

# Electrical and Optical Interconnects for High-Performance Computing


Thesis by

Meisam Honarvar Nazari

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy



California Institute of Technology

Pasadena, California

2013

(Defended April 17, 2013)

© 2013

Meisam Honarvar Nazari

All Rights Reserved

To my parents and wife

# Acknowledgements

During the course of my graduate studies at Caltech, I was able to experience an enormous amount of educational, professional, and personal growth due to the guidance of many good teachers and support of many friends. This work would not have been possible without either of these.

I have been very lucky to have had the opportunity to work with an excellent research advisor, Prof. Azita Emami-Neyestanak. Her technical expertise, teaching skills, and the unique ability to motivate me through challenging stages of my research have played a key role in my success. Every single meeting with Azita helped me to find my path and go another step forward in my research, and motivated me to learn more. Her enthusiasm and care were the driving forces for the creation of new ideas in this work. She also helped me to improve my public speaking and writing skills. I sincerely thank her for her kindness and patience.

I would also like to especially thank some of the key faculty and researchers that I had the pleasure of working with at Caltech. I am very grateful to my candidacy and oral defense committee, Prof. Ali Hajimiri, Prof. Axel Scherer, Prof. David Rutledge, Prof. Sander Weinreb, and Prof. Hyuck Choo. Particularly, I gratefully acknowledge Prof. Ali Hajimiri and Sander Weinreb for their support in providing me with access to their labs for experiments. I also thank Dr. Firouz Aflatooni for expanding my knowledge of optics and integrated photonics.

The faculty of the Department of Electrical Engineering at Caltech are undoubtedly among the most brilliant teachers in the world and I had the opportunity of learning from many of them. Particularly, I would like to express my gratitude to Professor Ali Hajimiri,

Prof. P. P. Vaidyanathan and Prof. Yaser Abu-Mostafa.

The highly academic environment at the University of Tehran helped me to build the required background in engineering and inspired me to pursue my graduate studies abroad. I sincerely thank Prof. Parviz Jabehdar-Maralani, Prof. Jalil Rashed, and Prof. Shams Mohajerzadeh, for their advice, encouragements and remarkable teaching. The strong knowledge base that I acquired during my Masters studies under supervision of Prof Roman Genov at the University of Toronto greatly prepared me for my Ph.D. research. I would like to thank Roman for his great support and patience in training me.

I have truly enjoyed and greatly benefited from the interaction with an amazing group of research colleagues at Caltech. I would like to thank my friends and colleagues Matthew Loh, Juhwan Yoo, Manuel Monge, Saman Saeedi, Mayank Raj, Krishna Sataluri, Kaveh Hosseini, and Laleh Rabieirad in Azita's group for their technical support. I extend my gratitude to my friends Amir Safaripour, Kaushik Sengupta, Kaushik Dasgupta, Steven Bowers, Alex Pai, Behrooz Abiri, and Constantine Sideris in Prof. Hajimiri's group.

I would like to thank my friends Peyman Tavalali, Masoud Farivar, Saeid Farivar, Amin Khajehnejad, Hamed Hamze, Ali Vakili and Sormeh Shadbakht for making Caltech a fun place to live and work.

I wish to thank the National Science Foundation, C2S2, and IFC for their financial support. Also, donated resources from ST Microsystems, CMP, and Cosemi Technologies enabled this research.

I sincerely thank my brothers, Mehdi and Masoud Honarvar Nazari and my sister, Narges Honarvar Nazari for their support that eased the hardship of being away from home. I also would like to thank my other brother and sister (in-law) Amin-Qasem Safarian and Neda Bozorgkhan for providing a wholesome environment and making me feel at home.

My parents have been a source of encouragement and support throughout my life. I would like to thank my mother Fatemeh Tahmasebi and father Mohammad Honarvar Nazari for their unconditional love, strength and sacrifices. I also thank my other parents (in-law), Fatemeh Hashemi and Akbar Safarian, for their unwavering love and support. Finally, I am forever indebted to my wife Zahra Safarian whose immeasurable love has supported me

throughout this journey. For that, I dedicate this thesis to them.

# Abstract

Technology scaling has enabled drastic growth in the computational and storage capacity of integrated circuits (ICs). This constant growth drives an increasing demand for high-bandwidth communication between and within ICs. In this dissertation we focus on low-power solutions that address this demand. We divide communication links into three sub-categories depending on the communication distance. Each category has a different set of challenges and requirements and is affected by CMOS technology scaling in a different manner. We start with short-range chip-to-chip links for board-level communication. Next we will discuss board-to-board links, which demand a longer communication range. Finally on-chip links with communication ranges of a few millimeters are discussed.

Electrical signaling is a natural choice for chip-to-chip communication due to efficient integration and low cost. IO data rates have increased to the point where electrical signaling is now limited by the channel bandwidth. In order to achieve multi-Gb/s data rates, complex designs that equalize the channel are necessary. In addition, a high level of parallelism is central to sustaining bandwidth growth. Decision feedback equalization (DFE) is one of the most commonly employed techniques to overcome the limited bandwidth problem of the electrical channels. A linear and low-power summer is the central block of a DFE. Conventional approaches employ current-mode techniques to implement the summer, which require high power consumption. In order to achieve low-power operation we propose performing the summation in the charge domain. This approach enables a low-power and compact realization of the DFE as well as crosstalk cancellation. A prototype receiver was fabricated in 45nm SOI CMOS to validate the functionality of the proposed technique and was tested over channels with different levels of loss and coupling. Measurement results show that the

receiver can equalize channels with maximum 21dB loss while consuming about 7.5mW from a 1.2V supply. We also introduce a compact, low-power transmitter employing passive equalization. The efficacy of the proposed technique is demonstrated through implementation of a prototype in 65nm CMOS. The design achieves up to 20Gb/s data rate while consuming less than 10mW.

An alternative to electrical signaling is to employ optical signaling for chip-to-chip interconnections, which offers low channel loss and cross-talk while providing high communication bandwidth. In this work we demonstrate the possibility of building compact and low-power optical receivers. A novel RC front-end is proposed that combines dynamic offset modulation and double-sampling techniques to eliminate the need for a short time constant at the input of the receiver. Unlike conventional designs, this receiver does not require a high-gain stage that runs at the data rate, making it suitable for low-power implementations. In addition, it allows time-division multiplexing to support very high data rates. A prototype was implemented in 65nm CMOS and achieved up to 24Gb/s with less than 0.4pJ/b power efficiency per channel. As the proposed design mainly employs digital blocks, it benefits greatly from technology scaling in terms of power and area saving.

As the technology scales, the number of transistors on the chip grows. This necessitates a corresponding increase in the bandwidth of the on-chip wires. In this dissertation, we take a close look at wire scaling and investigate its effect on wire performance metrics. We explore a novel on-chip communication link based on a double-sampling architecture and dynamic offset modulation technique that enables low power consumption and high data rates while achieving high bandwidth density in 28nm CMOS technology. The functionality of the link is demonstrated using different length minimum-pitch on-chip wires. Measurement results show that the link achieves up to 20Gb/s of data rate (12.5Gb/s/ $\mu\text{m}$ ) with better than 136fJ/b of power efficiency.



# Contents

<b>Acknowledgements</b>	<b>iv</b>
<b>Abstract</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Organization . . . . .	6
<b>2 Electrical Interconnects Background</b>	<b>8</b>
2.1 Metrics of Electrical Links . . . . .	9
2.1.1 Bit-Rate . . . . .	9
2.1.2 Bit-Error Rate . . . . .	10
2.1.2.1 Amplitude Noise . . . . .	13
2.1.2.2 Timing Noise . . . . .	15
2.2 Basics of Electrical Links . . . . .	15
2.2.1 Channel . . . . .	16
2.2.1.1 Inter Symbol Interference . . . . .	17
2.2.1.2 Crosstalk . . . . .	20
2.2.2 Transmitter . . . . .	23
2.2.2.1 Transmitter Equalization . . . . .	27
2.2.3 Receiver . . . . .	31
2.2.3.1 Receiver Equalization . . . . .	32
2.2.4 Timing Recovery . . . . .	37
2.2.5 Advanced Modulation Techniques . . . . .	39

2.3	Summary . . . . .	41
<b>3</b>	<b>High-Speed Low-Power Electrical Transceivers</b>	<b>43</b>
3.1	Receiver Architecture . . . . .	44
3.2	Decision Feedback Equalization Implementation . . . . .	46
3.2.1	Current-Mode Summers . . . . .	47
3.2.2	Charge-Mode Summers . . . . .	48
3.2.2.1	4-tap Realization . . . . .	50
3.2.2.2	Generalization of Switched Capacitor Summers . . . . .	57
3.2.3	Comparator Design . . . . .	61
3.3	Far-End Crosstalk Cancellation . . . . .	62
3.4	Experimental Results . . . . .	70
3.5	Transmitter Design . . . . .	77
3.5.1	Transmitter Architecture . . . . .	77
3.5.2	Experimental Results . . . . .	81
3.6	Summary . . . . .	85
<b>4</b>	<b>Overview of High-Speed Optical Links</b>	<b>88</b>
4.1	Optical Channels . . . . .	90
4.2	Optical Receivers . . . . .	92
4.2.1	Photodiodes . . . . .	92
4.2.2	Receiver Front-End . . . . .	93
4.2.2.1	Transimpedance Amplifiers . . . . .	94
4.2.2.2	Integrating Front-Ends . . . . .	100
4.3	Optical Transmitters . . . . .	105
4.3.1	Vertical Cavity Surface Emitting Laser . . . . .	105
4.3.2	Mach-Zehnder Modulator . . . . .	106
4.3.3	Ring Resonator . . . . .	108
4.4	Summary . . . . .	109

<b>5</b>	<b>Low-Power Optical Receiver Design</b>	<b>111</b>
5.1	Low-Power Double-Sampling RC Front-End . . . . .	112
5.1.1	Front-End Sensitivity Analysis and Implementation . . . . .	119
5.2	System-Level Design Considerations . . . . .	126
5.2.1	Adaptation of Dynamic Offset Modulation . . . . .	126
5.2.2	Double-Sampling Front-End Scaling . . . . .	130
5.2.2.1	Sensitivity . . . . .	131
5.2.2.2	Data Rate . . . . .	132
5.2.2.3	Power Consumption . . . . .	132
5.2.2.4	Dynamic Range . . . . .	133
5.2.3	Photodiode Capacitance Scaling . . . . .	134
5.2.4	Clocking . . . . .	136
5.3	Experimental Results . . . . .	140
5.4	Summary . . . . .	146
<b>6</b>	<b>On-Chip Wires: Characteristics, Models, and scaling</b>	<b>148</b>
6.1	Wire Characteristics . . . . .	149
6.1.1	Resistance . . . . .	150
6.1.2	Capacitance . . . . .	152
6.1.3	Inductance . . . . .	153
6.2	Wire Performance Metrics . . . . .	153
6.2.1	Delay . . . . .	154
6.2.2	Crosstalk . . . . .	156
6.2.3	Power . . . . .	160
6.3	Repeaters . . . . .	160
6.4	Wire Scaling . . . . .	169
6.5	Summary . . . . .	171

<b>7</b>	<b>On-Chip Interconnects</b>	<b>173</b>
7.1	On-Chip Communication Power Trend . . . . .	173
7.2	Prior Art in Design of On-Chip Links . . . . .	176
7.2.1	Low Voltage Signaling . . . . .	177
7.2.2	Current-Mode Drivers . . . . .	179
7.2.3	On-Chip Transmission Line . . . . .	179
7.2.4	Equalization . . . . .	181
7.3	Double-Sampling Link . . . . .	182
7.3.1	Receiver Design . . . . .	182
7.3.2	Transmitter Design . . . . .	183
7.4	Link Latency . . . . .	187
7.5	Circuit Implementation . . . . .	188
7.6	Experimental Results . . . . .	190
7.7	Summary . . . . .	195
<b>8</b>	<b>Conclusions</b>	<b>197</b>
8.1	Electrical Interconnects . . . . .	197
8.2	Optical Link Performance Summary . . . . .	199
8.3	On-Chip Interconnects . . . . .	201
	<b>References</b>	<b>202</b>
	<b>Bibliography</b>	<b>202</b>

