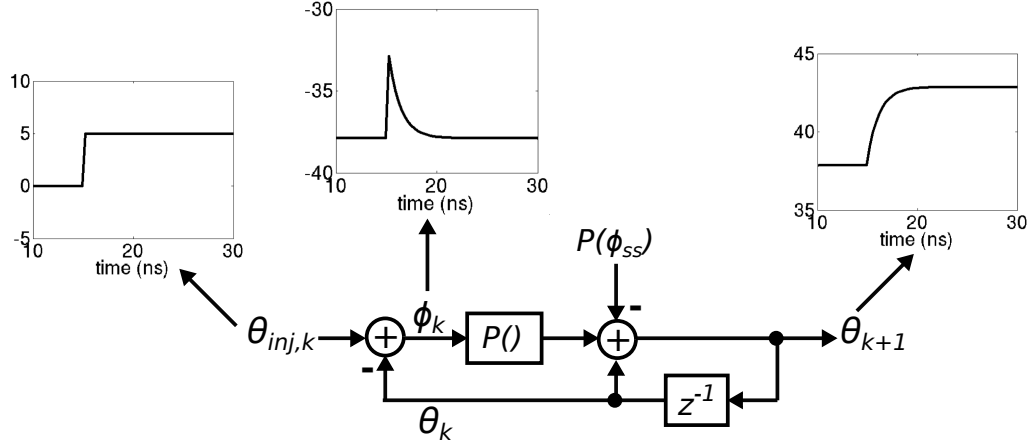


Figure 5.10: The jitter tracking bandwidth of an ILO can be determined by applying a step change to the phase of the injected signal and observing the resulting change in the output phase. The displayed results are from Simulink simulations of the proposed model, performed using the PTC of the ILO as determined from Spectre simulations.

output phase of the ILO can be determined using Equation (5.11) as

$$\theta_{k+1} = \theta_k + P(\theta_{inj,k} - \theta_k) - P(\phi_{ss}) \quad (5.21)$$

This equation shows that the rate of change of the output phase of the ILO in response to $\theta_{inj,k}$ is determined by P . To determine the jitter tracking bandwidth of the system, we apply a small step change to $\theta_{inj,k}$ and observe the system response as illustrated in Fig. 5.10. When this step is applied, the phase difference between the injected signal and the ILO output, ϕ_k , jumps by the value of the applied step, which then causes the output phase θ_{k+1} to change by $P(\phi_k)$, according to Equation (5.21).

Since the output phase follows a first order exponential settling, as shown in the previous section, the time constant τ_{TB} of this response can be found from the slope of θ_{k+1} . This can be determined by taking the derivative of Equation (5.21), resulting in

$$\frac{d\theta_{k+1}}{dk} = \frac{dP(\phi_k)}{dk}. \quad (5.22)$$

If we assume that for small perturbations in θ_{inj} the slope of P is a constant given by m




Figure 5.11: Comparison with direct transistor-level simulation using Spectre shows that the 3-dB jitter tracking bandwidth can be accurately predicted over a range of injected frequencies using the PTC model.

then

$$m = \left. \frac{dP(\phi_k)}{dk} \right|_{\phi_k = \phi_{ss}} \quad (5.23)$$

and the time constant, τ_{TB} , of the settling behaviour can be defined as

$$\tau_{TB} = \frac{1}{mf_{inj}} \quad (5.24)$$

where f_{inj} is included to convert τ_{TB} from injection cycles to seconds. This then means that the jitter tracking bandwidth of the first-order phase tracking model is given by

$$f_{TB} = \frac{1}{2\pi\tau_{TB}} = \frac{mf_{inj}}{2\pi} \quad (5.25)$$

where the injected frequency, f_{inj} , is related to Δf and therefore ϕ_{ss} through Equation (5.12).

The accuracy of this model is illustrated by Fig. 5.11 where the f_{TB} of the 4-stage ring oscillator discussed in the previous section is calculated using Equation (5.25) and compared to SPICE simulations of small phase perturbations over a range of injected frequencies.

5.2.5 Simulation Time

The time required for any spice-level transient simulation varies depending on a number of factors including the processing power of the machine used, the size of the circuit being simulated, the time step used by the simulator and the level of accuracy required by the designer. Nonetheless, by making all of these factors as constant as possible, this section attempts to quantify the reduction in simulation time achieved by using the PTC in place of traditional ILO simulations.

To this end, simulations of a four-stage, 4.7-GHz ring oscillator with narrow pulse injection into each stage of the oscillator [21] were performed using Cadence Spectre. With the accuracy set to “conservative” for small step sizes, transient simulations were performed to determine the PTC, as described previously. By using initial conditions to help the free-running oscillations begin, the ILO settles to its steady-state in approximately 1 ns and the injected signal was applied 2 ns after beginning the transient simulation. By delaying the application of this pulse from 2 ns to 2.22 ns (in increments of 10 ps), the $P(\phi)$ of the ILO was determined for a range of ϕ spanning 360° of the 213-ps output signal period. For this comparison, the peak-to-peak value of the PTC was then translated to a lock range using Equation 5.17. Total simulation time required to determine the lock range in this way was measured at 445 seconds.

In contrast, it is possible to determine the ILO characteristics directly using similar simulations. For example, the ability to lock to an injected signal can be determined by applying the injected signal to the ILO for an extended period of time and observing the output phase of the oscillator (relative to a signal oscillating at the ideal output frequency) to see if the phase settles to some steady state value. By repeating this simulation over a range of injected frequencies, the lock range of the ILO can be determined. Unfortunately, in order to obtain a reasonable level of accuracy in the lock range when it is determined in this way, a large number of injected frequencies must be attempted. Further complicating the situation is the fact that each transient must be run for a large number of periods of




Figure 5.12: Spectre transient simulations of an ILO settle to a constant phase if the ILO is correctly locked (dotted line). If the injected signal is beyond the ILO’s lock range then this can be identified by slipping of the output phase (solid line), which may not become apparent until the simulation has been run for many output clock cycles.

the output clock since the phase slipping that identifies an unlocked ILO may take many cycles to present itself for frequencies close to the edge of the lock range. This case is illustrated by the transistor-level simulation results in Fig. 5.12 where the phase of the 4.86-GHz injected signal (solid line) appears to settle to a steady state but is revealed to slip only after the simulation has been running for over 60 cycles of the output clock.

Therefore, in order to determine the lock range of the 4.7-GHz oscillator discussed above, 20-ns transient simulations were run for frequencies from 3.9 GHz to 5.1 GHz, in 10 MHz steps. Although the actual lock range of the ILO is from 4.15 GHz to 4.85 GHz, this information is not known before the simulations are performed and wider range of frequencies must therefore be attempted. A resolution of 10 MHz was found to provide comparable accuracy to that obtained using the PTC, as will be shown later in Table 5.2.

Due to the length of each transient, and the fact that 121 frequencies need to be simulated to cover 3.9 to 5.1 GHz in 10 MHz steps, this simulation takes considerably longer than the PTC simulation. This is apparent from the comparison shown in Table 5.1. Since the PTC can be used to quickly determine not only the lock range of an ILO but also its steady-state phase, lock time and jitter tracking bandwidth, as was shown in the previous sections, the advantage of using the PTC model is apparent.

simulation	simulation time
PTC lock range	445 seconds
transient lock range	34180 seconds

Table 5.1: Comparison of simulation time required to determine the lock range of an ILO to that required to determine the PTC, which can be used to determine lock range in addition to other ILO parameters.

5.3 Wide Lock Range ILO Design

This section presents a design example applying the proposed model to a multiplying ILO (MILO) that generates a 4-GHz output signal from a 1-GHz reference clock. In order to demonstrate the usefulness of the PTC model, the MILO is designed to have a very wide lock range, which is difficult to model using other methods. Wide lock range is typically difficult to achieve for ILOs, with reported lock ranges commonly less than 5% of the free-running frequency [45], [21]. Hence, an unconventional circuit architecture is required that does not fit conventional ILO models, but the PTC-based model can be applied easily and is shown to accurately predict performance.

A ring oscillator topology was chosen as it generally provides wider lock range than LC-based ILOs [58]. The frequency domain model presented in [24] states that the lock range of a ring oscillator-based ILO is

$$\Delta\omega_{max} = \frac{2\omega_0}{n \sin \frac{2\pi}{n}} \frac{K}{1 - K^2} \quad (5.26)$$

where n is the number of stages in the ring and K is the relative injection strength given as I_{inj}/I_{osc} . Although this model indicates that the number of oscillator stages should be decreased and that the injection strength should be increased in order to maximize $\Delta\omega_{max}$, the model provides very little insight into what the injected signal should look like and how it should be applied to the MILO. Further complicating the application of the frequency domain model is the fact that Equation (5.26) must be modified once the loosely defined boundary between weak and strong injection is crossed.




Figure 5.13: One stage of the four-stage CML injection locked ring oscillator. Applying the injected signal to a secondary differential pair provides a strong injection strength.

The PTC model, specifically Equation (5.17), indicates that the lock range, Δf_{max} , can be increased by maximizing the peak-to-peak phase transfer characteristic, P_{pp} . Since P_{pp} can be efficiently determined through simulation, the lock range of different MILO topologies can be quickly evaluated and compared. In this design a lock range of 1 GHz, or 25% of the 4-GHz f_0 was targeted, which translates to a target P_{pp} of 360° .

To serve as a starting point in the design, a four-stage CML ring oscillator was simulated in a standard 65-nm GP CMOS process. To create a strong injected signal strength, V_{inj} was applied to a secondary input differential pair with the drain nodes connected to those of the original CML stage as shown in Fig. 5.13. Transistor sizes in μm are shown, with all gate lengths implemented as minimum sizes. The tail current of the main pair was set to $744 \mu\text{A}$, while that of the injection pair was chosen to be $1/5$ of this in order to ensure that the ILO continues to oscillate when there is no injected signal present. Varactor load capacitances are used to tune the MILO's free-running frequency if necessary.

The PTC of the four-stage MILO is then determined by simulation using injected pulses with an amplitude of 300 mV_{pp} (differential) and a width equal to approximately half of one period of a 4-GHz clock signal. These pulses were applied to the first stage of the oscillator, as shown in Fig. 5.14. In order to create a realistic pulse shape in the simulation environment, an ideal pulse is first applied to a CML differential pair before it is applied to the MILO. The secondary differential pair shown in Fig. 5.13 is included




Figure 5.14: The PTC is determined by measuring the output phase change created when a single pulse is applied to the MILO at different phases relative to the oscillator’s output signal.

in each stage of the oscillator to provide a consistent load at the output of each stage. Where these injection pairs are unused, their gates have been grounded.

By applying this pulse at various times spanning one period of the clock signal and observing the resulting change in the output phase of the MILO, the PTC was determined and is plotted as the “1 inj” curve in Fig. 5.16. It exhibits a P_{pp} of 26^0 , which corresponds to a lock range of 54 MHz and highlights the difficulty of achieving a wide lock range for an ILO. Although the strength of the injected signal has been maximized relative to the available headroom in the 65-nm CMOS process, other strategies that attempt to increase the MILO’s sensitivity to injected signals are required in order to increase lock range.

5.3.1 Injection Point Selection

Injection into multiple locations of an oscillator has been shown to increase the lock range of injection-locked frequency dividers by applying the injected signal to the tail currents of two [26] or three different stages of an n -stage ring oscillator [58]. In both cases off-chip controls were used to modify the phase relationship between the injected signals, demonstrating that injected signals should be applied with successive phase delays of π/n in order to achieve the widest lock range. In other words, the injected signal should experience the same delay as is created by one oscillator stage before being injected into the subsequent stage.




Figure 5.15: Injection into multiple ILO locations can increase locking range if the injected signal experiences a delay that is equal to that created by each stage of the ring oscillator.

This means that the required phase shift in the injected signal can be easily created on-chip by passing the injecting signal through delay elements that are identical to those that make up the ring oscillator. This concept is illustrated in Fig. 5.15 where injection delay stages I_1 to I_4 are identical to ring oscillator stages R_1 to R_4 with their unused injection input ports grounded. Small buffer stages are also added to the output of each stage of the ring oscillator in order to ensure that the load seen at the outputs of R_1 to R_4 are the same as those seen at the output of stages I_1 to I_4 .

The addition of each new injection site increases the ILO's sensitivity to an injected pulse. This is illustrated by the Spectre simulation results in Fig. 5.16 where the peak-to-peak value of the PTC is increased as the number of injection sites increases from single injection into stage R_1 ("1 inj") to injection into all four ring oscillator stages ("4 inj"). Although the effects of this multi-stage injection would be difficult to predict using conventional ILO models, the PTC is readily determined through schematic-level simulation and the results can easily be translated into the resulting lock range, lock time, or tracking bandwidth, as discussed in the previous sections.

It should be noted, however, that the limited bandwidth of each element in the delay line created by I_1 to I_4 results in the loss of some high frequency content of the injected pulse as it travels through each successive stage. This means that the pulse, which began




Figure 5.16: Spectre simulation results show that the peak-to-peak amplitude of the PTC increases as more injection sites are added.

with a width equal to half the bit period of the output clock will become wider by the time it reaches R_4 . These wide pulses are therefore able to produce a larger positive phase change in the MILO output, which corresponds to improved locking to frequencies lower than that of the free-running oscillator. They are, however, not able to produce a more negative phase in the MILO output, meaning that there is no improvement in locking to higher frequencies. Although this effect is typically not addressed by existing ILO models, it is clearly visible in the difference between positive and negative peak values in Fig. 5.16.