# List of Figures

1.1	Examples of complex electronic systems. A high-performance multi-core processor (a). A computer server (b).	2
1.2	History of microprocessor performance scaling.	3
1.3	IO bandwidth requirement of microprocessors in recent years (a). Constant growth of the required IO bandwidth according to ITRS (b).	3
2.1	Fan-out of four inverter chain	10
2.2	FO-4 delay metric for different technology nodes (a), and $f_T^{-1}$ versus FO4 metric, which shows a linear relation (b).	11
2.3	Illustration of a data eye. (a) shows a noisy data stream, and (b) shows the folding of the data stream into a data eye.	11
2.4	Effect of amplitude noise in (a) and timing noise in (b) on the eye diagram quality.	12
2.5	Translation of the eye-diagram to bathtub curve as a tool to measure signal integrity.	13
2.6	Diagram of bit error rate versus signal to noise ratio.	14
2.7	Components of an electrical communication link.	16
2.8	A typical backplane link and its components [22].	16
2.9	Transfer characteristics of typical backplane channels with and without via stubs at different lengths [22].	17
2.10	Pulse response of a typical bandwidth-limited backplane channel illustrating dispersion.	18
2.11	Channel dispersion resulting in pre- and post-cursor ISI.	19

2.12	Far-end and near-end crosstalk in coupled transmission lines. . . . .	20
2.13	Lumped equivalent model for coupled transmission lines. . . . .	21
2.14	Effect of coupling on even and odd modes, which results in data dependent jitter. . . . .	24
2.15	Effect of coupling on the victim signal amplitude when the transmitted aggressor and victim signals are not synchronized. . . . .	24
2.16	Voltage-mode and current-mode transmitter design. . . . .	25
2.17	Different transmitter examples, (a) voltage-mode, (b) single-ended current mode, (c) differential current-mode. . . . .	26
2.18	Transmitter pre-emphasis using high-frequency boosting. . . . .	27
2.19	Block diagram of a transmitter with m-tap FIR-based equalization. . . . .	28
2.20	Pulse response of a channel before and after 2-tap transmitter FIR equalization. . . . .	29
2.21	Eye diagram at the end of a lossy channel using a transmitter with no tap of FIR equalization as well as 1-tap and 2-tap of equalization. . . . .	30
2.22	Transmitter equalization through high frequency boosting to achieve flat response. . . . .	30
2.23	Different high-frequency boosting configurations. . . . .	31
2.24	A typical electrical receiver block diagram. . . . .	32
2.25	Simplified block diagram of a decision feedback equalizer. . . . .	33
2.26	Loop-unrolling of the first tap of the DFE to remove the critical feedback loop. . . . .	34
2.27	Feed-forward equalization at the receiver. . . . .	35
2.28	Continuous-time receiver equalization using frequency dependent source degeneration. . . . .	36
2.29	Passive high-pass filter equalizer (a). Dual path continuous-time equalizer (b). . . . .	36
2.30	Clock recovery at the receiver. . . . .	38
2.31	CDR phase detectors: (a) linear [51], (b) binary [52] . . . . .	39
2.32	Simple binary (1bit/symbol) and PAM-4 (2bits/symbol) modulations. . . . .	40
3.1	A typical parallel communication link. Signal integrity is degraded due to channel dispersion and crosstalk noise of the neighboring channels. . . . .	45

3.2	Top level architecture of the proposed receiver employing half-rate clocking and loop-unrolling. . . . .	46
3.3	Current-mode summer employed in conventional DFE (a). Current-integrating summer (b). Operation of the current-integrating summer (c). . . . .	49
3.4	(a) Circuit-level implementation of the front-end S/H/summer operating in two phases: (b) sample/sum phase, (c) sum/hold phase (single-ended version is shown for simplicity). . . . .	51
3.5	S/H/summer performance in the case of $n=2$ . S/H/summer voltage loss, $A_{summer}$ (a). SNR at the output of the S/H/summer (b). SNR normalized to clocking power consumption (c). To achieve high SNR (hence high BER) and power efficiency while maintaining the DFE speed, $C_{S1}$ and $C_{S2}$ are chosen to be equal to 19fF and 14fF, respectively. . . . .	54
3.6	Simulated input eye diagram for a 5" FR-4 trace at 15Gb/s. Red and green circles show the sampled input before and after 2-tap decision feedback equalization, respectively (a). Histogram of the sampled input before and after equalization. 2-tap DFE improves eye opening from 2mV to 40mV (b). . . . .	56
3.7	4-bit current steering DAC that generates DFE tap coefficients. . . . .	56
3.8	Switched capacitor summer for 2n-tap, sample/sum phase (a), sum/hold phase (b). . . . .	57
3.9	Signal gain comparison between SC and current-mode summer (a). Normalized SNR for the SC and current-mode summer (b). Tap coefficient gain for different number of taps (c). . . . .	59
3.10	Linearity of the post-cursor taps for an eight-tap (a) switched capacitor and (b) current-mode summer. . . . .	60
3.11	2-tap DFE architecture with loop-unrolling. . . . .	62
3.12	Combined analog MUX and latch with cross-coupled capacitors to reduce kick-back. . . . .	63
3.13	Crosstalk cancellation technique employing a high-pass filter as a differentiator to emulate FEXT. . . . .	63

3.14	Measured transfer characteristics of a 5" long, 32mil wide coupled trace with 40mil separation (a). Simulated FEXT noise due to a 15Gb/s pulse and the emulated FEXT employing the differentiator along with the residual FEXT noise (b). . . . .	64
3.15	Simulated output of the S/H/summer before and after applying the aggressor as well as when the crosstalk cancellation is enabled. . . . .	65
3.16	Crosstalk cancellation technique for multiple coupled channels. . . . .	67
3.17	Simulated channel, adjacent FEXT, and distant FEXT response for width=32mil, spacing=48mil, and length=10", (d) for length=20". (e) Residual crosstalk from the adjacent aggressor after cancellation normalized to the FEXT energy along with the distant FEXT for 10" channel, and (f) 20" channel.. . . .	68
3.18	Effect of loading due to the crosstalk cancellation circuitry on the overall channel insertion loss for a 10" channel (a) and a 20" channel (b). . . . .	69
3.19	The die micrograph of the receiver with major blocks highlighted. . . . .	71
3.20	Receiver DFE and crosstalk cancellation test set-up. . . . .	71
3.21	CMOS clock generation through CML-to-CMOS conversion followed by duty cycle correction. . . . .	72
3.22	Channel transfer characteristics for 5", 10" and 18" PCB traces. . . . .	72
3.23	The PRBS7 eye diagram at the receiver input and the bathtub curve after equalization for, (a) 11Gb/s data over 18" trace, (b) 13Gb/s data over 10" trace, (c) 15Gbs data over 5" trace. . . . .	74
3.24	Receiver power breakdown. . . . .	75
3.25	The receiver bathtub curve without and with crosstalk noise, and after crosstalk cancellation for, (a) 8Gb/s, (b) 10Gb/s, (c) 11Gb/s, and (d) 12.5Gb/s victim and aggressor data. . . . .	76
3.26	Shunt and double-series bandwidth enhancement technique, which requires three inductors. A T-coil can also perform the same functionality. . . . .	78
3.27	Segmented T-coil layout 3D and top views. The five top metals are employed. . . . .	78
3.28	Segmented T-coil lumped model generated using IE3D electromagnetic simulator. . . . .	79



3.29	Transistor-level schematic of the transmitter employing segmented T-coils. . .	80
3.30	Simulation results showing the programmable high frequency peaking achieved by RC source degeneration. . . . .	81
3.31	Simulated transfer characteristics of the transmitter and channel for different levels of pre-emphasis, (a) 5" FR4 channel, (b) 10" FR4 channel. . . . .	82
3.32	Transmitter die micrograph along with the core layout. . . . .	82
3.33	Transmitter measurement setup for characterizing the performance over different channels . . . . .	82
3.34	20Gb/s on-chip PRBS-7 generator. . . . .	83
3.35	Channels' transfer characteristic. . . . .	84
3.36	The transmitted 10Gb/s PRBS7 data over 5" FR4 channel with about 10dB loss at Nyquist, before equalization (a), over equalized (b) and optimally equalized (c). . . . .	84
3.37	Transmitter output at 15Gb/s over lossy channel with 7dB loss. . . . .	85
3.38	Output of the channel before and after equalization with the 2-tap FIR equalizer and continuous-time equalizer for 15Gb/s data over 5" channel (a), 15Gb/s data over 10" channel (b), and 20Gb/s data over lossy channel (c). . . . .	86
4.1	Optical signal transmission over fiber. . . . .	89
4.2	An optical fiber cross-section with the core and cladding having refractive index of $n_1$ and $n_2$ , respectively to allow for total internal reflection. . . . .	90
4.3	Cross section of a polymer waveguide. . . . .	92
4.4	Electrical model of a photodiode. . . . .	93
4.5	Simple resistive receiver front-end performing current to voltage conversion. .	94
4.6	Schematic of a common-gate TIA (a) and a regulated cascode TIA (b). . . . .	95
4.7	Typical shunt-shunt feedback TIA (a). Simplified small signal model of the TIA (b). . . . .	96
4.8	Schematic of a common source TIA. . . . .	98
4.9	Inverter-based TIA with one (a) and three (b) stages. . . . .	99

4.10	TIA data rate (a) and power consumption (b) as a function of transimpedance.	99
4.11	Sensitivity degradation as a result of data rate scaling.	100
4.12	Integrating optical front-end employing balanced photodiode and an inverter.	101
4.13	Integrate and reset front-end.	102
4.14	Double-sampling front-end with a DC current to provide bipolar voltage difference for a one and a zero.	103
4.15	Double-sampling integrating front-end.	104
4.16	Input demultiplexing receiver using multiple sampler clock phases.	104
4.17	Typical current-mode VCSEL driver	106
4.18	A Mach-Zehnder modulator comprising of two arms which introduce different phase shifts to the optical signal in order to perform amplitude modulation.	107
4.19	A ring resonator with pn structure for performing resonant wavelength shift enabling optical amplitude modulation.	109
5.1	Different optical receiver architectures. (a) simple resistive front-end, (b) transimpedance front-end with limiting amplifiers, (c) integrating double-sampling receiver.	112
5.2	The proposed RC double-sampling front-end architecture (a). The exponential input voltage and the corresponding double-sampled voltage for a long sequence of successive ones (b).	115
5.3	Modified RC front-end with DOM to resolve input dependent double-sampled voltage (a). The basic operation of DOM technique (b).	117
5.4	Block diagram of the offset modulation technique (a). The first sample is subtracted from the double-sampled voltage, $\Delta V[n]$ , to make it constant regardless of the input sequence. Simulated operation of the DOM for a long sequence of ones showing $\Delta V[n]$ before and after DOM (b).	118
5.5	Top level architecture of the RC double-sampling front-end.	120
5.6	Detailed schematic of the RC double-sampling front-end.	121

5.7	Schematic showing the noise sources in the front-end (a). This plot shows how the clock jitter is translated into the double-sampled voltage noise (b). There is an optimum range, 15-25fF, for the sampling capacitor to achieve maximum SNR (c).	123
5.8	Basic operation of the DOM gain, $\beta$ , adaptation algorithm. The error signal is generated for a certain pattern depending on the difference between $\Delta V[n]$ and $\Delta V[n - 1]$ .	127
5.9	The input waveform and $\beta$ error detection (a), modified sense-amplifier as the difference comparator (b), samplers and comparators for error detection (c), bang-bang $\beta$ adaptation loop (d).	128
5.10	Simulated performance of the front-end before (a) and after (b) DOM adaptation. Gaussian noise with $\sigma=10\text{mV}$ is applied at the sampler.	129
5.11	Optimum sampling capacitor size, $C_S$ versus the photodiode capacitance, $C_{PD}$ .	131
5.12	Receiver current sensitivity versus photodiode capacitance, with and without scaling sampling capacitor (a). Receiver data rate versus photodiode capacitance for $100\mu\text{A}$ sensitivity (b).	135
5.13	2x-oversampled phase detection for the proposed receiver.	136
5.14	The input waveform and baud-rate phase detection, for in-phase (a), and out of phase clock (b).	137
5.15	Electrical measurement setup.	138
5.16	Photodiode current emulator.	139
5.17	Receiver sensitivity characteristics for different data rates.	140
5.18	Current and voltage sensitivity versus data rate.	141
5.19	Power consumption and efficiency at different data rates.	141
5.20	The receiver power breakdown.	142
5.21	Optical test set-up.	142
5.22	Micrograph of the receiver with bonded photodiode (a). Coupling laser through fiber to the photodiode (b).	144
5.23	Optical input eye-diagram to the photodiode at 14Gb/s (a) and 24Gb/s (b).	144