The peak PTC values were translated to lock ranges using equation 5.17 and are reported in Table 5.2. These results are compared to lock ranges obtained using the ISF method [54] and to those obtained directly from SPICE-level simulation using Virtuoso Spectre. To obtain the lock range in this way the transient response of the oscillator was simulated over a range of frequencies and a locked condition is identified by the settling of the MILO's output phase (relative to a reference signal) to some steady-state value. In order to ensure that the MILO is locked and that there is no eventual slipping in

	# inj. sites	1	2	3	4
ISF theory [54]	Δf_{low}	101 MHz	249 MHz	466 MHz	475 MHz
	Δf_{high}	87 MHz	174 MHz	226 MHz	242 MHz
PTC model	P_{max}	14.7 ⁰	30.4 ⁰	44.4 ⁰	56.5 ⁰
	P_{min}	-11.6 ⁰	-20.61 ⁰	-22.69 ⁰	-21.64 ⁰
	Δf_{low}	32 MHz	67 MHz	97 MHz	124 MHz
	Δf_{high}	25 MHz	45 MHz	50 MHz	47 MHz
Spectre sim.*	Δf_{low}	30 MHz	60 MHz	90 MHz	120 MHz
	Δf_{high}	30 MHz	40 MHz	50 MHz	50 MHz

*performed with 10 MHz resolution

Table 5.2: Comparison of lock ranges calculated using the ISF model [54] and the PTC model to those obtained directly using extensive SPICE-level simulations.

the output phase, the transient simulation must be run for several hundred clock cycles. This, combined with the fact that the step size of the frequency sweep must be small in order to accurately determine the lock range, results in simulations that consume a significant amount of time and resources. This highlights the usefulness of determining this information using the PTC method instead.

5.3.2 Frequency Pre-Conditioning

Another method of increasing the lock range of a multiplying ILO is to introduce a frequency pre-conditioning circuit that emphasizes the input's desirable harmonic. For example, an edge detector comprising a delay and XOR gate is often used for this purpose [56], and results in the MILO design shown in Fig. 5.17. This topology was developed in collaboration with a team of engineers while on internship at Rambus Inc. and was published as part of the bidirectional link presented in [21]. By incorporating a lock range of approximately 7%, this MILO is able to lock to the desired 2.8-GHz clock frequency across changes in voltage and temperature. It should be noted, however, that the PTC model presented in this chapter was developed separately and involved no collaboration with Rambus.

While the impact of adding an edge detector to the MILO would be difficult to include




Figure 5.17: Creating pulses at the reference clock edges emphasizes the desirable harmonic of the input thereby improving the lock range of the MILO.




Figure 5.18: The addition of a second edge detector with wide pulse widths further emphasizes the desired harmonic of the input signal.

in existing ILO models, its inclusion in the PTC simulations is trivial. First, the delays and XORs shown in Fig. 5.17 were included prior to the ILO. Then a DC offset was added to the first amplification stage in the delay chain in order to create the return-to-zero pulses shown in this figure. Simulations of the MILO using this injection technique show that it increases P_{pp} to 114° . Further increases are then achieved by adding a second edge detector set to create pulse widths equal to twice that of the original edge detector, resulting in the MILO topology shown in Fig. 5.18. Using this technique increases P_{pp} to 204° .

Furthermore, with the injected pulses now arriving at the same frequency as the output clock, it becomes unnecessary for the injected signal to return to zero between injected pulses and a full swing sinusoid can now be used as the injection signal instead. When this is done P_{pp} increases beyond the target value of 360° . The PTC for




Figure 5.19: Spectre simulations determine the PTC for ILOs using (a) one edge detector, (b) two edge detectors and (c) two edge detectors used to produce a sinusoidal (NRZ) injection signal.

each of these options was determined through transistor-level simulation in Spectre and is shown in Fig. 5.19. The PTC values reported for the case of two edge detectors with sinusoidal injection indicate that a 4-GHz ILO using this topology should be able to achieve a lock range that extends 0.7 GHz above f_0 (using $P_{min} = -250.9^0$ in Equation 5.13) and 1.06 GHz below f_0 (using $P_{max} = 383.3^0$ Equation 5.14).

Using this method P_{pp} was found to be 634.2^0 , which translates to a lock range of 1.76 GHz or 44% of f_0 . This compares well with both direct simulation of the lock range (40% of f_0) and the measured lock range of this MILO in a prototype chip (42.5% of f_0), which will be explained in detail in the following chapter.

5.4 Summary

Modeling ILO behaviour by using conventional frequency domain models requires the use of several parameters which are difficult to define. Modeling using the ISF-based model is accurate only when the injected signal strength is low. As an alternative to these options, the proposed PTC model of an oscillator can be used in conjunction with simple transistor-level simulations to accurately predict the behaviour of any ILO under any

injected signal. This makes the PTC model useful during the design of an ILO, helping to optimize the circuit for a given set of requirements. This model has been submitted for publication in TCAS-I.

Using this PTC model a MILO was designed to multiply a 1-GHz reference clock by 4 to produce a 4-GHz output clock. By simulating the PTC of the MILO at various stages throughout the design process, it was possible to quickly evaluate the impact that each change in topology had on the lock range. This in turn made it possible to achieve a wide lock range in a logical progression of design steps, resulting in a MILO with a lock range equal to 44% of the free-running frequency.

Chapter 6

A Wide Lock Range, Fast Power-On Clock Multiplier

Many wireline communication links require peak bandwidth operation for only a small fraction of their operating time [20]. The power efficiency of many electronic devices has therefore been improved by either varying the interface baud rate as bandwidth requirements change [4], and by scaling the supply voltage of the link [59]. Such systems are referred to as dynamic voltage and frequency scaling (DVFS) and are commonplace in many commercially available processors [59].

More recently, work has focused on improving power efficiency by completely powering down the link during periods of inactivity [35], [21]. While these techniques have proven effective at significantly reducing standby power consumption, they have so far been unable to combine this advantage with the ability to also scale the frequency of the links when they are powered on. Greater link flexibility and thus further power savings can be achieved by combining both techniques in applications such as the DVFS system shown in Fig. 6.1. In this system, the digital logic operates at a frequency and supply voltage that varies according to demand. Any communication that must occur with other chips must first pass through a serializer/deserializer (SerDes) block, which increases the data

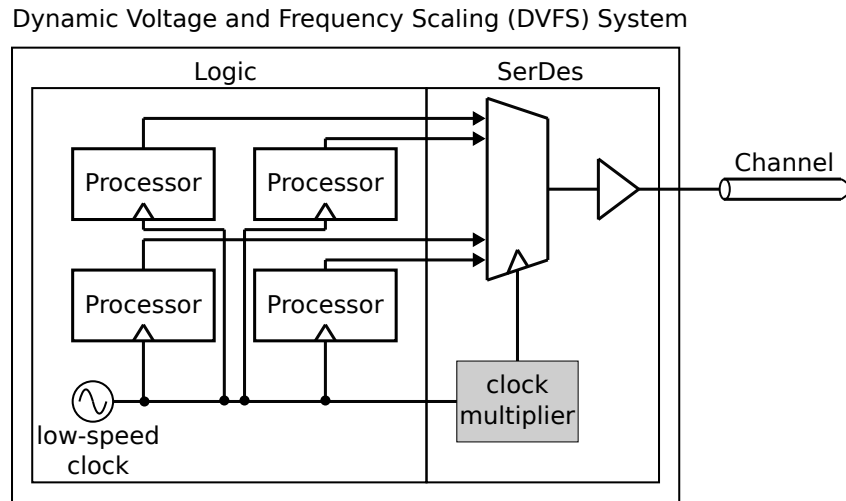


Figure 6.1: The design of a frequency agile clock multiplier that is suitable for fast power cycling can achieve link flexibility and power savings in DVFS applications.

rate of each link and thereby and limits of the number of I/O pins and interconnect traces that are required. The high-speed clock required for the SerDes operation is typically produced by a clock multiplier, as shown in Fig. 6.1 (note that only the serializer portion of the SerDes block is shown for simplicity). If this multiplier can be included in the blocks that are powered down when the link is inactive, further power savings can be achieved. However, the implementation of a multiplier that is both frequency agile and capable of fast power cycling presents significant design challenges, which will be explored in this chapter.

Multiplying injection-locked oscillators (MILOs) can be powered on quickly but have narrow lock ranges, typically less than 10% of the free-running frequency, even in applications where attempts have been made to maximize the frequency range [21]. PLLs [60], or MDLLs [61] can be tuned to accommodate a wide range of input frequencies, but their slow settling time makes them unsuitable for fast power-on architectures. While it is sometimes possible to adapt these loops for fast power-on applications by initializing their control voltages [35], this requires constant link speed between consecutive power-on cycles, limiting the usefulness of this technique for frequency agility. This chapter presents the first clock generator that is both frequency-agile and capable of fast power-




Figure 6.2: This chapter presents the first clock multiplier that is frequency-agile and has the ability to be powered on in under 10 cycles of the reference signal.

on. Measured results demonstrate an aggregate lock range of 55.7% and show that a valid clock output available within 10 cycles of the reference clock. The result is the first clock generator that is capable of performing the frequency shifting operation illustrated in Fig. 6.2 with no adjustments or tuning of any kind between power-up sequences.

6.1 MILOs with Adjacent Lock Ranges

As discussed in the previous chapter, the lock range of a MILO can be increased by either increasing the effective strength of the injected signal, or by increasing the sensitivity of the oscillator to this signal. Both of these effects can be captured by simulating the phase transfer characteristic of a MILO in response to one period of the injected signal. In the previous chapter, a MILO was designed using a ring oscillator with multiple injection points in combination with edge detectors to achieve a lock range of 45% of the free-running frequency. This MILO will serve as a starting point for the multiplier architecture presented in this chapter. While the lock range achieved by this design is a significant improvement over more traditional MILO architectures, some applications may require even greater frequency agility.

Since this can be difficult for a single MILO to achieve, a possible alternative is to employ multiple MILOs with adjacent lock ranges along with some control circuitry that is able to switch between MILO outputs when necessary. This idea is illustrated in Fig. 6.3 for the case of a clock multiplier with multiplication factor $N = 4$, designed to

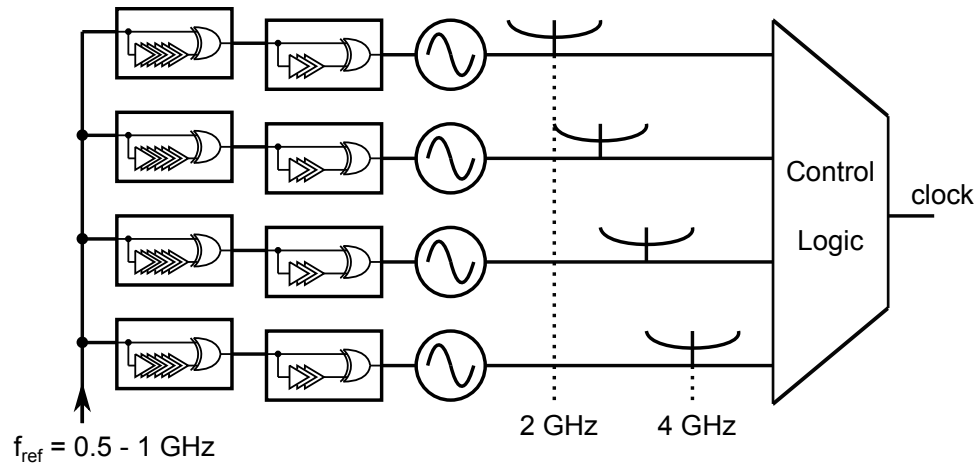


Figure 6.3: Four MILOs with adjacent lock ranges can cover an aggregate output frequency range from 2 to 4 GHz.

produce output clock frequencies ranging from 2 to 4 GHz. Although this example uses four MILOs to cover the desired frequency range, this number can easily be changed to accommodate any desired frequency range or required overlap of adjacent frequency ranges.

A critical component of the architecture shown in Fig. 6.3 is the control logic, which must monitor the output of each MILO in order to perform the following tasks:

1. Determine which MILOs are frequency locked.
2. Determine if a locked MILO's output frequency corresponds to $N = 4$ times the input frequency, which may not be the case if the MILO has locked to a different harmonic of the injected signal.
3. If two adjacent MILOs are locked to the correct frequency, then a single output should be selected based upon some desirable criteria.
4. After making a decision based on these criteria, the three unused MILOs should be powered down.