5.24	Optical sensitivity at different data rates. . . . .	145
5.25	Comparison between voltage sensitivity for electrical and optical measurement. . . . .	146
6.1	On-chip interconnect length trend. . . . .	149
6.2	On-chip metal stack in different technology nodes, 130nm (a), 65nm (b), and 32nm (c). . . . .	150
6.3	The schematic profile of diffusion barrier layer (a). SEM cross section of the diffusion barrier layer [171] (b). . . . .	151
6.4	A simple capacitance model for on-chip wires. . . . .	152
6.5	Transition time ( $t_r$ ) versus the length of the interconnect line ( $l$ ). The crosshatched area denotes the region where inductance is important. . . . .	154
6.6	Simple lumped model for an RC dominated wire. . . . .	155
6.7	Simple model for evaluating capacitive coupling. . . . .	157
6.8	Capacitive coupling model in a wire with considerable resistive loss. . . . .	157
6.9	Techniques to combat capacitive coupling in on-chip wires. . . . .	158
6.10	Differential signaling along with wire twisting to remove crosstalk. . . . .	159
6.11	Ground shield insertion to avoid croostalk. . . . .	160
6.12	Inserting repeaters to improve on-chip wire delay. . . . .	161
6.13	A segment of a repeated wire represented by a $\pi$ -model. . . . .	162
6.14	Delay of the repeated wire normalized to the optimal delay (a). The constant delay contours as a function of the repeater normalized width and wire segments relative to optimal values. . . . .	165
6.15	Constant power contours along with delay contours illustrating the trade-off between power and delay. . . . .	166
6.16	Constant power-delay product contours. . . . .	166
6.17	Eye diagram when 10-to-90% rise time equal to the bit time. . . . .	167
6.18	Constant rise time contours for a repeated wire. . . . .	168
6.19	Constant rise time along with power contours. . . . .	169
6.20	Two kinds of wire on a chip: local and global . . . . .	170

6.21	Unrepeated global and semi-global wires delay normalized to FO-4 inverter delay for different technologies. . . . .	170
6.22	Repeated global and semi-global wires delay normalized to FO-4 inverter delay for different technologies. . . . .	171
7.1	Repeated and un-repeated wire delay variation trend with CMOS technology scaling. . . . .	174
7.2	The repeater distance and number of repeaters for an optimally repeated wire in different technology nodes. . . . .	175
7.3	Projection of the repeated wire power consumption in different technology nodes.	176
7.4	Dynamic power breakdown for a single core processor [178]. . . . .	177
7.5	Charge-recycling stacked transmitter employed to reduce effective supply voltage.	178
7.6	Current-mode on-chip wire drivers. . . . .	180
7.7	Sense amplifier based current-mode on-chip wire driver. . . . .	180
7.8	Receiver top-level architecture, double-sampling technique and DOM. . . . .	184
7.9	Z-domain representation of the double-sampler and the dynamic offset modulation (a). Operation of the dynamic offset modulation (b). . . . .	185
7.10	Capacitively-driven transmitter (a), double-sampling technique to resolve the received data (b). . . . .	186
7.11	Frequency characteristics of a minimum-pitch 7mm wire along with the power spectral density of the double-sampled pulse. . . . .	187
7.12	Power spectral density of the transmitted pulse and the double-sampled pulse.	188
7.13	Transistor level schematic of the receiver front-end and the StrongArm sense amplifier with capacitive offset cancellation. . . . .	189
7.14	Shielded single-ended on-chip wire (a). Simulated and measured characteristics of the on-chip wires (b). . . . .	191
7.15	Die micrograph. . . . .	192
7.16	Total power consumption of the receiver and the transmitter for the 4mm, 5mm, and 7mm links. . . . .	192

7.17	Power breakdown for the 5mm, and 7mm links at different data rates. . . . .	193
7.18	Crosstalk measurement setup. . . . .	194
7.19	Power consumption of the 4mm link in the presence of an aggressor at different data rates. . . . .	194
7.20	Receiver output eye diagram at 10Gb/s (a) and 20Gb/s (b) input data rate.. .	195

# List of Tables

3.1	DFE PERFORMANCE SUMMARY . . . . .	74
3.2	CROSSTALK CANCELLATION PERFORMANCE SUMMARY . . . . .	75
3.3	TRANSMITTER PERFORMANCE SUMMARY . . . . .	85
5.1	OPTICAL RECEIVER PERFORMANCE SUMMARY. . . . .	146
6.1	PERFORMANCE OF ON-CHIP WIRES IN 28nm TECHNOLOGY. . . . .	156
7.1	OPTICAL RECEIVER PERFORMANCE SUMMARY. . . . .	195

# Chapter 1

## Introduction

Most advanced electronic systems today require complex architectures that consist of interconnected integrated circuits (IC). Examples of such systems are computer servers and high-performance multiprocessor systems, Figure 1.1. CMOS technology scaling has enabled a huge growth in processing capability. Figure 1.2 shows the history of the microprocessors. In the early age of the CMOS technology, the performance of the microprocessors was improved by both scaling the number of transistors per area [1] as well as the operating clock frequency. This approach provided a tremendous improvement in processing power until 2004, when designers ran into a new problem, power consumption. It turned out that by scaling clock frequency, marginal improvement in processing performance was achieved while a significant power penalty had to be paid [2]. As a result of the increased power, designers employed a parallel computing approach through multi-core processors. Nowadays we have microprocessors with over tens of cores on a single die and in the near future processors are expected to have hundreds of cores to enable exa-scale computing [3]. This increase in computing power necessitates a corresponding increase in inter-chip as well as on-chip communication bandwidth. As shown in Figure 1.3, the I/O bandwidth requirement for microprocessor products scales by a factor of 2-3 every two years [5] and this trend is anticipated to continue according to International Technology Roadmap for Semiconductors (ITRS) [4]. High aggregate bandwidth can be achieved by employing large numbers of inputs and outputs (IOs) per chip as well as high data rates per IO. As evident from Figure 1.3 the number of pins does not scale as fast due to the physical connection limitations. This



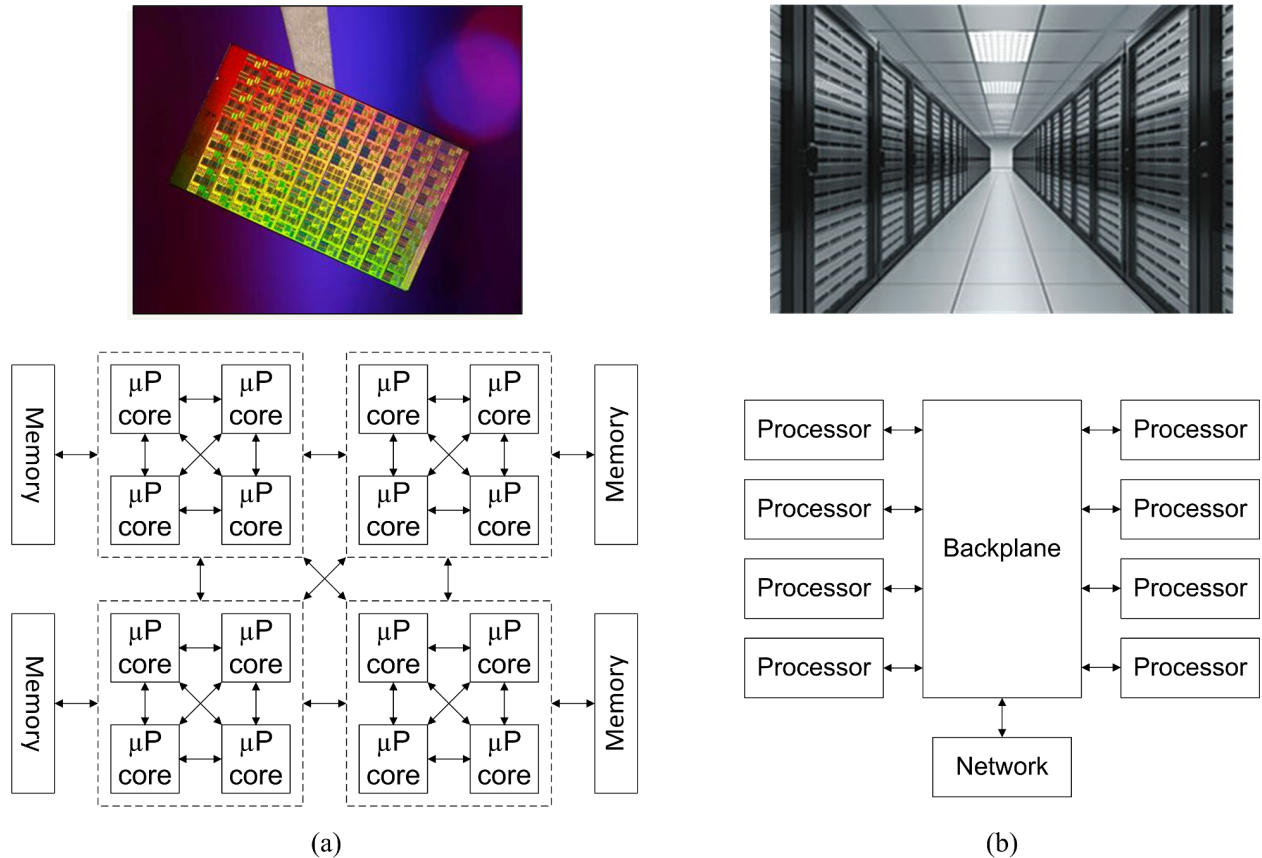


Figure 1.1: Examples of complex electronic systems. A high-performance multi-core processor (a). A computer server (b).

leaves the increase in per pin bandwidth as the only solution to the future I/O bandwidth problem.

Electrical interconnects are conventionally the main platform for data communication. However, due to their limited bandwidth, the scaling of data rate has proven to be very difficult. The dielectric and resistive losses of printed circuit board (PCB) traces increase as the operation frequency increases. Such frequency dependent attenuation causes inter-symbol interference (ISI) and ultimately signal-to-noise-ratio (SNR) degradation. In addition, reflections from discontinuities in the signal path due to transitions from chip-to-package and package-to-board generate more ISI and further reduce the SNR. These problems are exacerbated as the data rate increases. As the number of interconnects increases, the spacing between channels decreases to allow for accommodating more channels and hence achieving

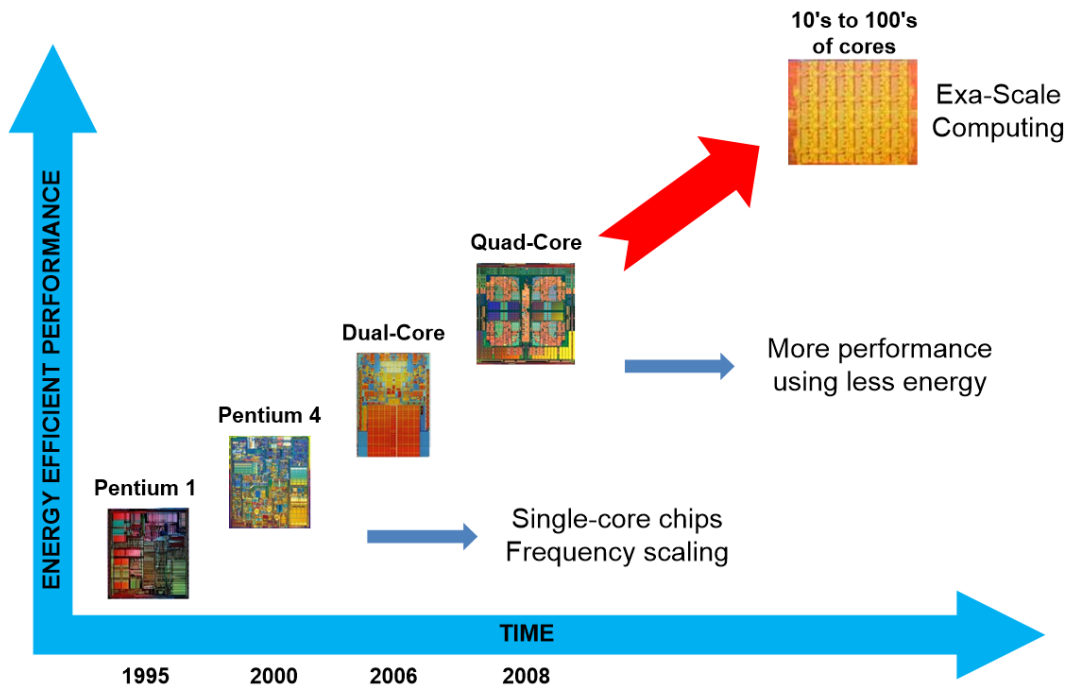


Figure 1.2: History of microprocessor performance scaling.