Further complicating these tasks is the fact that this monitoring and decision making must happen within only a few cycles of the reference clock in order to meet the fast

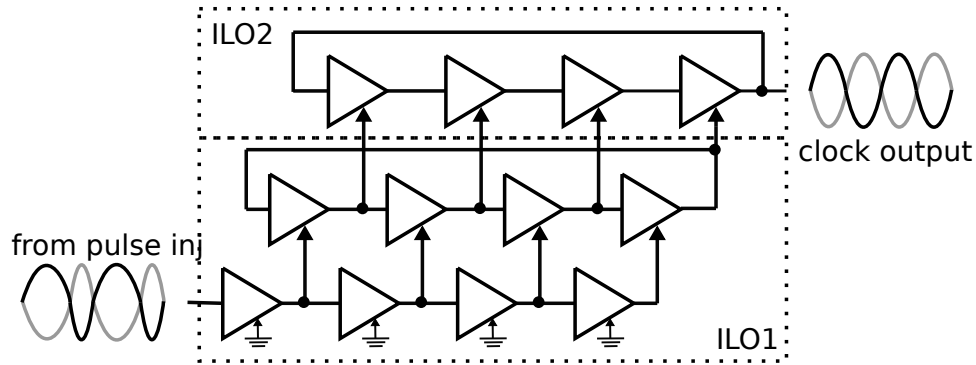


Figure 6.4: The addition of a second, identical ring oscillator compensates for DJ introduced by unequal pulse widths created by the edge detectors.

power-on requirement. The techniques used to perform these tasks in the allotted time are described in the following subsections.

6.1.1 Verifying Lock and Measuring Frequency Offset

Simulations of a single MILO show that, since the width of the pulses created by the edge detectors is fixed, the pulse widths of the injected signal can vary greatly as the frequency of the injected signal strays from the center of the MILO's lock range. This can create a significant amount of deterministic jitter (DJ) at the MILO output. To address this, a second ring oscillator was added to the output of the first, as shown in Fig. 6.4.

Making the second oscillator identical to the first and injecting signals into each oscillator stage ensures that the lock range will not be limited by the addition of the second ILO. The resulting improvement in DJ was simulated for this MILO structure in 65-nm GP CMOS and is displayed in Fig. 6.5.

In addition to its jitter filtering properties, the addition of a second ILO can also be used to verify that both oscillators are locked to an injected signal and also to measure their distance from the free-running frequency. As was illustrated in the previous chapter, the steady-state phase relationship between an ILO output and its injected signal is proportional to the frequency difference between this injected signal and the oscillator's free-running frequency (Equation (5.12)). This phase relationship can be quickly deter-




Figure 6.5: Spectre simulations show improvement in DJ performance without sacrificing lock range.

mined by using the outputs of each stage of the first ILO to trigger latches reading one output of the second ILO, as shown in Fig. 6.6. Details of the CML latches used in this circuit are given in a schematic diagram included in Appendix B, Fig. B.2. The inclusion of XOR gates is explained later in section 6.1.4.

These latches perform the function of a coarse time to digital converter (TDC). When the injection frequency is close to f_0 the phase of the highlighted injection signal (inj) will be approximately equal to that of the highlighted clock signal (clk), resulting in the latch outputs shown for case (b) in Fig. 6.6. As the injection frequency increases, the phase of “inj” will begin to lead that of “clk”, resulting in case (a) and vice versa as the injection frequency decreases, resulting in case (c). This means that bits D_1 to D_6 can be used to estimate how far the ILOs are from their shared free-running frequency. Since an ILO operating close to f_0 should produce less jitter and have good tolerance to any subsequent variations in voltage and temperature, this information can then be used to select a clock signal from two different MILOs which have locked to the correct multiple of the reference signal.

The addition of a sixth latch, which is triggered by the inverted output of stage 1 of

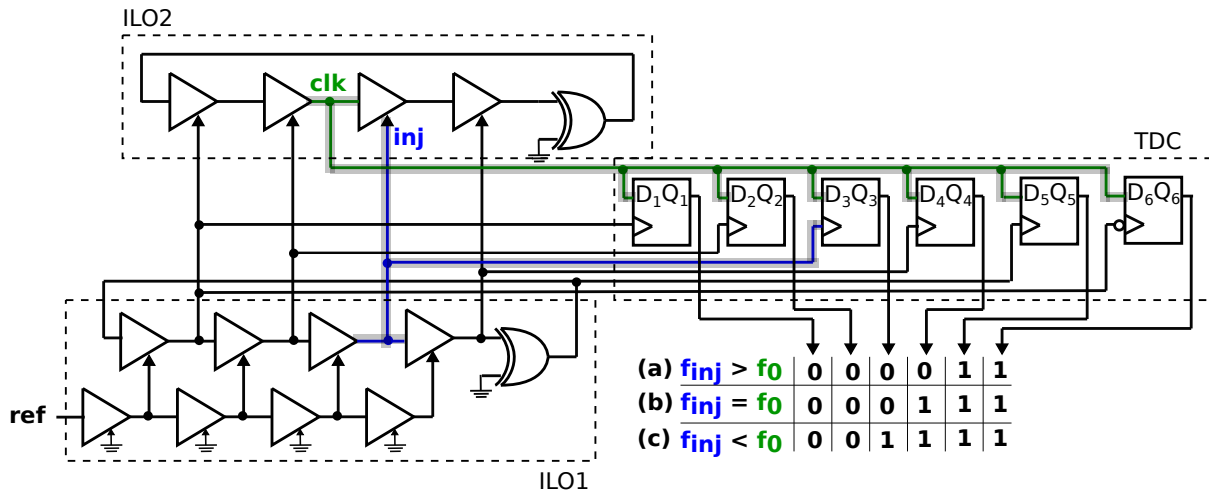


Figure 6.6: Latches at the output of each stage of the first ILO can be used to compare the phase relationship between the two ILOs.

ILO1, ensures that the latches are triggered over a range from 0 to 180° in steps of 30° , which is determined by the number of stages in the oscillator. This guarantees that D_1 and D_6 will have opposite values so long as each ILO is locked to the same frequency. Therefore the outputs of D_1 and D_6 can be used to indicate when the ILOs are in a locked condition and can trigger a power-down of any MILOs when this is not the case.

6.1.2 Verifying Multiplication Factor

Since the multiplier covers a broad range of frequencies, it is possible for two MILOs to be locked to frequencies at different multiples of the reference signal frequency. This effect is not captured by the frequency lock detection discussed in the previous section and, as a result, some additional logic is required to detect which MILOs are correctly locked to $N = 4$ times the reference frequency. This can be accomplished using the modified ripple counter architecture shown in Fig. 6.7.

In this structure the chain of latches exit the reset state when the reference clock signal goes high at t_1 and propagate a “logic high” signal down the chain as each successive output clock edge occurs. By the time the reference clock signal goes low at time t_2 , latch outputs Q_1 to Q_6 should correspond to the values shown in the table in Fig. 6.7 provided

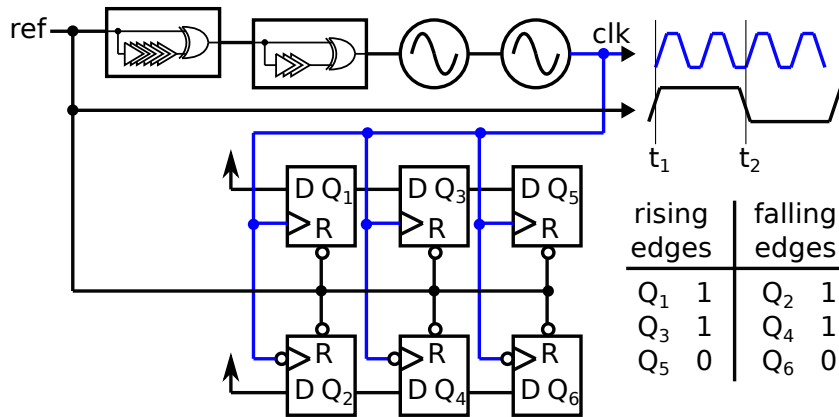


Figure 6.7: Latch chains verify the multiplication factor by ensuring that there are exactly 2 rising and 2 falling clock edges within one half period of the reference clock.

that the output clock frequency is 4 times that of the reference. Any deviation from this multiplication factor will result in different latch outputs, which will then trigger a power-down of the incorrectly locked MILO. It should be noted that, since the phase relationship between the output clock and the reference signal is unknown, it is impossible to know whether a rising or falling clock edge will arrive first after the reference signal goes high at t_1 , making it necessary to count rising and falling edges in separate counter chains, as shown in Fig. 6.7.

6.1.3 Power Down Unused MILOs

The power required to operate several MILOs continuously would likely outweigh any savings gained through frequency agility and the ability to power down the system during idle periods. As a result, it is necessary to power down unused MILOs as quickly as possible. If the edge counter is able to detect frequency lock to an incorrect multiple of the reference frequency or the TDC is able to detect an out-of-lock condition, as described previously, then the corresponding MILO can be powered down immediately. A block diagram of the system used to implement this strategy is illustrated in Fig. 6.8.

If neither of these conditions are found then it becomes necessary to compare the MILO's clock output to that of any adjacent MILO that may also be locked to the

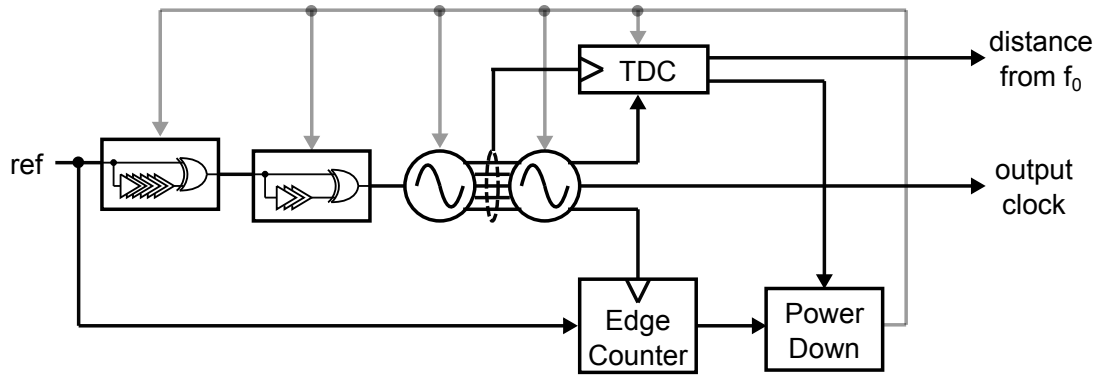


Figure 6.8: A MILO can be powered down immediately if the TDC detects an out-of-lock condition or if the edge counter detects an incorrect frequency multiplication ratio.

correct frequency. This system is pictured in Fig. 6.9 where data about the distance of each ILO’s output from its free-running frequency is passed to a “frequency offset compare” logic block that is external to the MILOs. This logic uses the output bits from the TDC to determine which multiplier is operating closest to its free-running frequency and sends a signal to the other MILO instructing it to power down. In the case where both TDC outputs are identical, the MILO with the lower free-running frequency is powered down. The reference clock signal is also applied to the power down logic where it is used to latch each of the TDC outputs for half cycles of the reference clock. This gives the frequency offset compare logic enough time to make its decision and power down a MILO without being susceptible to sudden changes in the TDC outputs, which can occur due to the deterministic jitter present in the ILO outputs.

If any of the signals from either the TDC, edge counter, or frequency offset comparison logic trigger a power-down then the circuits within the unused multiplier are turned off by applying a logic signal to switches in the tails of each differential pair in the MILO, as shown in Fig. 6.10. Since all parts of the MILO employ CML signaling, this ensures that the power drawn by the multiplier goes to almost zero when powered down. It should be noted that this idea can be easily extended to cut power to the monitoring circuitry (TDC, edge counter, external power-down logic) of the MILO that remains active in order to achieve further power savings. However, this technique was not implemented in

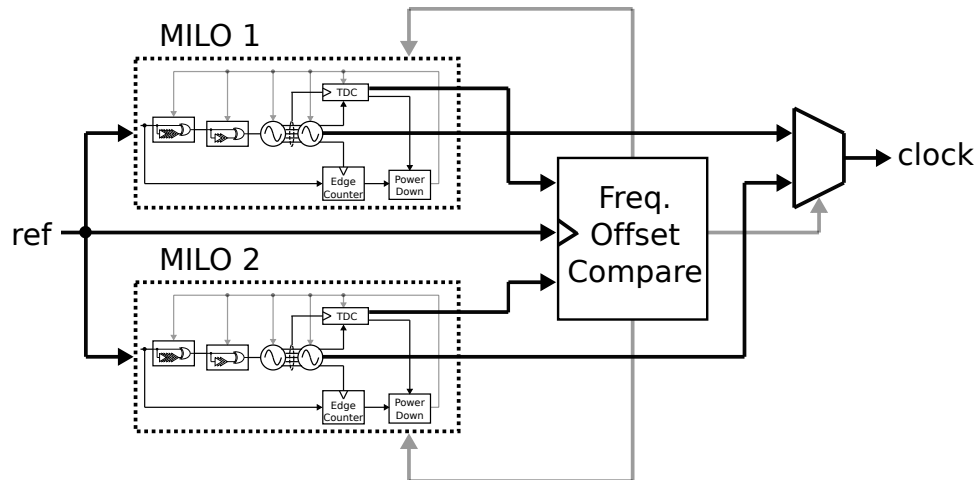


Figure 6.9: If two MILOs are locked to the correct frequency, the power-down decision is made by determining which MILO is operating closest to its free-running frequency.

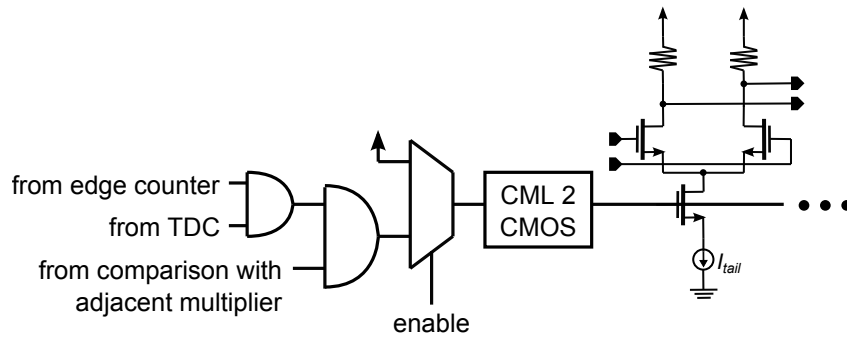


Figure 6.10: Power down of any unused MILOs is accomplished by blocking tail currents in all CML stages.

this prototype in the interests of simplicity and testability.

To ensure that the CML stages are powered down effectively and that any leakage current present after power-down remains as small as possible, the power-down logic signal is converted to full-swing CMOS logic levels by the circuit shown in Fig. 6.11. This circuit is based on the CML-to-CMOS converter presented in [62] with the modification that the width of transistor M_2 is made twice that of transistor M_1 . This ensures that if both V_{CML} inputs go to V_{dd} , which occurs when the preceding CML stages are turned off, the V_{CMOS} output remains low, keeping the MILO circuits powered down.

The decision to power down a multiplier can only be made after all bias voltages and oscillator outputs have settled to their steady-state operating conditions. This is




Figure 6.11: Converting the power down signal to CMOS logic levels ensures successful power-down and minimizes leakage power.




Figure 6.12: Timing diagram of the power-on sequence for a 1-GHz reference signal.

accomplished by adding an enable signal to the power down logic in Fig. 6.10. This signal can be created automatically by counting some number of reference clock cycles using a chain of latches similar to that used to measure the frequency multiplication ratio. By waiting for 8 reference clock cycles before enabling power down from the edge counter and TDC logic, and an additional 2 clock cycles before enabling power down from frequency offset comparison with adjacent multipliers, it is possible to avoid making premature power-down decisions. This sequence is shown in Fig. 6.12 for a possible power-up scenario using a 1-GHz reference signal.

Counting of reference clock periods is performed by a series of latches as shown in Fig. 6.13. Upon receiving a “Start” signal from an off-chip signal source, the latches




Figure 6.13: Power down of individual MILOs is enabled after 8 cycles and power down resulting from comparison of two correctly locked MILOs is enabled after 10 cycles.

exit the reset state and begin to pass the value of “Pdown_en” down the latch chain. Pdown_en is set off-chip in order to provide the ability to manually disable the power down circuitry for testing purposes.

6.1.4 Creating Adjacent Lock Ranges

The task of creating two ILOs with identical free-running frequencies, along with two edge detectors—one with a pulse width equal to $1/f_0$ and the other with a pulse width of $2/f_0$ —is non-trivial. While it is possible to design some method of tuning these blocks so that calibration can be used to individually tune each ILO to a desired frequency and each edge detector to a desired pulse width, this calibration can quickly become cumbersome as the number of MILOs in the system is increased.

To avoid this problem it is possible to take the four delay stages and an XOR gate used in the edge detector structure and use this circuit as the basis for the second edge detector and both ILOs in the MILO as shown in Fig. 6.14. This technique ensures that the pulse widths and ILO operating frequencies will be well matched regardless of PVT variations, so long as care is taken to match layout parasitics. A schematic diagram of the CML XOR gate used in each of these blocks is included in Appendix B, Fig. B.3. It should be noted that buffers were used at the output of each ILO stage to ensure that the load seen by each delay element in the multiplier is identical, but these were omitted from Fig. 6.14 for simplicity.




Figure 6.14: Using four delay stages and an XOR gate as the building block for each component in the MILO ensures good matching between pulse widths and ILO free-running frequencies.

Creating MILO lock ranges that are adjacent to each other with enough overlap to ensure that no gaps are present in the overall lock range, while also keeping this overlap small enough to not compromise the achievable lock range can be challenging. In order to accomplish this, load capacitances were added to the loads of each delay element shown in Fig. 6.14. Then, by increasing the physical size of these capacitances from one MILO to the next, the desired overall lock range and amount of overlap can be achieved. In order to account for any discrepancies between simulated and measured behaviour, these capacitances were implemented using NMOS varactors to provide some tunability. Since the delay stages and their associated varactors are identical in each stage of the MILO, it is therefore possible to adjust the MILO's operating frequency by manipulating a single varactor control voltage, which is applied to both the ILOs and the edge detectors. Once this initial calibration is performed to set each MILO to a desired operating frequency, no further adjustment is required.

6.2 Measured Results

A prototype of the clock multiplier described in this chapter was fabricated in a 65-nm GP CMOS process. The multiplier contains four parallel MILOs with all of their associated

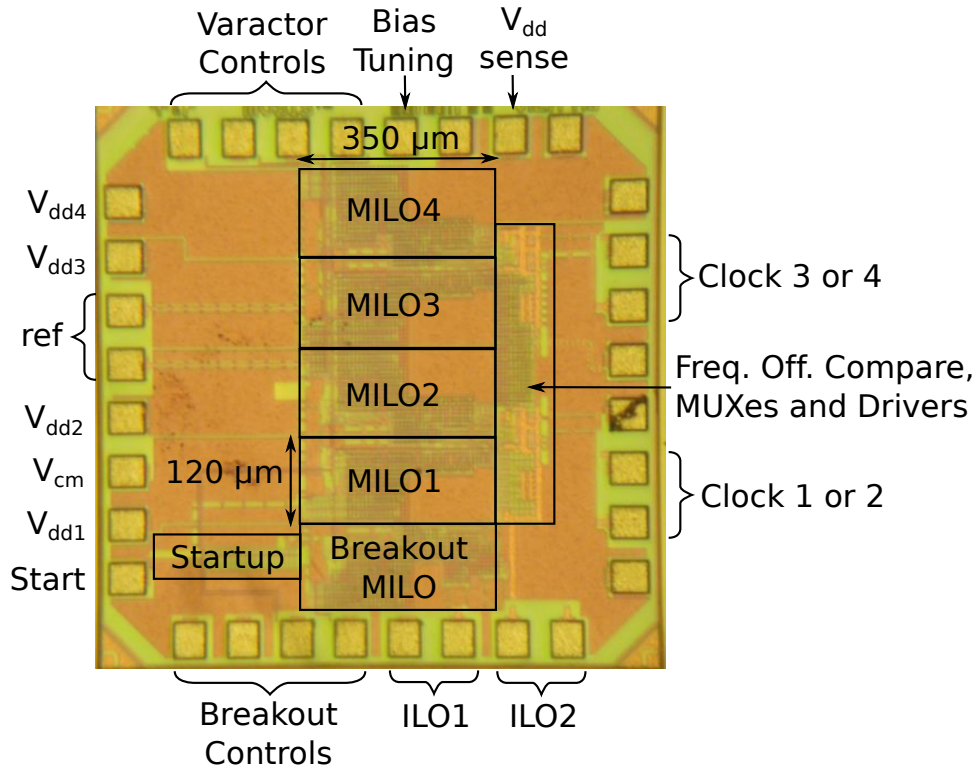


Figure 6.15: Die photo of the clock multiplier in 65-nm GP CMOS.

control logic, as well as a breakout MILO that provides off-chip output from each ILO and is isolated from any frequency offset power-down logic. Besides the absence of this logic, the breakout is a replica of MILO3. A die photo of the 1 mm x 1mm chip is shown in Fig. 6.15. Of special note is the fact that each MILO is powered by a dedicated supply voltage in order to provide isolation from variations in power supply which can be created when one or more MILOs are quickly powered down. This also improves testability of the circuit by making it possible to test MILO lanes either individually or in any desired combination.

6.2.1 Free-Running Frequencies

With no injected signal applied and all power-down logic disabled, the free-running frequency of each MILO was measured over a range of varactor voltages. The results are plotted in Fig. 6.16 and show that each MILO can be tuned to a range of frequencies in




Figure 6.16: Measured free-running frequency of each MILO shows reasonable spacing between adjacent lanes and, if necessary, these free-running frequencies can be adjusted using the varactor control voltages.

order to achieve the appropriate amount of overlap between adjacent lock ranges. This figure also shows that if all MILOs are operated using the nominal setting of 0.9 V for all varactors, which is the desired case for normal operation, the MILO operating frequencies should be close enough to avoid gaps in the aggregate lock range. Furthermore, this aggregate range should span from at least 2.5 to 5 GHz.

6.2.2 Phase Transfer Characteristic

Direct measurement of the PTC of an ILO is impractical since the application of a single period of the injected signal is non-trivial, as is the ability to measure the real-time phase response of the MILO output to such a stimulus. Instead, it is possible to obtain part of the PTC using the relationship shown in Equation (5.12), which is repeated here for convenience.

$$P(\phi_{ss}) = \frac{N\Delta\omega}{\omega_0} 360^\circ \quad (5.12)$$

Since the fabricated ILO was designed to have a multiplication factor of $N = 4$ and was measured to have $f_0 = 4$ GHz, it is possible to measure the resulting values of ϕ_{ss} as




Figure 6.17: Measured values of (a) ϕ_{ss} for various frequency offsets can be translated to (b) the PTC of the ILO. These measurements show good agreement with SPICE-level simulations.

various frequency offsets, Δf , are applied. However, since changes to the frequency of the injected signal also introduce an unknown amount of phase shift in the signal that arrives at the MILO, Δf cannot be varied by changing f_{inj} . It is instead possible to achieve the same effect by keeping f_{inj} , and therefore the phase of the injected signal, constant and sweeping the varactor control voltage, thereby effectively sweeping f_0 .

This technique is able to reproduce a portion of the PTC between the peaks of the P function. Due to the fact that the varactors in the breakout MILO are only able to achieve free-running frequencies ranging from 3.4 to 4.7 GHz, this limits the $\Delta\omega$ that can be applied, thereby limiting the PSF values that can be observed. Despite this, Fig. 6.17(a) compares the ϕ_{ss} values measured in this way to those obtained from SPICE-level simulations. Fig. 6.17(b) then translates these measurements to the PTC using Equation (5.12) and compares them to the PTC that is obtained through simulation as discussed in the previous chapter. In both cases there is good agreement between simulated and measured results.




Figure 6.18: Measurements show wide lock ranges for MILO1, MILO2 and the Breakout MILO, but problems with the reference clock distribution cause a reduction in the lock ranges of MILO3 and MILO4.

6.2.3 Lock Range

With all power down logic disabled and varactor voltages set to their mid-range values, the lock range of each MILO was individually measured and the results are shown in Fig. 6.18. MILO1 and MILO2 show reasonable overlap with wide lock ranges that increase with f_0 , resulting in ranges of 32.2% and 36.3% respectively. This is not true, however, for MILO3 and MILO4 whose lock ranges measure only 18.7% and 8.6% of f_0 , respectively. Since the only designed difference between each MILO is the size of the varactor load capacitance implemented in the CML delay stages, it is likely that this drop in lock range for MILO3 and MILO4 is due to factors external to the MILOs themselves.

The likely cause of this problem is the clock distribution network which is used to apply the reference clock signal to each MILO as well as the startup circuit and the frequency offset compare logic. Although a $50\ \Omega$ termination is placed near its application to the chip at the “ref” pads in Fig. 6.15, the remainder of the distribution network consists of long lengths of unmodeled transmission lines that terminate at MOSFET gates. These are essentially open-circuit terminations, which can cause reflections that

PTC model	P_{max}	383.3 ⁰
	P_{min}	-250.9 ⁰
	Δf_{low}	1.06 GHz
	Δf_{high}	0.7 GHz
Measured	Δf_{low}	1 GHz
	Δf_{high}	0.7 GHz

Table 6.1: Comparison of measured lock range of the breakout MILO to those calculated using PSF simulations.

may create undesired signals at the MILO inputs. This conclusion is reinforced by the fact that the Breakout MILO, which is a copy of MILO3 but is attached to a different point of the clock distribution network, shows no such lock range limitation, achieving a range of 42.5% around a free-running frequency of 4 GHz. This lock range shows excellent agreement with simulated PSF values as can be seen in Table 6.1.

Despite this problem, correct operation of MILO1 and MILO2 is sufficient to demonstrate the feasibility of the multiple-MILO technique to create a frequency-agile, fast power-on clock multiplier. With all power-down logic enabled, the top portion of Fig. 6.19 shows the overall lock range of a multiplier consisting of MILO1 and MILO2. By combining the 0.84 GHz lock range of MILO1 with the 1.24 GHz lock range of MILO2 and incorporating a 320 MHz overlap region, a 1.76 GHz lock range is created, which is equivalent to 55.7% of the 3.16 GHz f_0 . The overlap region can be initially calibrated to other desired values in order to tolerate variations in voltage and temperature by using the varactor controls. Also, although not shown in this plot, tests were conducted with reference clock frequencies equal to $\frac{1}{16}$, $\frac{1}{8}$, $\frac{1}{2}$ and 1 times a frequency within the valid lock range. In every case the resulting incorrect frequency multiplication factor was detected and the MILO was powered down.