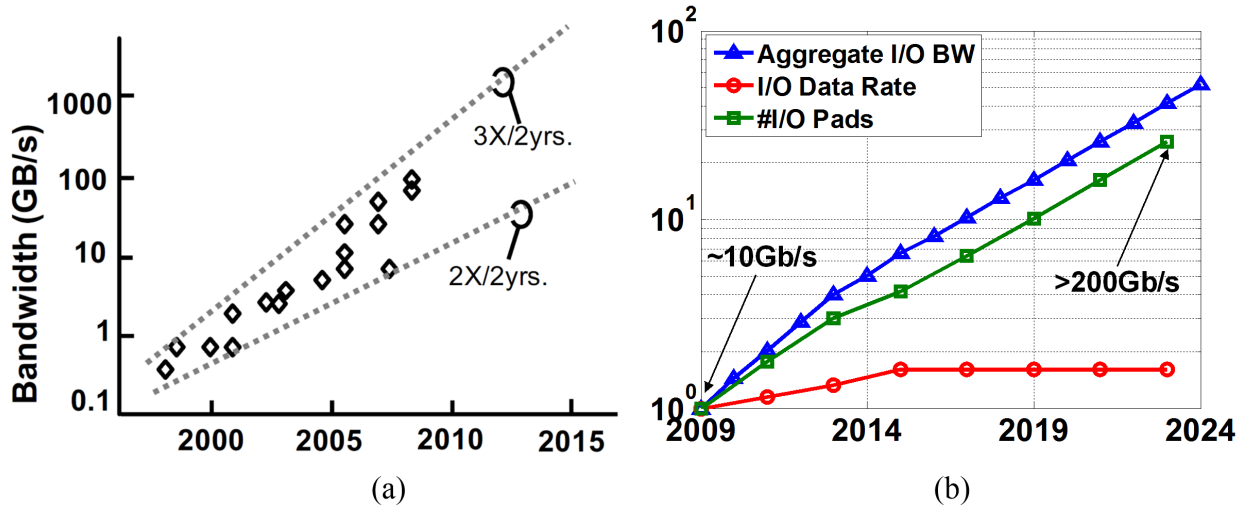


Figure 1.3: IO bandwidth requirement of microprocessors in recent years (a). Constant growth of the required IO bandwidth according to ITRS (b).

higher aggregate data rate. One of the consequences of this approach is the excessive capacitive and inductive coupling between adjacent channels, which manifests itself as crosstalk noise. Crosstalk can also severely degrade signal integrity. The effect of ISI and crosstalk will be explained in more detail in the next chapter. A common approach in the design of high-speed serial links over bandwidth-limited channels is to employ equalization techniques to cancel destructive effects of ISI. Typical equalization techniques include decision feedback equalization (DFE) [64–67], feed-forward equalization (FFE) [43–47] and continuous time linear equalization [68–71] at the receiver and FFE at the transmitter [33–37]. These techniques can be used in parallel links with many IOs to increase the aggregate data rate. A number of techniques have been also proposed to remove the effects of crosstalk. The design in [73] employs an FFE equalizer and [76] uses crosstalk-induced jitter equalization at the receiver. Other approaches to compensate for crosstalk noise include the use of staggered I/Os [77] or a finite-impulse response (FIR) filter at the transmitter [78]. All these schemes result in significant power consumption and are not suitable for parallel data links. This dissertation provides a compact low-power electrical link architecture with crosstalk cancellation capability to enable scaling of chip-to-chip communication bandwidth. The main emphasis of the proposed design is low-power consumption, small area and scalability to future technologies.

As data rates scale to meet increasing bandwidth requirements, the shortcomings of copper channels are becoming more severe. As mentioned earlier, in order to continue scaling data rates, equalization techniques can be employed to compensate for the ISI. However, the power and area overhead associated with equalization make it difficult to achieve target bandwidth with a realistic power budget. As a result, rather than being technology limited, current high-speed I/O link designs are becoming channel and power limited. A promising solution to the I/O bandwidth problem is the use of optical inter-chip communication links. The negligible frequency dependent loss of optical channels provides the potential for optical link designs to fully utilize increased data rates provided through CMOS technology scaling without excessive equalization complexity. Optics also allow very high information density through wavelength division multiplexing (WDM). By providing a high-capacity channel,

optical signaling can potentially close the gap between the interconnect speed and on-chip data processing speed. This dissertation investigates the challenges of designing electronics for short-haul optical links and proposes a number of solutions to enable optical IOs. We focus on techniques to design simple, compact and low-power receivers suitable for dense parallel optical interconnects.

Hybrid integration of optical devices with electronics has been demonstrated to achieve high performance [89, 129–134], and recent advances in silicon photonics have led to fully integrated optical signaling [135, 136]. These approaches pave the way to massively parallel optical communications. In order for optical interconnects to become viable alternatives to established electrical links, they must be low cost and have competitive energy and area efficiency metrics. Dense arrays of optical detectors require very low-power, sensitive, and compact optical receiver circuits. Existing designs for the input receiver, such as TIA, require large power consumption to achieve high bandwidth and low noise, and can occupy large area due to bandwidth enhancement inductors. Moreover, these analog circuits require extensive efforts to migrate and scale to future technologies. Therefore, in this thesis we develop techniques to implement low-power, compact receiver circuits for highly parallel optical communication.

As VLSI technologies continue to scale, on-chip wires will present increasing latency and energy problems. While circuit performance benefits from technology scaling, the shrinking cross-sectional area of the on-chip wires increases electrical resistance and hence latency. Repeaters mitigate the latency problem but do little to improve the energy cost. Moreover, as technology scales, the number of repeaters grows significantly, which increases power consumption and adds complexity. A CMOS wire driver running at an effective frequency  $f$  must switch a total wire capacitance  $C_w$  through the voltage  $V_{dd}$ , leading to a power cost proportional to  $C_w f V_{dd}^2$ . Under technology scaling, wire capacitance remains largely constant (for global wires spanning constant-sized die),  $V_{dd}$  scales down only slowly, while  $f$  scales up, leading to nearly constant power per wire. With chips containing more and more devices, this constant power per wire gets multiplied by an ever-increasing number of wires. For instance, each chip in a data-routing grid for a high-performance processor

may carry in excess of 250 meters of wiring interconnect, which would burn nearly 50 W of wire and repeater power at 4 GHz [183]. Designers have proposed different techniques to mitigate the power and latency problem of the on-chip wires. Low-swing differential signaling [179, 182–185, 191], current-mode signaling [186, 187, 192], equalization [185, 191] and transmission lines [179, 186] have been employed to resolve the energy and latency problem of the repeated links. However, these techniques are becoming less adequate in meeting bandwidth density and power requirements. In this thesis we have attempted to provide a solution that maximizes bandwidth density for on-chip communication in a power and area efficient manner.

In summary, this dissertation provides solutions to enable increasing data rate both for chip-to-chip and on-chip communications with the emphasis on low power consumption to meet the ever-increasing demand for the bandwidth required by future microprocessors.

## 1.1 Organization

This thesis is divided into three closely related topics. Chapter 2 provides a background in high speed data transmission systems. Shortcomings of electrical communication channels are introduced and their effects on the signal integrity and maximum achievable data rate are investigated. It is shown that for short channels, electrical signaling can still be a viable solution especially due to low cost and efficient integration. In chapter 3 we introduce a solution to achieve high bandwidth over short electrical communication channels using high performance low-power equalization techniques. A low-power high data rate electrical transmitter is presented that employs passive pre-emphasis to enhance the overall channel bandwidth and a DFE receiver to compensate for the residual post-cursor ISI. The proposed 2-tap DFE receiver employs a novel switched capacitor technique to perform equalization in a power efficient manner. It has been also shown that this technique can be readily generalized to many equalization taps with minimal power overhead compared to conventional techniques. In addition, a novel crosstalk cancellation is integrated in the receiver to allow for highly parallel communication over high loss and coupled channels to achieve high aggregate data

rate. It is shown that the proposed crosstalk cancellation technique operates efficiently with minimal power overhead.

As the communication distance increases to tens of inches, the efficiency of equalization technique drastically degrades due to the increase in the number of equalization taps required to compensate for the excessive channel loss. This makes electrical links inadequate in meeting the bandwidth requirement for long communication distances. In addition optical links can have significant advantages over electrical links if a large number of parallel optical channels can interface with each IC. For parallel optical interconnects, the design of a low-power receiver front-end is particularly challenging. In Chapter 4 we provide an introduction to chip-to-chip optical communication links with an emphasis on optical receivers. This chapter continues with exploring the prior art and investigating their challenges in meeting the requirements for highly parallel high data rate optical links. Chapter 5 of this thesis focuses on the receiver design for chip-to-chip optical interconnects. It is shown that a promising solution to these challenges is an RC front-end which employs double-sampling as well as dynamic offset modulation techniques. Unlike most prior designs, this receiver avoids high-gain analog blocks that operate at the the input data-rate. The eventual goal in optical interconnect design is to have thousands of transceivers in a single chip. Scaling of parallel optical interconnects to hundreds and thousands of links on a single chip requires receiver and transmitter circuitries that are very compact and have very low power consumption at high data rates. The proposed design was implemented in a test-chip fabricated in 65nm CMOS technology. Measurement results from this test-chip show that the proposed design achieves low-power consumption at high data rates while achieving good sensitivity for hybrid integrated solutions which offer moderate parasitic capacitance in the range of 100-200fF.

In Chapter 6 of this dissertation we take a close look at on-chip wires scaling and investigates the challenges of on-chip signaling in highly-scaled technologies. Then in Chapter 7, we will introduce a novel technique inspired by the optical receiver introduced in Chapter 5 in conjunction with low-swing signaling techniques to mitigate these challenges in a power and area efficient manner.

Finally, in Chapter 8 we summarize the conclusions of this work.

# Chapter 2

## Electrical Interconnects Background

High-speed electrical links are commonly used in short distance chip-to-chip communication applications such as Internet routers [7–9], multi-processor systems [10–13], and processor-memory interfaces [14–16]. The rapid scaling of CMOS technology continues to increase the processing power of microprocessors and the storage volume of memories. This increases the need for high bandwidth interconnection between chips, which can be achieved by employing large numbers of inputs and outputs (IOs) per chip as well as high data rates per IO. In order to achieve high data rates, these systems employ specialized IO circuitry that performs signaling over carefully designed controlled-impedance channels. As will be described later in this section, the electrical channel’s frequency-dependent loss, crosstalk and impedance discontinuities become major challenges in data rate scaling. This section begins by describing the three major electrical link circuit components, the transmitter, receiver, and timing generation. Next, it discusses the electrical channel properties that impact the transmitted signal. Basic receiver and transmitter design and challenges will be described later in this chapter and it concludes by providing an overview of common equalization schemes and advanced modulation techniques that designers implement in order to extend data rates over the band-limited electrical channels.

## 2.1 Metrics of Electrical Links

A link's performance can be evaluated based on several factors. For exploring the maximum data rate of a given technology, two metrics in particular are used to evaluate various designs: the bit-rate (or its inverse, the bit-time) and the bit-error rate (BER). To get a better sense of how bit-rate changes with scaling we normalize it by the speed of the CMOS technology. The normalization factor is the delay of a loaded inverter, as described in the next section.

Since a link's receiver needs to convert an analog signal back into digital data, there is always a probability for errors to occur. Bit-error rate (BER) indicates the reliability of the link in data communication links. A link's maximum data rate is usually specified at a specific BER (e.g.  $10^{-12}$ ) to guarantee the robustness of the overall system. The following section describes how errors occur due to voltage or timing noise. To illustrate the effect of noise on the system performance, the BER is shown as a function of the signal-to-noise ratio (SNR) in a simplified analysis.

### 2.1.1 Bit-Rate

The minimum bit-time varies with the CMOS process technology. It is useful to employ a metric to represent the bit-time that is independent of technology to facilitate performance extrapolation to future technologies. An appropriate metric is the delay of a buffer driving a normalized load. A "FO-4" delay is the delay of one stage in a chain of inverters, in which each inverter is driving another inverter with 4X larger size, as shown in Figure 2.1. In other words, each of the inverters in the chain drives a capacitive load (fan-out) that is 4X larger than its input capacitance. The delay of various circuits can be normalized to FO-4 delay in a certain technology.

This metric is applicable based on the observation that the delays of topologically different CMOS digital circuits scale by approximately the same factor [17]. Figure 2.2 shows the actual FO-4 delay for various technology nodes. It is interesting to note that a reasonable estimate for FO-4 delay is roughly  $400\text{ps}/\mu\text{m}$  of effective channel length and, conversely, the  $f_T^{-1}$  in a certain process can be approximated to be  $\frac{1}{3}$  FO-4 [18]. There is a meaningful






Figure 2.1: Fan-out of four inverter chain


relation between this metric and the maximum achievable data rate in a certain technology node. As a result, by migrating from one technology to another, in theory the maximum data rate improves due to the smaller FO-4. It should be also noted that the increase in data rate makes data transmission difficult over electrical channels. This will be discussed in the remainder of this chapter in more detail.

## 2.1.2 Bit-Error Rate


One of the most important metrics for performance of a link is the bit-error rate, BER, which indicates the reliability of the link. This reliability ties closely with the data-rate as excessive errors may force a link to operate at a lower data rate. The errors are due to noise on the signal that is transmitted and the noise in the receiving circuits as well as the noise introduced by the channel. The noise can be divided into timing noise and amplitude noise.