The lower portion of Fig. 6.19 shows the steady-state (SS) current measured from the dedicated MILO supplies, illustrating the point at which the Frequency Offset Compare logic switches the multiplier output from MILO1 to MILO2 as the reference frequency is increased. This transition occurs well within the overlap region although not quite

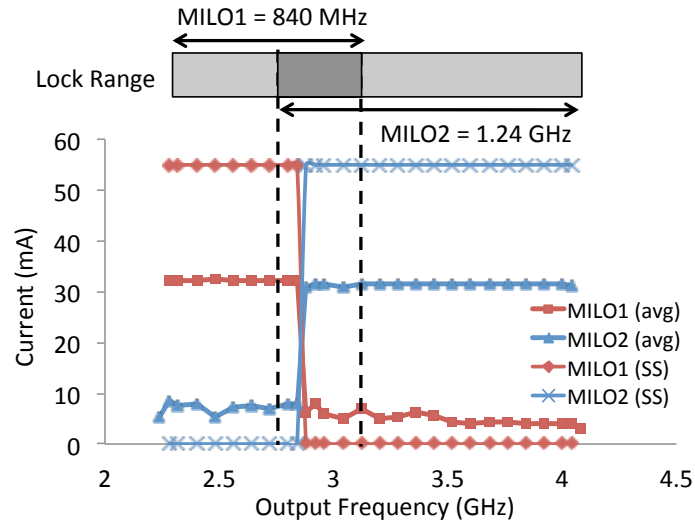


Figure 6.19: Two MILOs are able to increase the overall multiplier lock range to 55.7%. The point at which logic switches between MILOs is illustrated by measurements of the steady-state (SS) current drawn by each MILO, as well as the average (avg) current drawn when the multiplier is active for 50% of the time in 50-ns bursts.

in its center because of the relatively coarse resolution of the TDCs in each MILO. To highlight the power savings that are realized by using the fast power on/off capability, the average (avg) current is also plotted showing a reduction to 32.2 mA when the multiplier circuit is active for 50% of the time for 50-ns bursts. Although not shown in this plot, the average continues to shrink as active time decreases or as burst length increases. This is in contrast to traditional clock synthesizers which must either be left on at all times [60] or sacrifice their frequency agility in order to achieve fast power-on [21]. Power consumption is analyzed in more detail in section 6.2.7.

6.2.4 Deterministic Jitter

The output of both ILOs in the Breakout MILO structure (“ILO1” and “ILO2” in Fig. 6.15) were taken off-chip in order to allow for analysis of their jitter behaviour. The peak-to-peak DJ was obtained by using oscilloscope cursors to measure clock half-periods and then taking the difference between the maximum and minimum values. The results of this measurement are displayed in Fig. 6.20 and show that the shape of the




Figure 6.20: Measured results show that the addition of ILO2 provides a reduction in DJ created by unequal pulse widths in the injected signal.

curves agrees well with the simulated results in Fig. 6.5. The resulting improvement in DJ achieved through the addition of ILO2 ranges from 2.5 to 30.5 ps across the lock range.

DJ at the output of both ILOs is minimized when the reference clock frequency is well-matched to the fixed pulse widths generated by the edge detectors, resulting in a low-jitter signal being injected into the ring oscillators. At frequencies far from this point the periodic variations in pulse widths create significant DJ. Due to the dual edge detector structure, this DJ occurs in repeating 4-bit patterns that can be detected and isolated from the random jitter (RJ) using the jitter decomposition functionality of the Agilent DCA-J oscilloscope. Fig. 6.21 shows this repeating pattern by plotting the DJ per bit for a 4-bit pattern. This pattern is further illustrated by Fig. 6.22, which shows a histogram of the total measured jitter, showing four distinct peaks due to the DJ.

It should be noted that the DJ values reported in Fig. 6.20 are higher than the simulated values (shown previously in Fig. 6.5) due to the fact that inadequate models of the QFN package parasitics were used when designing the $50\ \Omega$ output buffers used to send the multiplier's clock outputs off-chip. Although these buffers were designed to

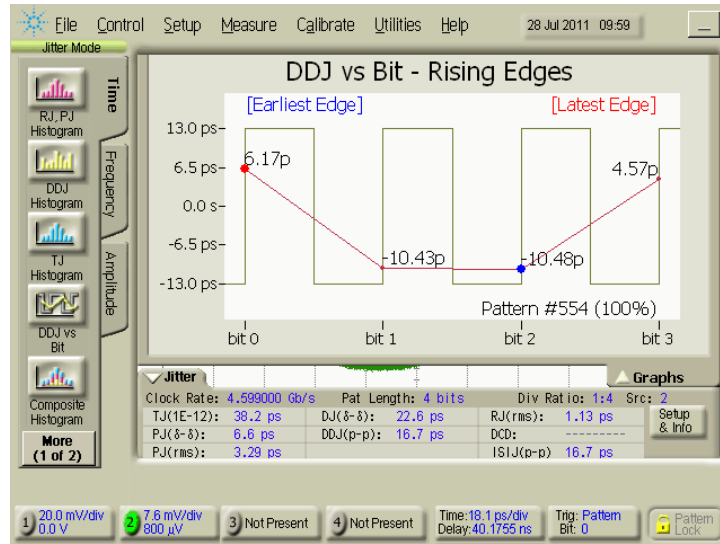


Figure 6.21: Decomposition of the measured jitter shows a DJ pattern that repeats in a 4-bit pattern due to the pulse widths of the injected signal.

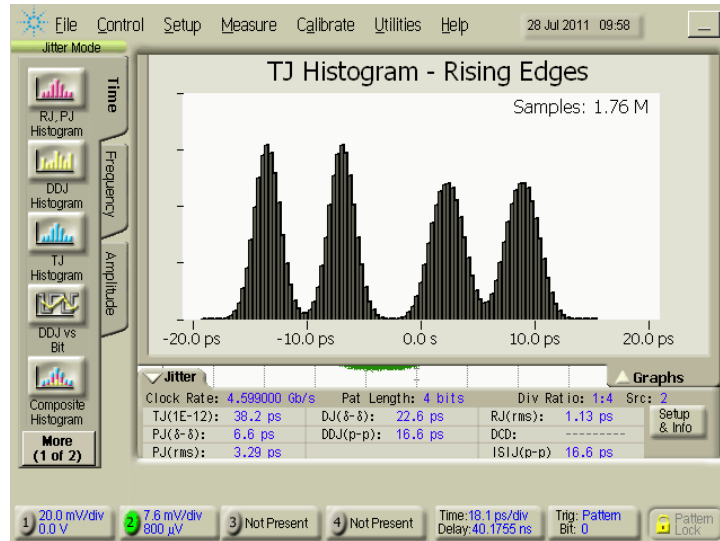


Figure 6.22: A histogram of the total measured jitter shows four distinct peaks representing the DJ.




Figure 6.23: Including package parasitics at the simulated clock outputs identifies the cause of small measured clock amplitude.

supply a 350 mV peak-to-peak per side differential signal to a $50\ \Omega$ load, output clock amplitudes measured only approximately 50 mV peak-to-peak per side. This effect was duplicated in simulation by adding the package parasitic structure shown in Fig. 6.23 to the chip outputs. These parasitics attenuate the high-frequency output to such a large degree that the subharmonics due to the injected reference signal become relatively more prominent, thereby creating an increase in DJ.

The effects of these parasitics can be partially overcome by increasing the power supplied to the output drivers. A schematic diagram of the two-stage, $50\ \Omega$ output drivers used in this application is shown in Appendix B, Fig. B.4. When the power delivered to the output drivers is increased by 50%, Fig. 6.24.(a) shows that the resulting output swing is increased to approximately 100 mV peak-to-peak per side. This improved output swing then translates to a reduction in DJ at the output of the MILO as shown in Fig. 6.24.(b). Unfortunately, since driver power could not be controlled independently from the rest of the circuit in the prototype chip, increasing driver power by 50% also results in a similar increase in the power consumption of the entire circuit. As a result, this high power setting was not used to measure any other circuit performance metrics.

6.2.5 Random Jitter

Although the DJ suffers from the reduced output signal amplitude, RJ measurements show no similar degradation due to this effect. Fig. 6.25.(a) shows a histogram of the random jitter obtained using the jitter decomposition function of the oscilloscope. This




Figure 6.24: Increasing driver power (a) increases output swing to approximately 100 mV_{pp} per side and (b) results in a decrease in DJ.

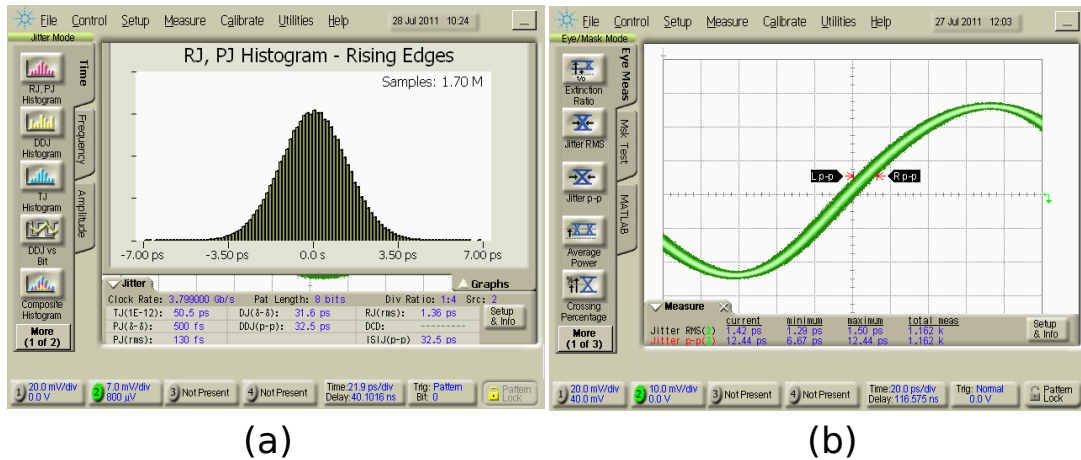


Figure 6.25: Histogram (a) and oscilloscope capture (b) showing that the RJ of the clock signal is approximately 1.4 ps-rms.

shows a measured RJ of 1.36 ps-rms, which agrees well with the 1.42 ps-rms and 12.44 ps-pp measurements obtained without jitter decomposition as shown in Fig. 6.25.(b). This measurement was performed by triggering the oscilloscope with the injected reference signal at $f_0/4$, making it possible to examine clock edges in isolation from the DJ.

In contrast to the deterministic jitter, RJ is independent of the injected frequency as the measured RJ of both ILOs, as shown in Fig. 6.26, remains at approximately 1.4 ps-rms over the entire lock range.




Figure 6.26: Measurements at the output of both ILOs shows that the RJ remains near 1.4 ps-rms across the entire lock range.

6.2.6 Fast Power On/Off Transients

Obtaining measurements of transient startup behaviour in response to an applied “Start” signal is non-trivial since any delays experienced by this signal between its source and its application to the chip must be deembedded from the measurement results. The test setup used to accomplish this is shown in Fig. 6.27. After being applied to the PCB, the Start signal is passed through a series of inverters in order to shorten its transition time before being applied to the prototype chip (DUT). The last stage of this inverter chain, along with its associated PCB trace length is duplicated and sent off the PCB to be measured by an oscilloscope. Using identical SMA cables to connect both this signal and the clock output to the oscilloscope ensures that their delays are reasonably well matched and that the two signals received by the scope present an accurate representation of the chip’s behaviour.

By using a 10 MHz signal source to repeatedly apply a Start signal to the circuit in 50-ns bursts, the power-on transient behaviour of the multiplier can be captured by a sampling oscilloscope. The results of this measurement technique are shown in Fig. 6.28(a) where the Start signal is shown in lower curve and the MILO output is shown

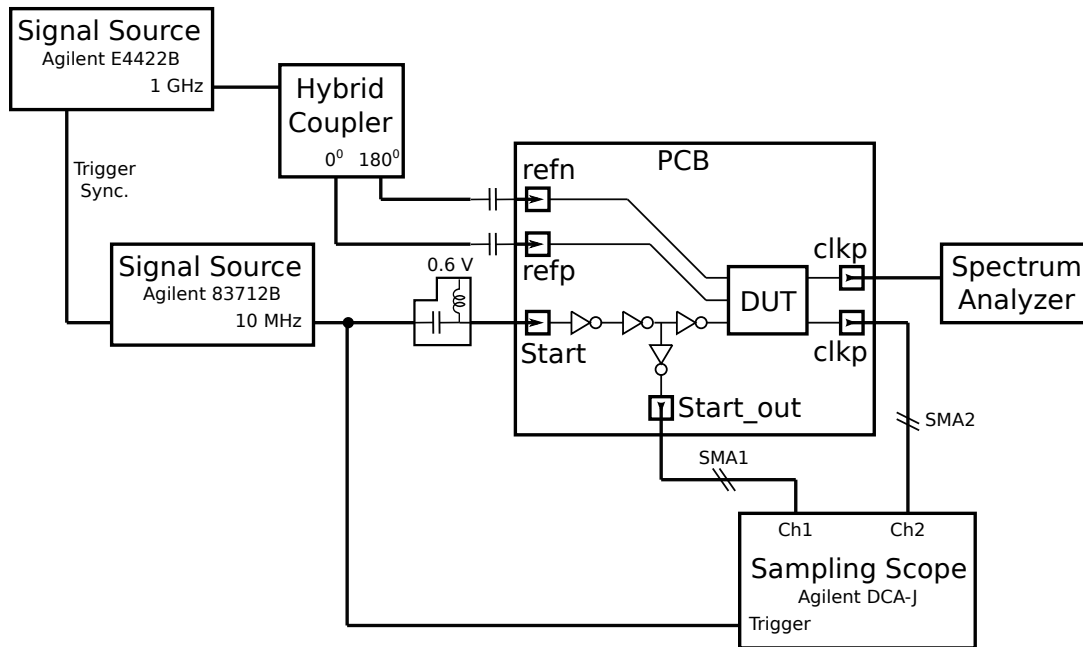


Figure 6.27: Test setup used to capture the transient response to an applied Start signal.

in the upper curve. The power-up behaviour of the circuit is examined more closely in Fig. 6.28(b), which shows that the delay between application of the Start signal and the beginning of oscillations at the clock output is approximately 3 ns.

To measure the transient behaviour of multiple MILO outputs simultaneously the test setup is modified as shown in Fig. 6.29. Since the prototype chip is capable of sending one of either MILO1 or MILO2 and one of either MILO3 or MILO4 off-chip, it is possible to apply the outputs from MILO2 and MILO3 to the oscilloscope simultaneously.

By again applying a Start signal to the circuit in 50-ns bursts, the transient behaviour of both of these clock outputs was measured and is shown in Fig. 6.30(a). This plot shows how the lower, unlocked clock (MILO3) is correctly identified by the power-down logic in 8.5 ns, or approximately 8 cycles of the 0.95 GHz reference signal used in this test. In Fig. 6.30(b) MILO1 is also enabled, providing two correctly locked clock outputs in the upper signal, and an unlocked clock output in the lower signal. In this case the unlocked MILO is identified and powered down in 9 ns. Then, after approximately 10 reference clock cycles (11 ns) the locked MILO that is furthest from its free-running frequency is

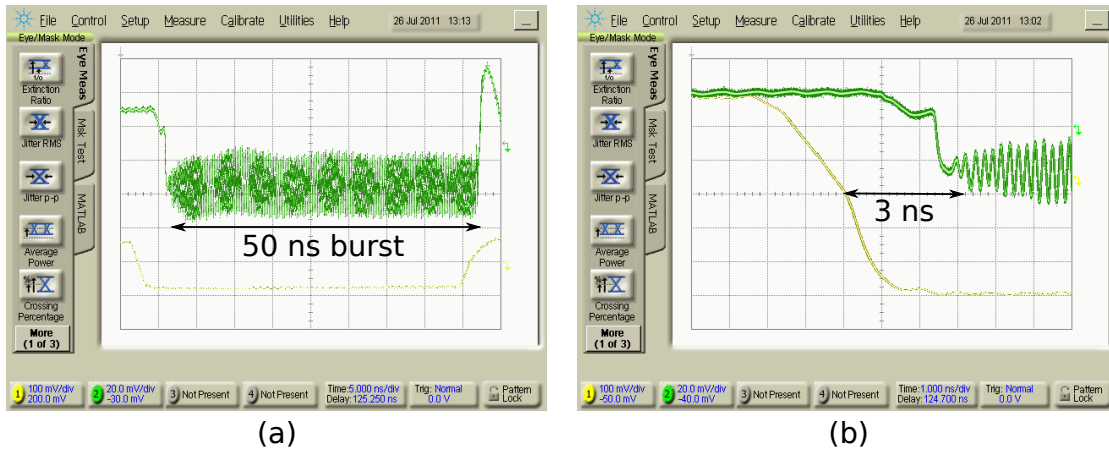


Figure 6.28: Repeated 50-ns bursts allow for power-on transient behaviour of the multiplier to be captured in (a). Zooming in on the start of these bursts (b) shows 3 ns delay between Start signal and output oscillations

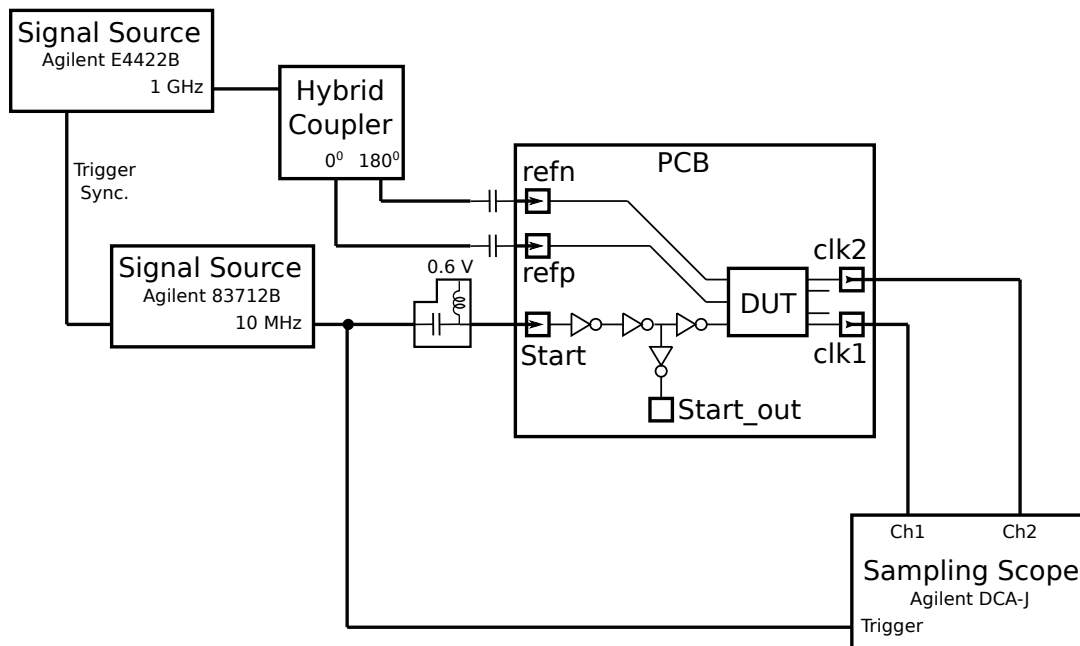


Figure 6.29: Test setup used to capture the transient response of two output clocks.

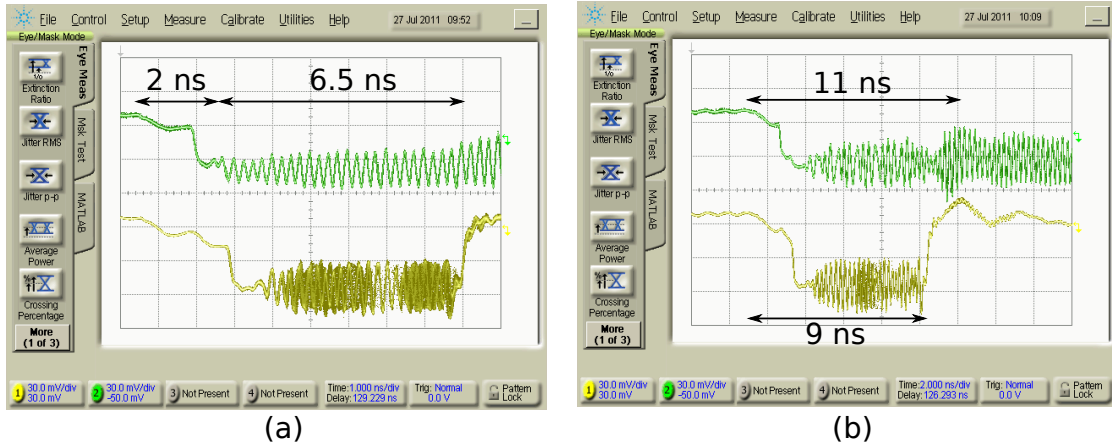


Figure 6.30: Transient startup behaviour when (a) two and (b) three MILOs are enabled shows identification and power down of unlocked clock(s) in 10 reference clock cycles.

identified and powered down, leaving only a single valid output clock.

It should be noted that output of MILO3, which is shown in the bottom of these figures experiences an additional 1 ns delay in the start of its oscillations when compared to the output of MILO2. This is the result of the increased physical separation between each MILO and the startup circuitry, as shown by the die photo in Fig. 6.15, and indicates that care must be taken to minimize this distance in the layout.

6.2.7 Power Consumption

Since this architecture involves rapid powering on and off of several blocks, the power that is consumed from the 1.1 V supply can not be adequately described by a simple, steady-state metric. During power up, each MILO that is turned on consumes 55 mA. For the MILOs that are powered down, this power is only consumed for a maximum of 10 reference clock cycles before it returns to 0 mA, leaving only 1 MILO consuming 60.5 mW during steady-state operation. In addition, steady-state power is also consumed by the startup circuit as well as the frequency offset comparison logic block, the output multiplexer and the $50\ \Omega$ output driver. These contribute an additional 35.5 mW of power consumption during steady-state operation.

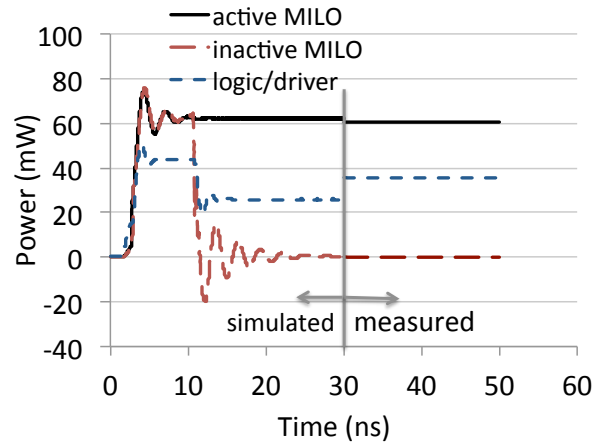


Figure 6.31: Simulated transient power consumption during one power-up cycle of the multiplier is compared to measured steady-state power consumption.

Although these steady-state power consumption values were measured, transient power consumption during power-up sequences could not be measured due to the short duration of these events. Instead, simulation results of a power-up event are shown in Fig. 6.31 along with the measured steady-state values stated above in order to lend confidence to these simulated results. Ringing seen around the steady-state levels at power on or off are caused by the package parasitics at the chip supply pads.

Since the output clock cannot be used until the power-up sequence is complete, power consumed during these 10 cycles of the reference clock penalizes the efficiency of the multiplier. The extent of this penalty can be calculated using the simulated transient power consumption values. For a 1-GHz reference clock, which corresponds to a power-up sequence duration of 10 ns, each MILO consumes approximately 620 pJ of energy while the logic circuits consume approximately 430 pJ. As a result, the efficiency of the multiplier depends strongly on the number of bits that are transmitted with each power-up sequence. This effect is captured by Fig. 6.32, which shows how total energy consumed remains below 25 pJ/bit for burst lengths of at least 250 bits but increases for shorter burst lengths.

Regardless of the burst length, however, the average power consumed by the multi-




Figure 6.32: Energy consumed by the power-up sequence means that efficiency of the multiplier suffers when used in short bursts.

plier scales nearly linearly with the percentage of active time. This is illustrated using the simulated data shown in Fig. 6.33 for burst lengths of 400 bits. This dependence highlights the fact that the anticipated percentage of active time of the multiplier is a key factor to be considered in the multiplier’s implementation.

In order to compare the performance of the multiplier to the state-of-the-art, Fig. 6.34 reports the steady-state power consumption in relation to the power-on time and frequency range of other reported clock multipliers. Since it was not reported, power-on time was estimated in [22] and [60] from their PLL settling times, τ_s according to the equation

$$\tau_s = \frac{10}{0.1f_{ref}} \quad (6.1)$$

Further details, such as the technology and area used as well as the jitter performance of these publications are included in Table 6.2.





Figure 6.33: Average power consumption of the multiplier scales linearly with the percentage of active time.



* estimated power-on time

Figure 6.34: Comparison to other clock multipliers shows the improvement in power-on time and frequency range achieved by this work.

	[21]	[22]	[60]	[61]	This Work
Technology	45-nm	90-nm	65-nm	65-nm	65-nm
Area	-	0.325 mm ²	1.1 mm ²	29.33 mm ²	0.15 mm ²
Power-on	8 ns	100 ns	277 ns	146 ns	10 ns
f_0	2.8 GHz	20 GHz	19.2 GHz	0.8 GHz	3.16 GHz
Freq. Range	7.1 %	4.7 %	35.8 %	103 %	55.7 %
Power	16.2 mW/GHz	5.3 mW/GHz	4.2 mW/GHz	133.8 mW/GHz	30.4 mW/GHz
Active Power	45.3	105 mW	80 mW	107 mW	96 mW
Standby Power	0 mW	105 mW	80 mW	107 mW	0 mW
Jitter	11 ps-pp (DJ)	0.17 ps-pp (TJ)	-	38 ps-pp (TJ)	30 ps-pp (DJ)

Table 6.2: Performance comparison with other recent publications.

6.3 Summary

Fast power cycling and frequency agility can increase the flexibility of communication circuits, allowing them to improve their power efficiency. However, the design of a clock generating circuit that is capable of meeting both of these requirements presents a significant challenge. The clock multiplier presented in this chapter introduces a MILO structure that addresses these challenges and shows how the frequency agility of the multiplier can be further improved by using parallel MILOs with adjacent lock ranges.