The effect of noise is often illustrated using a data eye. Figure 2.3 shows how a data eye is created by folding a signal waveform into a single bit-time. For a noise-less signal, the horizontal and vertical eye opening are at their maximum. Noise on the bit stream results in a reduced eye opening, making the signal more difficult to be resolved at the receiver.

Figure 2.4 illustrates ideal data eyes to demonstrate how the sources of errors impact signal reception while assuming a single decision threshold and a single sampling time. As shown in Figure 2.4(a), sufficiently large noise causes the signal to not cross the decision level or to accidentally cross it, which results in a wrong decision. Similarly, a voltage offset in



(a)



(b)

Figure 2.2: FO-4 delay metric for different technology nodes (a), and  $f_T^{-1}$  versus FO4 metric, which shows a linear relation (b).




Figure 2.3: Illustration of a data eye. (a) shows a noisy data stream, and (b) shows the folding of the data stream into a data eye.




Figure 2.4: Effect of amplitude noise in (a) and timing noise in (b) on the eye diagram quality.

the decision level could make data resolution more sensitive to error by reducing the voltage noise margin for a one versus a zero, or vice versa. Figure 2.4(b) shows the effect of both timing offset and timing jitter. Static timing error is the fixed offset in the sample position from the ideal position,  $T_{os}$ . Dynamic timing jitter is due to the random noise in the phase position. The resulting timing margin,  $T_{margin}$ , can be calculated as

$$T_{margin} = T_b - T_{os} - T_{jd} - T_{jc}, \quad (2.1)$$

where  $T_b$  is the bit time, and  $T_{jd}$  and  $T_{jc}$  are the jitter on the data transitions and the sampling clock, respectively. Since the sampling position is defined with respect to the data transition and the clock and data jitter are two independent random processes, jitter on both the clock and the data additively reduces the timing margin. With ideal square pulses, as long as the sum of the magnitudes of the static and dynamic timing errors is less than a bit-time, the sampled value will always be the correct bit. However, because of finite signal slew rates, timing errors that are less than a bit-time can reduce the amplitude of the signal at the sample point hence affecting the BER.

Amplitude and timing noises impact the bit-error rate and hence the performance of the




Figure 2.5: Translation of the eye-diagram to bathtub curve as a tool to measure signal integrity.

system. We will discuss the effect of these noise sources on the performance of the system in the next two sections.

In addition to the eye diagram, the bathtub curve is another diagnostic tool for performing signal integrity analysis. Bathtub curves are usually created by measuring the BER while sweeping the sampling clock over the bit time. Figure 2.5 shows a typical bathtub curve. Bathtub curves are useful tools for characterizing the performance of the receiver and show how tolerant the system is to the sampling clock jitter noise, as well as the amount of horizontal and vertical eye opening. We will be using bathtub curves to characterize electrical receivers in the next chapter.

### 2.1.2.1 Amplitude Noise

In a communication link, the received signal is the sum of the transmitted values and noise which appears as an added signal with random value. At the sample point, there is a finite probability for the noise amplitude to be greater than the signal amplitude, causing a wrong




Figure 2.6: Diagram of bit error rate versus signal to noise ratio.

decision. This probability determines the bit-error rate. The BER indicates how many errors are likely to occur for a certain number of resolved bits. This rate often depends on the amount of signal power and the amount of noise power.

For instance, the probability of error due to a noise with Gaussian distribution can be expressed as a function of the signal-to-noise ratio (SNR) in case of an equiprobable one or zero [19]:

$$P_{error} = \int_A^{\infty} \frac{1}{2\pi\sigma_n^2} \exp\left(-\frac{x^2}{2\sigma_n^2}\right) dx = 1 - Q\left(\frac{A}{\sigma_n}\right) = 1 - Q(SNR), \quad (2.2)$$

where  $A$  is the signal amplitude,  $\sigma_n$  is the standard deviation of the noise, and  $Q(x)$  is the Q-function and represents the tail probability of the standard normal distribution. Figure 2.6 shows how the BER changes with SNR considering the above model. Increasing the signal amplitude will increase SNR, and hence improve the BER.

Other than the white noise, there are other sources of noise that can degrade the overall SNR, such as supply and substrate noise. These noise sources, unlike the white noise, are bounded in amplitude and usually scale with the signal amplitude as well as signal activity. As a result, the absolute BER can not be solely related to the total noise power as shown in Equation 2.2. Nevertheless, the SNR versus BER analysis serves as a useful tool. This method can be used to illustrate how a dc offset in the decision level can degrade performance. The offset can be considered as a reduction of the signal amplitude for one of the two signal

levels. A decision level shifted higher by  $\alpha A$  would reduce the amplitude of a one and increase the amplitude of a zero. In case of equiprobable one and zero, the probability of error is the average of the two probabilities. Since the probability increases exponentially with decreasing SNR, the error rate is dominated by the signal value with the lower SNR which, in this case, is due to the one. This reduction in performance can be expressed as an SNR penalty of  $20\log_{10}(1 - \alpha)$ .

### 2.1.2.2 Timing Noise

Timing noise can similarly affect the link's performance. Timing noise can have a noise distribution similar to the amplitude noise. If the magnitude of the timing error exceeds half the bit-time, the receiver would sample the previous or next bit instead of the current bit, resulting an error. The probability of the noise exceeding  $\frac{T_b}{2}$  determines the minimum BER independent of the signal amplitude. Typically the timing noise does not scale with decreasing the bit-time, which results in increasing the minimum BER.

In addition to affecting the minimum BER, phase errors can affect SNR. As seen in Figure 2.4(b), sampling away from the peak of the waveform results in a sampled value that is less than the peak signal amplitude. Also, due to the finite rise and fall time of the input data, timing noise effectively results into amplitude noise and hence reduces the SNR.

## 2.2 Basics of Electrical Links

Electrical links can provide high communication bandwidths between chips, and consist of three major components as shown in Figure 2.7. The transmitter converts the digital data into an electrical signal that travels through the channel. The electrical channel is the complete electrical path from one die to the other. This channel can consist of traces on a printed circuit board (PCB), coaxial cables, shielded or unshielded twisted pairs of wires, traces within chip packages, and the connectors that join these various parts together. A receiver then converts the incoming electrical signal back into digital data.

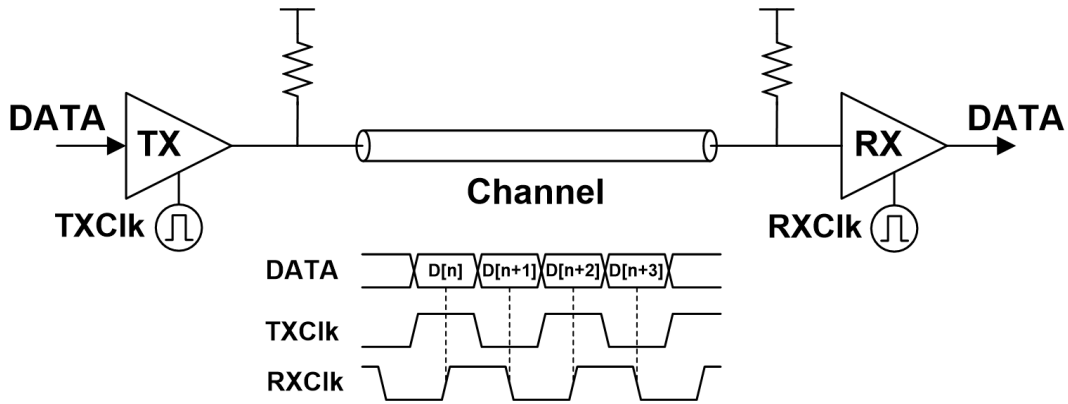


Figure 2.7: Components of an electrical communication link.




Figure 2.8: A typical backplane link and its components [22].

## 2.2.1 Channel

As the data rate keeps increasing to meet the bandwidth requirement for today's systems, the filtering imposed by the electrical channel becomes the most challenging problem. The performance of the channel strongly depends on the application. As an example a typical backplane link and its components are shown in Figure 2.8 [20, 21]. Loss per unit length of PCB-traces increases with the frequency due to the dielectric loss and skin effect. Different trace lengths and backplane material properties, as well as types of connectors, vias and routing layers, cause significant variation in channel transfer characteristics.

The typical transfer characteristics for channels within a single backplane are illustrated in Figure 2.9 [22]. Since the loss in the channel increases with the frequency, the channel acts




Figure 2.9: Transfer characteristics of typical backplane channels with and without via stubs at different lengths [22].

as a low-pass filter. The filtering effects lead to a spread of narrow pulses originally confined to a bit-period as shown in Figure 2.10, which is referred to as pulse dispersion. The tail of the pulse acts as an additive noise for the next bits and is referred to as inter-symbol interference or ISI. Dispersion is enhanced by the filters formed by unintended transmission line impedance discontinuities caused by via stubs and connections. In the time domain these discontinuities cause reflections, which also lead to ISI. Crosstalk is the other problem that occurs in dense interconnects. Both far and near end cross-talk (FEXT and NEXT), are important in such systems.

### 2.2.1.1 Inter Symbol Interference

At frequencies well into the gigahertz range, the electrical channels start behaving like lossy transmission lines. As mentioned earlier, skin-effect and dielectric loss are two contributing effects causing the loss to increase with frequency. Skin-effect is manifested as crowding of the higher-frequency current toward the surface of the conductor. This results in a smaller effective cross-section area for the current flow and hence increase in the effective resistance. In fact, this resistance is proportional to the square root of the frequency and can be expressed






Figure 2.10: Pulse response of a typical bandwidth-limited backplane channel illustrating dispersion.

as

$$R_{AC}(f) = \frac{2.16 \times 10^{-6}}{\pi D} \sqrt{p_r f}, \quad (2.3)$$

where  $D$  is the wire diameter and  $p_r$  is the relative resistivity of the wire material compared to copper [23].

Dielectric loss is related to the energy loss in the dielectric surrounding the transmission line. This loss increases proportionally to signal frequency [23]

$$\sigma_D = \frac{\pi s q r t \epsilon_r}{c} f \tan \delta, \quad (2.4)$$

where  $\tan \delta$  is the loss tangent,  $c$  is the speed of light and  $\epsilon_r$  is the relative permittivity. Due to the linear dependence on frequency, the dielectric loss dominates over the skin-effect at high frequencies. The crossover frequency depends on the material properties and dimensions of the trace, for instance, the crossover occurs at around 500 MHz for FR4 material [24].

The frequency dependent loss due to the dielectric loss and skin effect distorts the transmitted pulses with sharp rise and fall times and causes pulse broadening. This phenomenon is illustrated in Figure 2.11. As a result, every bit is extended to previous and next bits, which is referred to as post- and pre-cursors. Post- and pre-cursors cause ISI, effectively

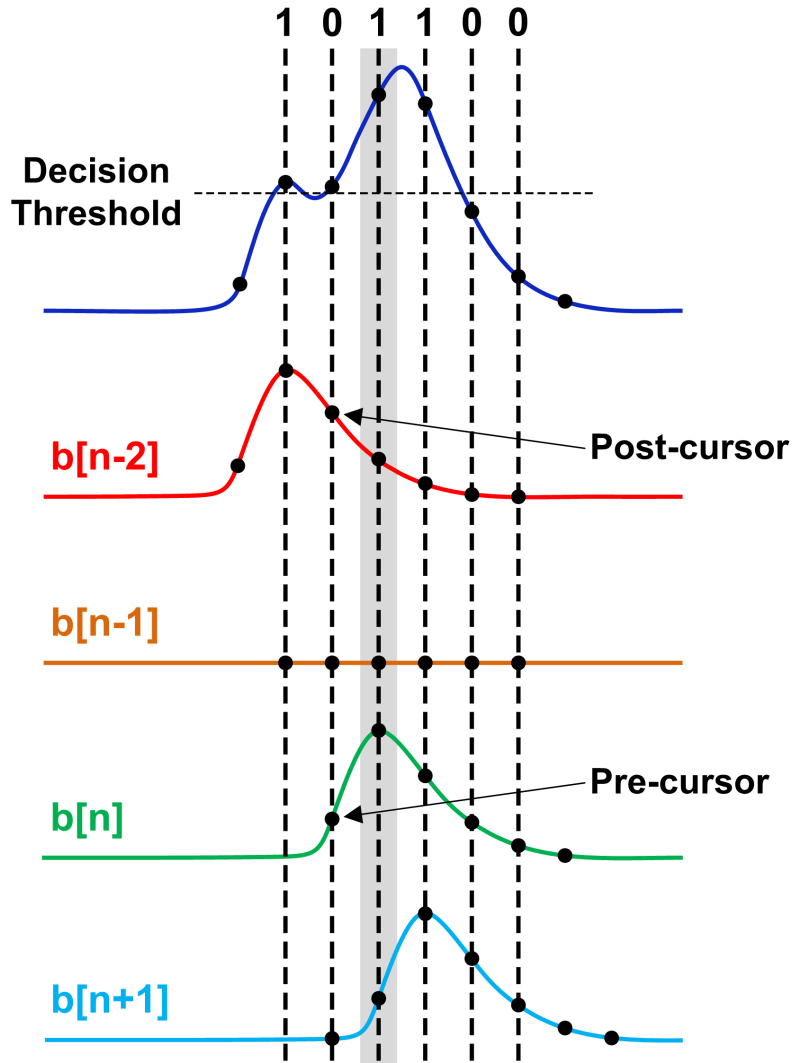


Figure 2.11: Channel dispersion resulting in pre- and post-cursor ISI.

reducing the eye opening at the receiver, hence degrading the overall BER. This problem is exacerbated as the data rate increases.