The fabricated multiplier is able to use two MILOs to achieve a frequency lock range of 55.7% of the 3.16 GHz centre frequency. Although the pulse injection scheme used introduces some DJ, this effect is mitigated by the addition of a second ILO, which improves DJ by anywhere from 2.5 to 30.5 ps, depending on the injected frequency. Transient tests of the power-on behaviour of the multiplier verify that the control logic is able to identify the correct output clock and power down the other MILOs within 10 reference clock cycles of the arrival of a start signal. The active MILO and its associated logic consume a total of 96 mW during steady-state operation, but this power scales with the amount of time during which the link is active. This multiplier has been submitted for publication at CICC in September 2012.

Chapter 7

Conclusion

In order to operate efficiently under a variety of channel conditions and data rates during both active and idle periods, wireline receivers must incorporate a great deal of flexibility into their design. This thesis has examined several aspects of adaptive receiver design in an effort to advance the state of the art in this field.

The preamplifier presented in Chapter 2 provides 10 dB of gain control in the input stage of an AFE while simultaneously maintaining a good impedance match to the wireline channel over a wide range of frequencies. This impedance match is achieved without the use of a purely resistive matching network, which helps to improve the noise performance of the receiver. In addition, this circuit responds to varying requirements of noise performance or linearity by automatically adapting its power consumption in response to changes in the received signal strength, which are seen as changes to the preamplifier gain settings. This helps to improve both the dynamic range and the efficiency of the receiver.

Automatic adaptation of both this gain setting and the EQ peaking is performed using the scheme presented in Chapter 2. This method shows that both of these AFE components can be configured intelligently to operate correctly over a wide variety of channels by adding only a modest amount of overhead to the receiver. Using this method,

one slicer, a low-speed ADC, and the circuitry required by the optimization algorithm are all that are required to optimize the vertical eye opening.

Another important feature of a flexible receiver is the ability to adapt to a variety of data rates and power levels as needed. This requirement can be particularly challenging when it comes to the design of the clock circuitry, which traditionally operates over a narrow range of frequencies and experiences significant latency in transitioning between power levels. These challenges were addressed by the frequency agile clock multiplier reported in Chapter 4. The use of MILOs in this design gives the multiplier the ability to transition from being completely inactive to fully operational in 10 cycles of the reference clock. Furthermore, the use of several MILOs with adjacent lock ranges was demonstrated as a way to increase the frequency range of the multiplier. By monitoring the outputs of these MILOs, it is possible to quickly identify the optimal output clock signal and then power down the unused MILOs in order to maintain receiver power efficiency.

In order to achieve frequency agility, this clock multiplier design also requires the optimization of individual MILO lock ranges. This requires extensive circuit simulation, which was made possible using the PTC modeling described in Chapter 3. By relating the lock time and lock range to the change in output phase caused by a single period of the injected signal, it becomes possible to efficiently simulate an ILO's behaviour under a variety of injected signal shapes, strengths and frequencies. Using this technique, Chapter 3 describes how changes to the way in which an injected signal is applied to an ILO can have a significant impact on its lock range and shows how this can in turn be used to optimize an ILO for frequency agile applications.

7.1 Suggestions for Future Work

The ability of a wireline receiver to sense and adapt to a variety of operating conditions was explored in this work. In order for these techniques to find wide-spread usefulness

in industrial applications, there is still some work that can be done in order to minimize the overhead required in their implementation.

In particular, the clock multiplier presented here relies on a clock forwarded architecture along with some additional control signal needed to implement the "wake up" transition from zero power to fully operational. Future work could be done to obtain either the clock signal or this "wake up" signal from an existing data channel, thereby reducing the number of pins required and maximizing the aggregate data rate.

Furthermore, although the clock multiplier quickly identifies and powers down unused MILO lanes, the power drawn by the MILO that remains active is still quite high. This is due in large part to the fact that the monitoring circuitry within the active MILO remains on even after the other MILOs have been powered down and no further decisions need to be made. As a result, significant power savings could be achieved by powering down this circuitry when all decisions have been made after 10 cycles of the reference clock.

Along these same lines, the monitoring circuitry used in this clock multiplier uses CML almost exclusively. This decision was made in order to ensure correct operation during variations in the supply voltage caused by sudden changes in the power consumption as parts of the multiplier are powered on and off quickly. If these supply voltage variations could be limited through the use of voltage regulators or through staggered progression of the power on/off transitions then it may be possible to replace several of these CML blocks with CMOS logic, which would translate to a significant reduction in static power.

Future work could also be done to expand the PTC-based ILO model presented in Chapter 3. Although this model has been shown to accurately predict the lock range, lock time and jitter tracking bandwidth of an ILO, the PTC should also be able to determine the amount of deterministic jitter present at the ILO output, as well as the overall phase noise. Obtaining all of this information about the behaviour of an ILO from a single set of simulation results would increase the usefulness of the PTC model but was outside the scope of this thesis.

Appendix A

Summary of Contributions

Chapter	Contribution	Publication
2	<p>AFE adaptation algorithm to optimize the vertical eye opening.</p> <ul style="list-style-type: none"> • Controls gain and equalizer peaking settings using the vertical eye opening and the amplitude noise. • Adds minimal overhead to the receiver. 	[19]
2.2 - 2.3	<p>Variable-gain 15 Gb/s preamplifier with output common mode control.</p> <ul style="list-style-type: none"> • Adjusts power consumption based on gain/linearity requirements. • Variable power helps maintain impedance match for different gain settings. 	[13]
3	<p>A simulation-based ILO model using the proposed phase transfer characteristic (PTC).</p> <ul style="list-style-type: none"> • Easily applicable using simple simulation methods. • Accurate for any kind of injection into any type of ILO. 	submitted to TCAS-I
3.4	<p>Fast power-on MILO with multiple injection points to create a moderate lock range.</p> <ul style="list-style-type: none"> • Able to lock to 2.8 GHz despite variations in voltage or temperature. • Developed as part of a team while on internship at Ramus Inc. 	[21]
4	<p>Multi-MILO clock multiplier with wide lock range and fast power-on.</p> <ul style="list-style-type: none"> • MILO lock range extended using frequency preconditioning. • Unused MILOs identified and powered down within 10 cycles of the reference clock. • Able to lock to different frequencies between power-on cycles with no tuning. 	accepted at CICC

Table A.1: Summary of contributions made along with related locations in the thesis.

Appendix B

Schematic Diagrams




Figure B.1: Schematic of the variable threshold slicer used in the AFE adaptation scheme. Transistor sizes are given in μm with all gate lengths set to minimum size.

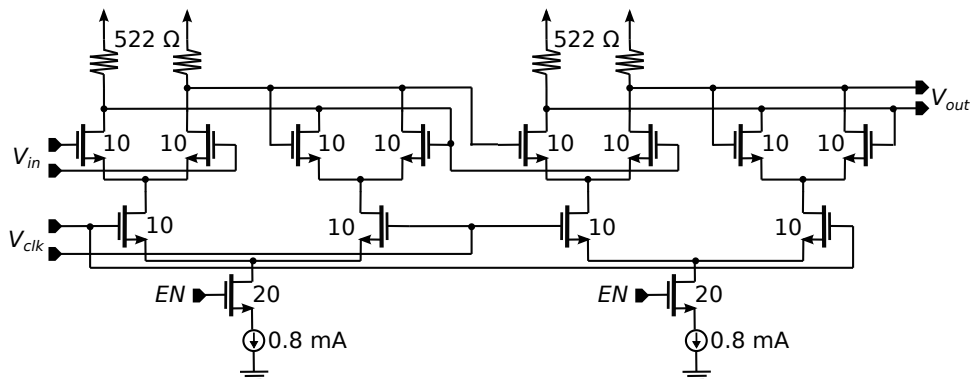


Figure B.2: Schematic of the CML latch used in the TDC and power down logic of the MILO. Transistor sizes are given in μm with all gate lengths set to minimum size.




Figure B.3: Schematic of the CML XOR gate used in the pulse generators and ring oscillators of the MILO. Transistor sizes are given in μm with all gate lengths set to minimum size.




Figure B.4: Schematic of the two-stage, $50\ \Omega$ output driver used to send the MILO output off-chip. Transistor sizes are given in μm with all gate lengths set to minimum size.

Bibliography

- [1] T. Dickson, K. Yau, T. Chalvatzis, A. Mangan, E. Laskin, R. Beerkens, P. Westergaard, M. Tazlauanu, M.-T. Yang, and S. Voinigescu, “The Invariance of Characteristic Current Densities in Nanoscale MOSFETs and Its Impact on Algorithmic Design Methodologies and Design Porting of Si(Ge) (Bi)CMOS High-Speed Building Blocks,” *IEEE J. Solid-State Circuits*, vol. 41, no. 8, pp. 1830–1845, Aug. 2006.
- [2] F. Lu, J. Min, S. Liu, K. Cameron, C. Jones, O. Lee, J. Li, A. Buchwald, S. Jantzi, C. Ward, K. Choi, J. Searle, and H. Samueli, “A Single-Chip Universal Burst Receiver for Cable Modem/Digital Cable-TV Applications,” in *IEEE Custom Integrated Circuits Conference*, 2000, pp. 311–314.
- [3] L. Lehtikoinen and T. Rätty, “Adaptive Real-Time Video Streaming System for Best-Effort IP Networks,” in *International Conference on Systems*, 2009, pp. 131–137.
- [4] B. Leibowitz, R. Palmer, J. Poulton, Y. Frans, S. Li, J. Wilson, M. Bucher, A. Fuller, J. Eyles, M. Aleksic, T. Greer, and N. Nguyen, “A 4.3 GB/s Mobile Memory Interface With Power-Efficient Bandwidth Scaling,” *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 889–898, Apr. 2010.
- [5] T. Ebuchi, Y. Komatsu, M. Miura, T. Chiba, and T. Iwata, “An Ultra-Wide Range Bi-Directional Transceiver With Adaptive Power Control Using Background Replica VCO Gain Calibration,” *IEEE J. Solid-State Circuits*, vol. 46, no. 4, pp. 986–991, 2011.

- [6] K. Kaviani, T. Wu, J. Wei, A. Amirkhany, J. Shen, T. J. Chin, C. Thakkar, S. Member, W. T. Beyene, S. Member, N. Chan, C. Chen, B. R. Chuang, D. Dressler, V. P. Gadde, M. Hekmat, E. Ho, C. Huang, P. Le, C. Madden, N. K. Mishra, L. Raghavan, K. Saito, R. Schmitt, D. Secker, X. Shi, S. Fazeel, G. S. Srinivas, S. Zhang, and C. Tran, "A Tri-Modal 20-Gbps/Link Differential/DDR3/GDDR5 Memory Interface," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 926–937, 2012.
- [7] J. Zerbe, C. Werner, V. Stojanovic, F. Chen, A. Ho, T. Thrush, R. Kollipara, M. Horowitz, and K. Donnelly, "Equalization and Clock Recovery for a 2.5-10-Gb/s 2-PAM/4-PAM Backplane Transceiver Cell," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2121–2130, Dec. 2003.
- [8] B. Garlepp, A. Ho, V. Stojanovic, F. Chen, C. Werner, G. Tsang, T. Thrush, A. Agarwal, and J. Zerbe, "A 1-10Gbps PAM2, PAM4, PAM2 Partial Response Receiver Analog Front End with Dynamic Sampler Swapping Capability for Backplane Serial Communications," in *IEEE Symposium on VLSI Circuits*, 2005, pp. 376–379.
- [9] M. Hossain and A. Chan Carusone, "510 Gb/s 70 mW Burst Mode AC Coupled Receiver in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 3, pp. 524–537, 2010.
- [10] M. Ramezani, M. Abdalla, A. Shoval, M. van Ierssel, A. Rezayee, A. McLaren, C. Holdenried, J. Pham, E. So, D. Cassan, and S. Sadr, "An 8.4 mW/Gb/s 4-lane 48Gb/s multi-standard-compliant transceiver in 40nm digital CMOS technology," in *IEEE J. Solid-State Circuits*, 2011, p. 2011.
- [11] F. Gerfers, G. W. D. Besten, P. V. Petkov, J. E. Conder, and A. J. Koellmann, "A 0.22 Gb/s 6x OSR Receiver Using a Digitally Self-Adaptive Equalizer," *IEEE J. Solid-State Circuits*, vol. 43, no. 6, pp. 1436–1448, 2008.