As discussed earlier, another source of ISI is reflection due to impedance discontinuities. A signal transitioning from one transmission line with characteristic impedance of  $Z_1$  to another line with different impedance,  $Z_2$  experiences a reflection of magnitude

$$R = \frac{Z_1 - Z_2}{Z_1 + Z_2}. \quad (2.5)$$

Impedance discontinuity can be caused by inaccurate termination of the transmission line




Figure 2.12: Far-end and near-end crosstalk in coupled transmission lines.

at the receiver or the transmitter side. It can be also created by via stubs on the board that carry signal from one metal layer to another. The stub acts as a capacitor, which reflects high frequency energy. Another dominant source of reflection is the frequency dependent impedance discontinuity due to parasitic device capacitance at both the transmitter and receiver.

### 2.2.1.2 Crosstalk

As the demand for communication bandwidth increases, link designers start to use more and more channels in parallel to increase the aggregate data rate. Placing channels in close proximity causes electromagnetic coupling between them, which can result in crosstalk interference. Crosstalk can be divided into far-end (FEXT) and near-end (NEXT). As shown in Figure 2.12, FEXT occurs when the aggressor signal travels in the same direction as the victim. The NEXT occurs when the aggressor signal travels in the opposite direction, and can be much more critical as the victim signal is severely attenuated due to the channel loss. Since crosstalk is caused either by capacitive or inductive coupling of different signal lines it has high attenuation at low frequencies. Due to the low-pass filtering of the channel, FEXT is also attenuated at high-frequencies. Therefore FEXT is mostly band-pass, while NEXT is high-pass.

Figure 2.13 shows the lumped equivalent circuit of a segment of coupled transmission

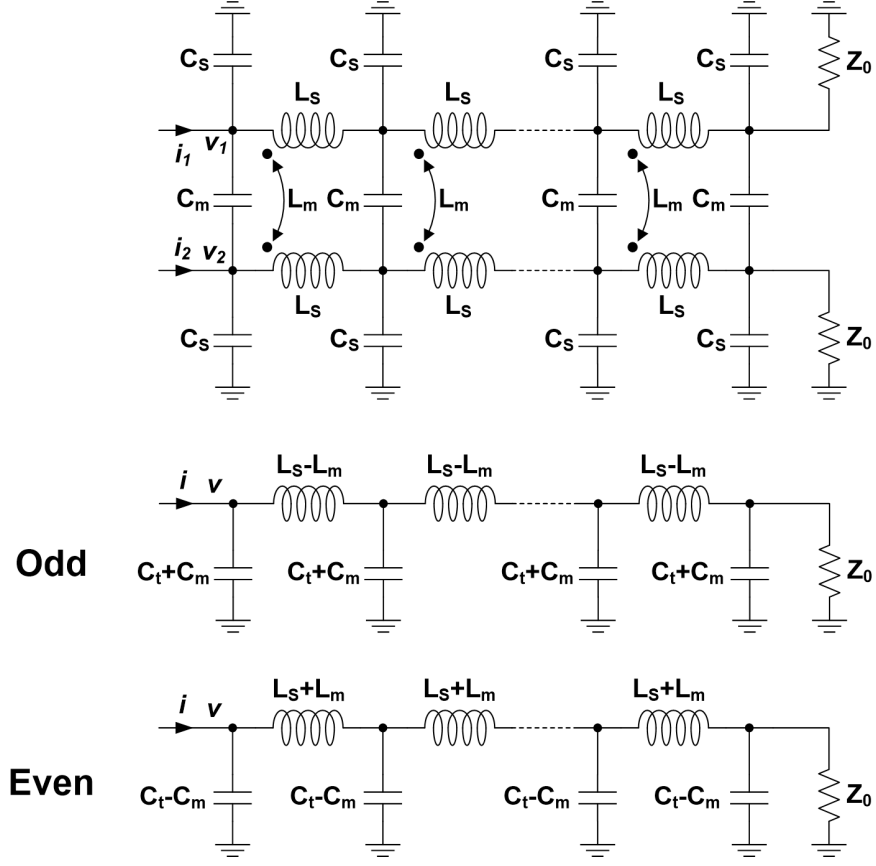


Figure 2.13: Lumped equivalent model for coupled transmission lines.

lines, where  $C_S$  and  $C_m$  represent the self and mutual capacitances of the transmission line per unit length, respectively, and  $L_S$  and  $L_m$  represent the self and mutual inductances per unit length, respectively. From the equivalent circuit, the voltage and current relation of the coupled lines can be described as

$$\frac{d}{dz} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} L_S & L_m \\ L_m & L_S \end{bmatrix} \times \frac{d}{dt} \begin{bmatrix} i_1 \\ i_2 \end{bmatrix}, \quad (2.6)$$

$$\frac{d}{dz} \begin{bmatrix} i_1 \\ i_2 \end{bmatrix} = \begin{bmatrix} C_t & -C_m \\ -C_m & C_S \end{bmatrix} \times \frac{d}{dt} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}, \quad (2.7)$$

where  $C_t = C_S + C_m$ . Equations 2.6 and 2.7 show that the current change at the aggressor line causes the inductively-coupled voltage drop at the victim line and the voltage change at

the aggressor line causes the capacitively-coupled current change in the victim line. Using the weak coupling assumption [25], it can be shown that the near-end and far-end crosstalk can be calculated as

$$V_{FEXT} = \frac{1}{4} \left( \frac{C_m}{C_t} + \frac{L_m}{L_S} \right) (V(t) - V(t - 2t_f)), \quad (2.8)$$

$$V_{FEXT} = \frac{1}{2} \left( \frac{C_m}{C_t} - \frac{L_m}{L_S} \right) t_f \frac{dV(t - t_f)}{dt}, \quad (2.9)$$

where  $V(t)$  is the driven pulse, and  $t_f$  is the time of flight along the coupled lines. It should be noted that these equations are valid in the case of matched terminated lossless transmission lines. For lossy transmission lines, the above equation can still be used when corrected by the attenuation. In addition, these equations only signify the effect of the crosstalk on the received amplitude of the voltage at both ends of the victim line. Another way that crosstalk manifest itself is through inducing timing jitter [25]. Solutions to Equations 2.8 and 2.9 can be decomposed into two modes, the even and odd modes. The even and odd modes effectively see two different coupling mechanisms. For the even mode, the coupling capacitance has no effect in the signal propagation and as the two lines are carrying the same current, the effective inductance is  $L_S + L_m$ , whereas for the odd mode, the effective capacitance and inductance are  $C_S + C_m$  and  $L_S - L_m$ , respectively. As a result, the propagation constant for these two modes can be expressed as [26]

$$\beta_o = \omega \sqrt{(L_S - L_m)(C_t + C_m)}, \quad (2.10)$$

$$\beta_e = \omega \sqrt{(L_S + L_m)(C_t - C_m)}, \quad (2.11)$$

where  $\omega$  is the angular frequency, and  $\beta_o$  and  $\beta_e$  are odd and even mode propagation constants, respectively. Using the relation between the propagation constant and the phase velocity, we can calculate the even and odd mode times of flight as

$$t_{fo} = l\sqrt{(L_S - L_m)(C_t + Cm)}, \quad (2.12)$$

$$t_{fe} = l\sqrt{(L_S + L_m)(C_t - Cm)}. \quad (2.13)$$

The speed of the two modes are equal if  $L_m/L_S=C_m/C_t$ . This condition is held in a homogeneous transmission line, such as a stripline. For microstrip lines, homogeneity is not guaranteed as the electric and magnetic fields are not symmetric above and below the line due to the surrounding air. The main effect of this difference in the time of flight for even and odd modes is the generation of crosstalk-induced jitter. This phenomena is shown in Figure 2.14. As seen for different cases of the aggressor and victim signal transition, the time of flight for both lines changes. For instance, when both signals transition in the same direction, it is different from when they transition in opposite directions.

It should be noted that in the above analysis we assumed that both victim and aggressor signals transition at the same time. This resulted in data dependent jitter due to the difference between the time of flight for the even and odd modes. However, in a real system it is likely that the transition on the victim and aggressor signal does not happen at the same time. For instance, Figure 2.15 illustrates the case in which the aggressor signal transitions in the middle of the bit time of the victim signal [27]. This will create amplitude noise in the victim signal, which according to Equation 2.9 is proportional to the derivative of the aggressor signal for a far-end channel. This type of crosstalk directly affect the SNR at the receiver, while in the other type, it affects the timing margin.

## 2.2.2 Transmitter

A data transmitter converts digital data into electrical signals that propagate through the channel to a receiver at the opposite end. For high-speed data communication links, this must be done with accurate signal levels and timing. The receiver chooses a threshold voltage to resolve the data sent from the transmitter. A common voltage reference is then required




Figure 2.14: Effect of coupling on even and odd modes, which results in data dependent jitter.

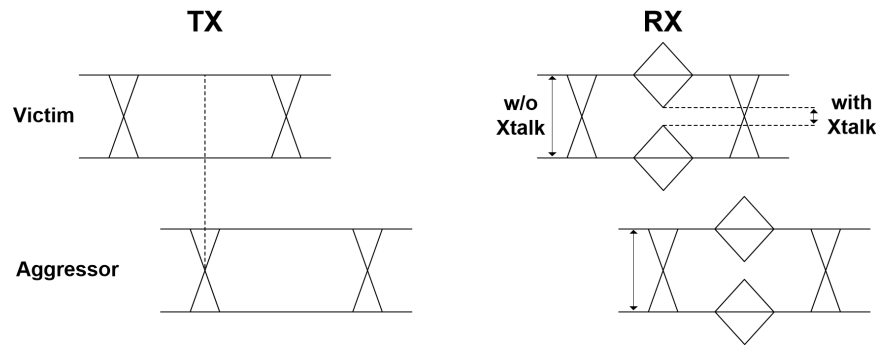


Figure 2.15: Effect of coupling on the victim signal amplitude when the transmitted aggressor and victim signals are not synchronized.




Figure 2.16: Voltage-mode and current-mode transmitter design.

to correctly resolve the data value. Ground level is usually used as the common voltage reference for data resolution. The signal transmission over the line can be done with either a voltage-mode driver or a current-mode driver, shown in Figure 2.16.

In voltage-mode drivers, switches are employed to alternate the line voltage between zero and one, as shown in Figure 2.17(a). Because the switches are implemented with transistors, the driver appears as a switched resistance. To switch the voltage fully, a small resistance is needed, which typically requires a large switching device. In contrast, current-mode drivers are switched current sources, where the output signals are generated via a current source that turns on and off depending on the transmitted data, Figure 2.17(b). The voltage swing at the output depends on the termination and the size of the current source. In order to control the voltage swing and proper termination, to avoid reflections, the current source should be kept in the saturation region and the termination resistor should be carefully designed and controlled [28–30].

Current-mode drivers are slightly better in terms of insensitivity to supply noise due to the high output impedance. The output current does not vary with ground noise as long as the bias signal is tightly coupled to the ground signal. The disadvantage of current-mode drivers is that, in order to keep the current sources in saturation, the transmitted voltage range must be well above ground which increases power dissipation and reduces the signal swing.



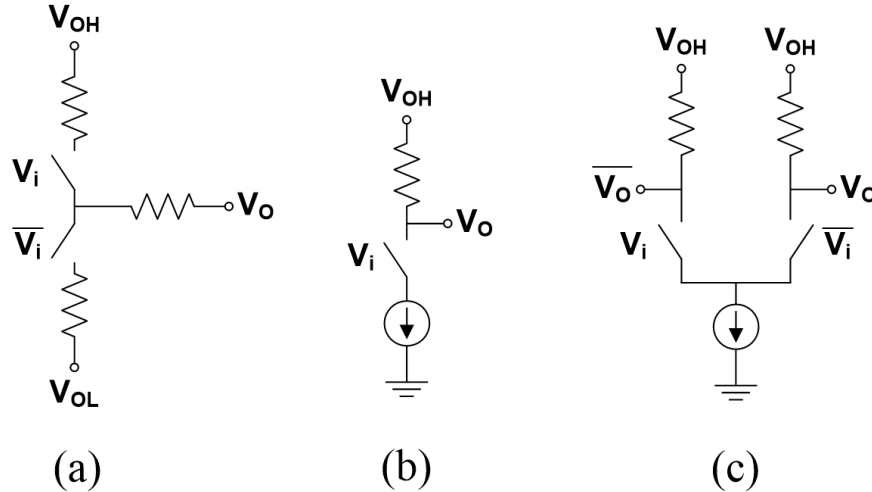


Figure 2.17: Different transmitter examples, (a) voltage-mode, (b) single-ended current mode, (c) differential current-mode.

For better supply-noise rejection, the outputs can be driven differentially, as shown in Figure 2.17(c), as it appears as a common-mode noise. Since the current remains roughly constant, the transmitter also induces less switching noise on the supply which could benefit other sensitive circuits on the same chip. However, the static current in this configuration makes it quite power hungry.

To reduce reflections at the transmitter side, it should be designed such that the output resistance properly serves as the termination resistor. An on-chip resistor can be incorporated with current-mode drivers to act as the source termination resistor [31]. Given the fact that the current source introduces a fairly large output resistance when biased in saturation, the overall output resistance would be almost equal to this termination resistor. In the case of voltage-mode drivers, the design is slightly more complex because the switch resistance should match the line impedance. This may be done either through proper sizing of the driver [32] or by oversizing the driver and compensating with an external series resistor, as shown in Figure 2.17(a) [29].

As mentioned in the previous section, reactive parasitics on the transmission lines, modeled as series inductors or shunt capacitors, introduce reflected or coupled noise that is frequency dependent: the higher the signal frequency, the worse the noise.




Figure 2.18: Transmitter pre-emphasis using high-frequency boosting.

### 2.2.2.1 Transmitter Equalization

As mentioned earlier, channel frequency dependent loss causes ISI in the transmitted data. Equalization techniques have been used extensively in high-speed links in recent years to remove ISI [33–36]. Basically, an equalizer subtracts the ISI in the time domain or equivalently flattens the frequency response of the channel. Equalization can be performed in the transmitter or the receiver side. Each of these approaches offer some advantages. In this section we will discuss different techniques for transmitter equalization and introduce their advantages and disadvantages.

**FIR-Based Pre-Emphasis.** Equalization eliminates the problem of frequency-dependent attenuation by filtering the transmitted or received waveform so that the overall system exhibits a flat frequency response. For instance, in a transmitter equalizer, if the transfer characteristics of the channel is expressed by  $A(z)$ , the transmitter equalization transfer function,  $P(z)$ , should be designed such that  $A(z) \times P(z) = 1$  or  $P(z) = 1/A(z)$ , as shown in Figure 2.18. Often times it is not possible to implement the exact required  $P(z)$ ; however, there are techniques to closely approximate the target transfer function. Transversal filters (FIR filters) are mainly used to perform the transmitter equalization [37]. The transfer function,  $H(z)$  can be written as

$$H(z) = 1 + a_1z^{-1} + \dots + a_nz^{-n}, \quad (2.14)$$

where  $a_i$ 's are called the tap coefficients and  $n$  is the total number of equalization taps.  $n$  determines how well  $H(z)$  matches the target transfer function  $P(z)$ . The larger the number




Figure 2.19: Block diagram of a transmitter with  $m$ -tap FIR-based equalization.

of taps in the equalizer, the better the approximation of  $P(z)$  is achieved. This technique is very well suited for digital communication techniques, in which generating a delay is very straightforward through use of latches and flip-flops as shown in Figure 2.19. Figure 2.20 illustrates how increasing the number of taps in an FIR-based transmitter reduces the pre- and post-cursor ISI. Figure 2.21 shows the improvement in the transmitted eye-diagram quality for one and two taps of transmitter equalization. As a case study, we consider a 2-tap realization in this section.

**Continuous-Time Equalization.** As all channels exhibit a low-pass behavior, the equalizer has to create a high-pass transfer function,  $P(z)$ , to achieve an overall flat response. As mentioned earlier, FIR filter is one way to realize the high-pass behavior. Another way to provide high frequency transfer function is to employ continuous analog filters. As shown in Figure 2.22, these filters create a high frequency boost that can compensate for the frequency roll-off in the channel.




Figure 2.20: Pulse response of a channel before and after 2-tap transmitter FIR equalization.




Figure 2.21: Eye diagram at the end of a lossy channel using a transmitter with no tap of FIR equalization as well as 1-tap and 2-tap of equalization.

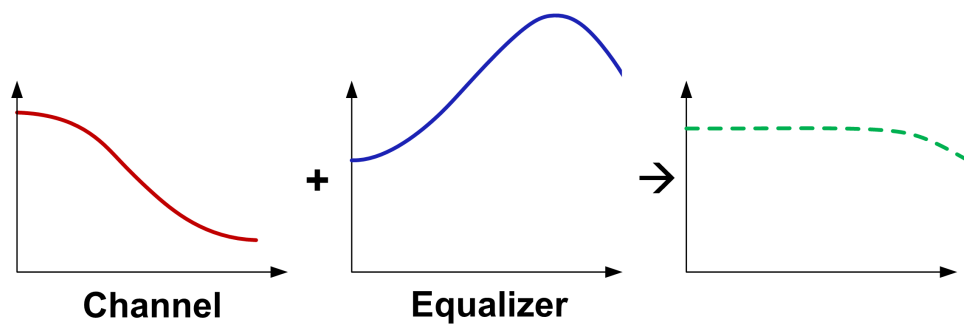


Figure 2.22: Transmitter equalization through high frequency boosting to achieve flat response.




Figure 2.23: Different high-frequency boosting configurations.

The main advantage of the continuous time equalization is that it mainly employs passive components and therefore has minimum power overhead. Figure 2.23 illustrates some circuit configurations employing inductors to perform high frequency boost. Detail analysis of these designs can be found in [38]. Passive continuous-time equalization is mainly employed in the receiver side, as will be explained in the next section, however transmitter equalization could be advantageous compared to receiver continuous equalization due to better noise performance. This is due to the fact that receiver high frequency peaking also amplifies high frequency noise where the signal strength is at its lowest, and hence degrades the overall SNR. The main drawback of this technique is the use of inductors, which inherently require large area. In the next chapter we will introduce a design that rectifies this problem.

### 2.2.3 Receiver

Receiver is one of the main building blocks in a communication link that resolves the data transmitted over the channel. As shown in Figure 2.24 a conventional receiver comprises of a pre-amplifier and a slicer that resolves the incoming data. The pre-amplifier is a low-noise stage with enough gain to provide a large enough SNR at the slicer input. In some designs, equalization is also incorporated in the pre-amplifier design to partly remove the ISI due to the channel dispersion. The slicer samples the output of the pre-amplifier in the middle of the bit-time and compares it with a threshold voltage chosen to be in the middle of the zero and one levels. The general requirements of a receiver in high-speed applications are high bandwidth, high gain, low noise and low power. The noise and offset of the comparator plus the coupled noise can increase the bit error rate significantly.

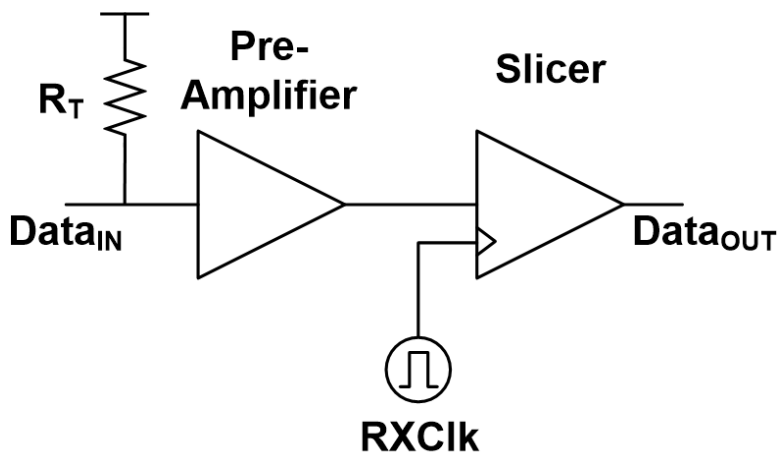


Figure 2.24: A typical electrical receiver block diagram.

Another important aspect of receiver design is the capability to compensate for the channel dispersion. We previously discussed the options for equalization at the transmitter side. However, as the data rate increases while the channel characteristics remain almost the same, the transmitter equalization becomes less than adequate for the excessive channel loss. Therefore designers have employed receiver equalization in conjunction with transmitter equalization to make high data rates possible over bandwidth limited channels. In the next section we will discuss different receiver equalization techniques.

### 2.2.3.1 Receiver Equalization

Receiver equalization is a powerful tool to compensate for the frequency dependent loss of the electrical channels. In this section we will introduce different techniques that are commonly employed by link designers to implement equalization at the receiver.

**Decision Feedback Equalization.** A common approach to remove ISI and enhance SNR is to employ decision feedback equalization (DFE). This technique helps to compensate for post-cursor ISI arising from the spread of a single pulse over time. The high level block diagram of a typical DFE is shown in Figure 2.25, [39]. The ISI from previous bits is compensated by adjusting the DFE taps:  $w_1, \dots, w_n$ . The delay elements taking on values of unit bit-times can be implemented using latches and flip-flops. In the DFE, weighted

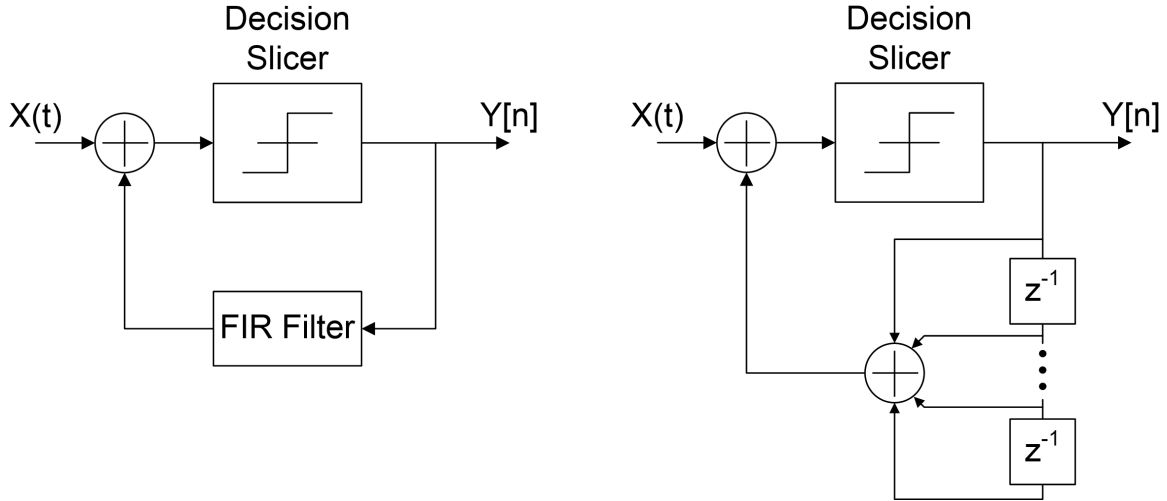


Figure 2.25: Simplified block diagram of a decision feedback equalizer.

versions of previous samples are added or subtracted to the main sample by a summer at the front-end.

The main problem in the design of a DFE is to meet the timing requirement for the first tap feedback loop. The constraint imposed by this critical signal path is that the sum of the slicer delay, the settling time of the summer, and the setup time of the latch needs to be less than one bit-time. Satisfying such a timing requirement becomes a difficult problem at increasingly high data rates. The problem can be mitigated by the application of speculative techniques, known also as loop-unrolling [40]. A block diagram of speculative DFE and the new critical path associated with this approach is shown in Figure 2.26. Basically, in this approach the outcomes for both possible values (one and zero) of the previous bit are fully resolved to digital levels. A multiplexer then chooses the correct value based on the resolved previous bit. Therefore, two latches are required to restore the digital levels. This technique can be further extended to the second equalization tap by resolving the incoming bit for four different cases possible for the previous two bits. Applying this technique to more taps, increases the complexity of the system exponentially. As a result, it is common to only apply it to the first tap. Recently, in order to enable data rates in excess of 25Gb/s, 2-tap [41] and 3-tap [42] loop-unrolling has been also employed.