- [12] P. Park and A. Chan Carusone, "A 20-Gb/s Coaxial Cable Receiver Analog Front-End in 90-nm CMOS Technology," in *IEEE Asian Solid-State Circuits Conference*, 2008, pp. 225–228.
- [13] D. Dunwell and A. Chan Carusone, "A 15-Gb/s Preamplifier with 10-dB Gain Control and 8-mV Sensitivity in 65-nm CMOS," in *IEEE International Symposium on Circuits and Systems*, 2010, pp. 205–208.
- [14] J. R. Miller, Y. Li, K. Hinckley, G. Blando, B. Guenin, A. Dengi, A. Rebelo, and S. McMorrow, "Temperature and Moisture Dependence of PCB and Package Traces and the Impact on Signal Performance," in *DesignCon 2012*, 2012.
- [15] J. Choi, M. Hwang, and D. Jeong, "A 0.18-um CMOS 3.5-Gb/s Continuous-Time Adaptive Cable Equalizer Using Enhanced Low-Frequency Gain Control Method," *IEEE J. Solid-State Circuits*, vol. 39, no. 3, pp. 419–425, 2004.
- [16] F. Buchali, S. Lanne, J. Thiery, W. Baumert, and H. Bulow, "Fast Eye Monitor for 10 Gbit/s and its Application for Optical PMD Compensation," in *Optical Fiber Communication Conference and Exhibit*, vol. 2, 2001, pp. 5–1 – 5–3.
- [17] T. Ellermeyer, U. Langmann, B. Wedding, and W. Pohlmann, "A 10-Gb/s Eye-Opening Monitor IC for Decision-Guided Adaptation of the Frequency Response of an Optical Receiver," *IEEE J. Solid-State Circuits*, vol. 35, no. 12, pp. 1958–1963, 2000.
- [18] B. Analui, A. Rylyakov, S. Rylov, M. Meghelli, and A. Hajimiri, "A 10-Gb/s Two-Dimensional Eye-Opening Monitor in 0.13-um Standard CMOS," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2689–2699, Dec. 2005.
- [19] D. Dunwell and A. Chan Carusone, "Gain and Equalization Adaptation to Optimize the Vertical Eye Opening in a Wireline Receiver," in *IEEE Custom Integrated Circuits Conference*, 2010, pp. 1–4.

- [20] L. A. Barroso and U. Hölzle, “The Case for Energy-Proportional Computing,” *IEEE Computer Magazine*, vol. 40, no. 12, pp. 33–37, 2007.
- [21] J. Zerbe, B. Daly, W. Dettloff, T. Stone, W. Stonecypher, P. Venkatesan, K. Prabhu, B. Su, J. Ren, B. Tsang, B. Leibowitz, D. Dunwell, A. Chan Carusone, and J. Eble, “A 5.6Gb/s 2.4mW/Gb/s Bidirectional Link With 8ns Power-On,” in *IEEE Symposium on VLSI Circuits*, 2011, pp. 82–83.
- [22] J. Lee, H. Wang, W. Chen, and Y. Lee, “Subharmonically Injection-Locked PLLs for Ultra-Low-Noise Clock Generation,” in *IEEE International Solid-State Circuits Conference*, 2009, pp. 92–93.
- [23] R. Adler, “A Study of Locking Phenomena in Oscillators,” *Proceedings of the IRE*, vol. 34, no. 6, pp. 351–357, 1946.
- [24] M. Hossain and A. Chan Carusone, “CMOS Oscillators for Clock Distribution and Injection-Locked Deskew,” *IEEE J. Solid-State Circuits*, vol. 44, no. 8, pp. 2138–2153, Aug. 2009.
- [25] M. Mansuri, F. O’Mahony, G. Balamurugan, J. Jaussi, J. Kennedy, S. Shekhar, R. Mooney, and B. Casper, “Strong Injection Locking of Low-Q LC Oscillators,” in *IEEE Custom Integrated Circuits Conference*, 2008, pp. 699–702.
- [26] A. Mirzaei, M. Heidari, R. Bagheri, S. Chehrazi, and A. Abidi, “Injection-Locked Frequency Dividers based on Ring Oscillators with Optimum Injection for Wide Lock Range,” in *IEEE Symposium on VLSI Circuits*, 2006, pp. 174–175.
- [27] A. Hajimiri and T. Lee, “A General Theory of Phase Noise in Electrical Oscillators,” *IEEE J. Solid-State Circuits*, vol. 33, no. 2, pp. 179–194, 1998.
- [28] B. Razavi, *Design of Analog CMOS Integrated Circuits*. McGraw-Hill, 2001.

- [29] O. Agazzi, M. Hueda, D. Crivelli, H. Carrer, A. Nazemi, G. Luna, F. Ramos, R. López, C. Grace, B. Kobeissy, C. Abidin, M. Kazemi, M. Kargar, C. Marquez, S. Ramprasad, F. Bollo, V. Posse, S. Wang, G. Asmanis, G. Eaton, N. Swenson, T. Lindsay, and P. Voois, “A 90 nm CMOS DSP MLSD Transceiver With Integrated AFE for Electronic Dispersion Compensation of Multimode Optical Fibers at 10 Gb/s,” *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2939–2957, 2008.
- [30] J. Cao, B. Zhang, U. Singh, D. Cui, A. Vasani, A. Garg, W. Zhang, N. Kocaman, D. Pi, B. Raghavan, H. Pan, I. Fujimori, and A. Momtaz, “A 500mW Digitally Calibrated AFE in 65nm CMOS for 10Gb/s Serial Links over Backplane and Multimode Fiber,” in *IEEE International Solid-State Circuits Conference*, 2009, pp. 370–371.
- [31] R. G. Meyer and W. D. Mack, “A Wideband Low-Noise Variable-Gain BiCMOS Transimpedance Amplifier,” *IEEE J. Solid-State Circuits*, vol. 29, no. 6, pp. 701–706, 1994.
- [32] E. Sackinger, *Broadband Circuits for Optical Fiber Communication*. Hoboken, NJ, USA: John Wiley and Sons, Inc., Mar. 2005.
- [33] C. Liao and S. Liu, “40 Gb/s Transimpedance-AGC Amplifier and CDR Circuit for Broadband Data Receivers in 90 nm CMOS,” *IEEE J. Solid-State Circuits*, vol. 43, no. 3, pp. 642–655, Mar. 2008.
- [34] S. Ibrahim and B. Razavi, “Low-Power CMOS Equalizer Design for 20-Gb/s Systems,” *IEEE J. Solid-State Circuits*, vol. 46, no. 6, pp. 1321–1336, June 2011.
- [35] F. O’Mahony, J. Jaussi, J. Kennedy, G. Balamurugan, M. Mansuri, C. Roberts, S. Shekhar, R. Mooney, and B. Casper, “A 47 x 10 Gb/s 1.4 mW/Gb/s Parallel Interface in 45 nm CMOS,” *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2828–2837, 2010.

- [36] A. Agrawal, J. Bulzacchelli, T. Dickson, Y. Liu, J. Tierno, and D. Friedman, "A 19Gb / s Serial Link Receiver with Both 4-Tap FFE and 5-Tap DFE Functions in 45nm SOI CMOS," in *IEEE International Solid-State Circuits Conference*, 2012, pp. 134–136.
- [37] G. Zhang and M. Green, "A 10 Gb/s BiCMOS Adaptive Cable Equalizer," *IEEE J. Solid-State Circuits*, vol. 40, no. 11, pp. 2132–2140, Nov. 2005.
- [38] A. Baker, "An Adaptive Cable Equalizer for Serial Digital Video Rates to 400 Mb/s," in *IEEE International Solid-State Circuits Conference*, 1996, pp. 174–175.
- [39] E. Chen, J. Ren, B. Leibowitz, H. Lee, Q. Lin, K. Oh, F. Lambrecht, V. Stojanovic, J. Zerbe, and C. Yang, "Near-Optimal Equalizer and Timing Adaptation for I/O Links Using a BER-Based Metric," *IEEE J. Solid-State Circuits*, vol. 43, no. 9, pp. 2144–2156, 2008.
- [40] K. Cheng and H. Chang, "Test Strategies for Adaptive Equalizers," *IEEE Custom Integrated Circuits Conference*, pp. 597–604, Sept. 2009.
- [41] F. Musa and A. Chan Carusone, "Modeling and Design of Multilevel BangBang CDRs in the Presence of ISI and Noise," *IEEE Trans. Circuits Syst. I*, vol. 54, no. 10, pp. 2137–2147, Oct. 2007.
- [42] H. Noguchi, N. Yoshida, H. Uchida, M. Ozaki, S. Kanemitsu, and S. Wada, "A 40-Gb/s CDR Circuit With Adaptive Decision-Point Control Based on Eye-Opening Monitor Feedback," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2929–2938, 2008.
- [43] H. Cheng and A. Chan Carusone, "A 32/16 Gb/s 4/2-PAM Transmitter with PWM Pre-Emphasis and 1.2 V_{pp} per side Output Swing in 0.13- μ m CMOS," in *IEEE Custom Integrated Circuits Conference*, 2008, pp. 635–638.

- [44] H. Hwang, J. Chien, T. Chen, and L. Lu, "A CMOS Tunable Transimpedance Amplifier," *IEEE Microwave and Wireless Components Letters*, vol. 16, no. 12, pp. 693–695, Dec. 2006.
- [45] I.-T. Lee and S.-I. Liu, "G-Band Injection-Locked Frequency Dividers Using pi-type LC Networks," *IEEE Trans. Circuits Syst. I*, vol. 59, no. 2, pp. 315–323, 2012.
- [46] L. Liu, T. Sakurai, and M. Takamiya, "315MHz Energy-Efficient Injection-Locked OOK Transmitter and 8.4uW Power-Gated Receiver Front-End for Wireless Ad Hoc Network in 40nm CMOS," in *IEEE Symposium on VLSI Circuits*, 2011, pp. 278–279.
- [47] G. Beers, "A Frequency-Dividing Locked-In Oscillator Frequency-Modulation Receiver*," *Proceedings of the IRE*, vol. 32, no. 12, pp. 730–737, 1944.
- [48] B. Helal, C. Hsu, K. Johnson, and M. Perrott, "A Low Jitter Programmable Clock Multiplier Based on a Pulse Injection-Locked Oscillator With a Highly-Digital Tuning Loop," *IEEE J. Solid-State Circuits*, vol. 44, no. 5, pp. 1391–1400, May 2009.
- [49] L. Paciorek, "Injection Locking of Oscillators," *Proc. IEEE*, vol. 53, no. 11, pp. 1723–1727, Sept. 1965.
- [50] S. Shekhar, M. Mansuri, F. O. Mahony, G. Balamurugan, J. E. Jaussi, J. Kennedy, D. J. Allstot, R. Mooney, and B. Casper, "Strong Injection Locking in Low-Q LC Oscillators: Modeling and Application in a Forwarded-Clock I/O Receiver," *IEEE Trans. Circuits Syst. I*, vol. 56, no. 8, pp. 1818–1829, 2009.
- [51] A. Mirzaei, M. Heidari, and A. Abidi, "Analysis of Oscillators Locked by Large Injection Signals: Generalized Adler's Equation and Geometrical Interpretation," *IEEE Custom Integrated Circuits Conference*, pp. 737–740, Sept. 2006.
- [52] D. Leeson, "A simple model of feedback oscillator noise spectrum," *Proc. IEEE*, vol. 54, no. 2, pp. 329–330, 1966.

- [53] X. Lai and J. Roychowdhury, "Capturing Oscillator Injection Locking via Nonlinear Phase-Domain Macromodels," *IEEE Trans. Microwave Theory and Techniques*, vol. 52, no. 9, pp. 2251–2261, Sept. 2004.
- [54] P. Maffezzoni, "Analysis of Oscillator Injection Locking Through Phase-Domain Impulse-Response," *IEEE Trans. Circuits Syst. I*, vol. 55, no. 5, pp. 1297–1305, June 2008.
- [55] R. Harjani and N. Lanka, "High Speed Frequency Hopping Using Injection Locked Front-Ends," *Asilomar Conference on Signals, Systems and Computers*, pp. 1336–1340, Nov. 2007.
- [56] J. Lee and M. Liu, "A 20-Gb/s Burst-Mode Clock and Data Recovery Circuit Using Injection-Locking Technique," *IEEE J. Solid-State Circuits*, vol. 43, no. 3, pp. 619–630, Mar. 2008.
- [57] N. Lanka, S. Patnaik, and R. Harjani, "Understanding the Transient Behavior of Injection Locked LC Oscillators," in *IEEE Custom Integrated Circuits Conf.*, 2007, pp. 667–670.
- [58] J. Chien and L. Lu, "Analysis and Design of Wideband Injection-Locked Ring Oscillators With Multiple-Input Injection," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1906–1915, Sept. 2007.
- [59] J. Luo, N. K. Jha, and L.-s. Peh, "Simultaneous Dynamic Voltage Scaling of Processors and Communication Links in Real-Time Distributed Embedded Systems," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 15, no. 4, pp. 427–437, 2007.
- [60] O. Richard, A. Siligaris, F. Badets, C. Dehos, C. Dufis, P. Busson, P. Vincent, D. Belot, and P. Urard, "A 17.5-to-20.94GHz and 35-to-41.88GHz PLL in 65nm

- CMOS for Wireless HD Applications,” in *IEEE International Solid-State Circuits Conference*, 2010, pp. 252–253.
- [61] F. Lin, R. a. Royer, B. Johnson, and B. Keeth, “A Wide-Range Mixed-Mode DLL for a Combination 512 Mb 2.0 Gb/s/pin GDDR3 and 2.5 Gb/s/pin GDDR4 SDRAM,” *IEEE J. Solid-State Circuits*, vol. 43, no. 3, pp. 631–641, Mar. 2008.
- [62] M. Kossel, T. Morf, W. Baumberger, A. Biber, C. Menolfi, T. Toifl, and M. Schmatz, “A Multiphase PLL for 10 Gb/s Links in SOI CMOS Technology,” in *IEEE Radio Frequency Integrated Circuits Symposium*, 2004, pp. 207–210.