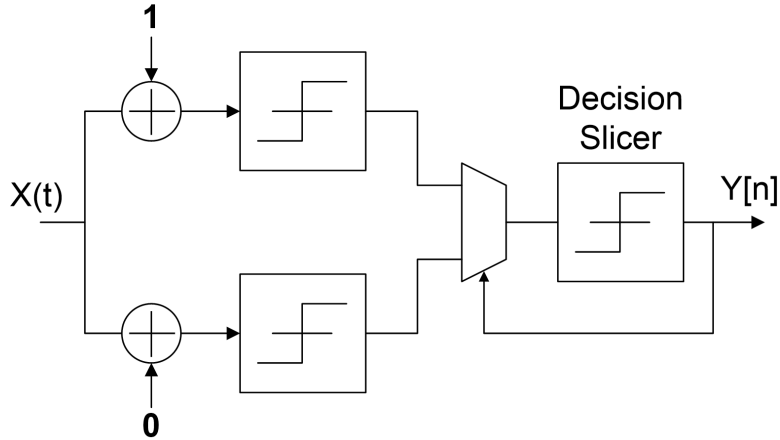


Figure 2.26: Loop-unrolling of the first tap of the DFE to remove the critical feedback loop.

**Feed-Forward Equalization (FFE).** Over the years one of the simplest and mostly used equalization techniques has been linear feed forward equalization (FFE). This technique usually involves the use of a linear transversal finite impulse response filter (FIR) as shown in figure 2.27. The FIR consists of adjustable tap coefficients  $w_1, \dots, w_n$  and a discrete or continuous unit delay,  $z^{-1}$  between each tap. The amount of delay,  $\tau$  that each delay cell represents can be as large as the bit-time [43], which is often referred as a symbol spaced equalizer. If  $\tau < T_b$  the equalizer is called a fractionally spaced equalizer (FSE) [44–46]. With proper selection of the tap gains this type of equalizer can be used to cancel precursor ISI, post-cursor ISI or both. Figure 2.22 illustrates how a weighted delayed version of the received pulse can be employed to remove post-cursor ISI. Similarly they can be used to compensate for pre-cursor ISI. Such a structure is very flexible, capable of handling any pulse response so long as a sufficient number of taps is provided. The choice of the delay cells could be either digital or analog. Digital FIR filtering requires a high-speed data converter. An alternative technique is to use sample and hold circuits as the tap delay line [43]. Another approach is to employ analog delay cells such as passive delay lines [44–46] or inductor-less active delay lines [47]. Passive delay lines employ inductor and variable capacitors to provide adjustability. They also achieve very high data rate. However, large on-chip inductors lead to a large die area. Inductor-less designs can achieve small area but they are limited to low data rates.

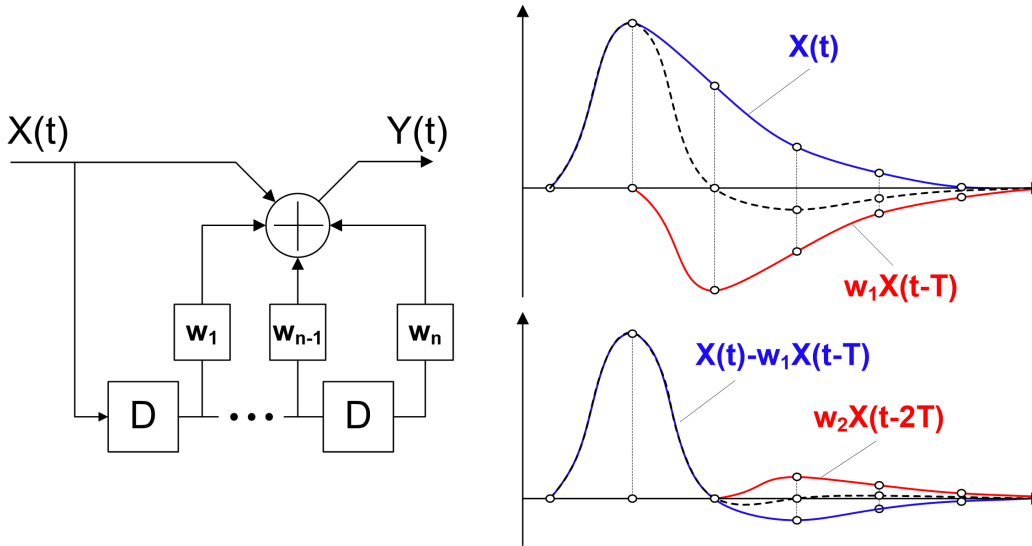


Figure 2.27: Feed-forward equalization at the receiver.

**Continuous Time Linear Equalization (CTLE).** Receiver equalization can also be implemented with a continuous-time amplifier that provides a high frequency boost. Usually the transfer function of these kinds of amplifier are adjustable to accommodate different channels. These amplifiers can be employed at the receiver front-end as the pre-amplifier to not only increase the signal level reaching the slicer, but to partly compensate for the channel loss. An example of such amplifier is shown in Figure 2.28. Here, programmable RC-degeneration in the differential amplifier creates a high-pass filter transfer function which compensates the low-pass channel [48, 49].

Another example of the continuous time linear equalizer is shown in Figure 2.29(a). This passive equalizer acts as a high-pass filter to compensate for the channel loss [49]. However, it introduces loss at low frequency which degrades the SNR at the slicer input. As a result, in most applications it is followed by an amplifier to improve SNR. The linear equalizer shown in Figure 2.29(b) employs a high frequency and a low frequency path to achieve high frequency boost [50].

While the implementation of the continuous-time linear equalizer is a simple and low-area solution, one issue is that the amplifier has to supply gain at frequencies close to the full signal data rate. This gain-bandwidth requirement potentially limits the maximum data

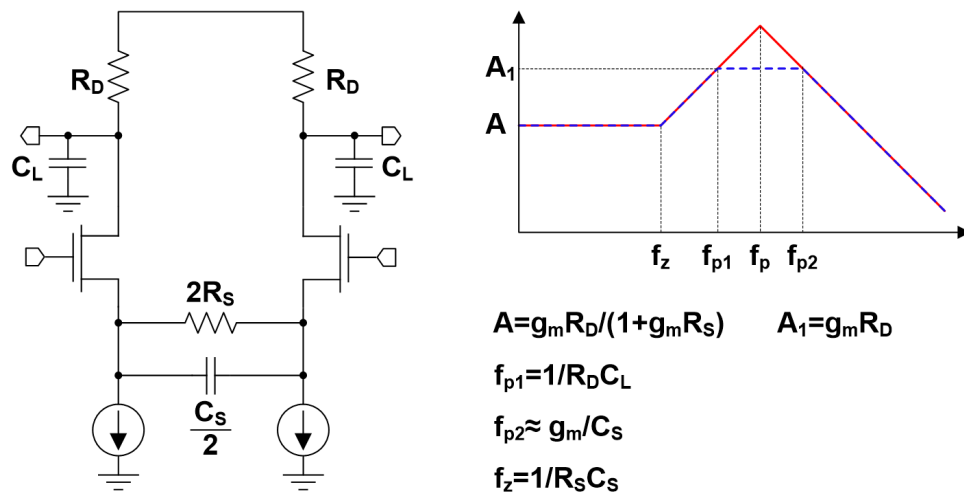


Figure 2.28: Continuous-time receiver equalization using frequency dependent source degeneration.

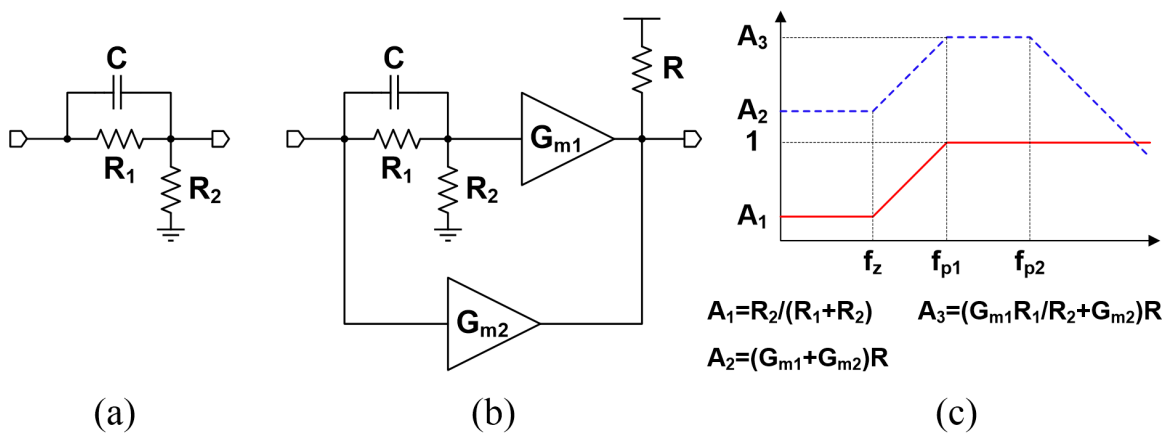


Figure 2.29: Passive high-pass filter equalizer (a). Dual path continuous-time equalizer (b).

rate, particularly in time-division demultiplexing receivers. In addition, the frequency boost introduced by the CTLE also amplifies the high frequency noise and degrades SNR at the receiver. This is especially problematic at high data rates, as the received signal is weak.

### 2.2.4 Timing Recovery

For maximum timing margin, the receiver should sample the bits in the middle of the data eye. The performance of the link is affected by how well the clock edge is positioned with respect to the incoming data stream. This clock position must be determined from the phase and frequency of the incoming data by the timing recovery circuit. In typical high-speed links, due to the process mismatches, time variations and undefined delays in the signal path, the received data can have an undefined phase and frequency. In a communication link either a separate clock signal can be transmitted along the data signal for timing information, which is known as forwarded-clock technique, or the clock should be recovered from the incoming data signal.

In a forwarded-clock link, the transmitted clock is employed in the receiver to perform data resolution. However, the mismatch between the phase of the data and received clock can cause detection error. As a result, a phase recovery system is required to adjust the sampling clock properly for best BER performance. On the other hand, in systems with clock recovery at the receiver, both clock frequency and phase need to be adjusted at the receiver.

As explained earlier, the transmitter's limited bandwidth and frequency dependent losses in the channel and reflections create inter-symbol interference (ISI). This ISI combined with the transmitter clock jitter and cross-talk adds amplitude noise and timing jitter to the received waveform. A clock recovery block should extract the clock component of the incoming signal and filter out the timing jitter. Finding the best sampling time with low variance and phase shift is particularly critical for bandwidth-limited channels to minimize the BER. In most systems direct estimation of a sampling time that minimizes the probability of error in the data resolution is not a practical task. Instead sub-optimal practical solutions are

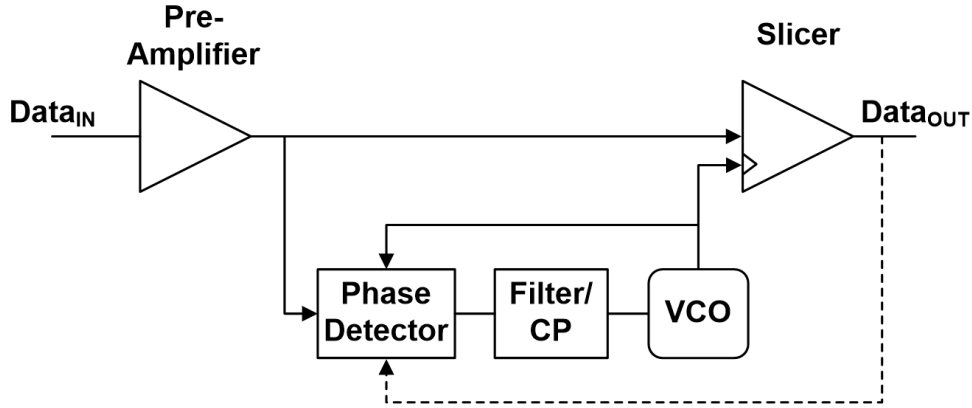


Figure 2.30: Clock recovery at the receiver.

adopted to define the best sampling point. The most common approach is to assume that the best sampling time is where the overall ISI is minimum. At this point the vertical eye-opening in the eye-diagram is maximum.

Figure 2.30 illustrates a typical timing recovery loop for a serial link. In case of phase mismatch between the clock and received data, the phase detector generates an error signal to adjust the VCO's phase and frequency. The output of the phase detector can be an analog signal for an analog loop or digital correction commands for a bangbang-controlled loop. In a bangbang-controlled PLL, the phase and frequency of the VCO are corrected by constant steps in two different directions depending on the decisions of the phase detector. The correction commands are called "Early" and "Late" signals. A decision-based phase detector is usually used in CDR techniques. In these CDRs, the data decisions are usually used to determine the type of transition that occurred and then use that information to find the correction needed. A number of different techniques for phase detection and definition of best sampling time have been proposed [51–54]. The over-sampling technique described in [53] takes three or more samples per each bit and performs the data resolution for all of them. Then, by looking at the sequence of resolved values, it decides which sample is the most reliable one, which is simply the one farthest from the transitions. This technique is very robust but requires a considerable hardware, power and area overhead for the oversampling. It also needs phase spacings that are at least three times shorter than a bit-period. For systems that run at data rates close to the technology limitations, such sampling schemes

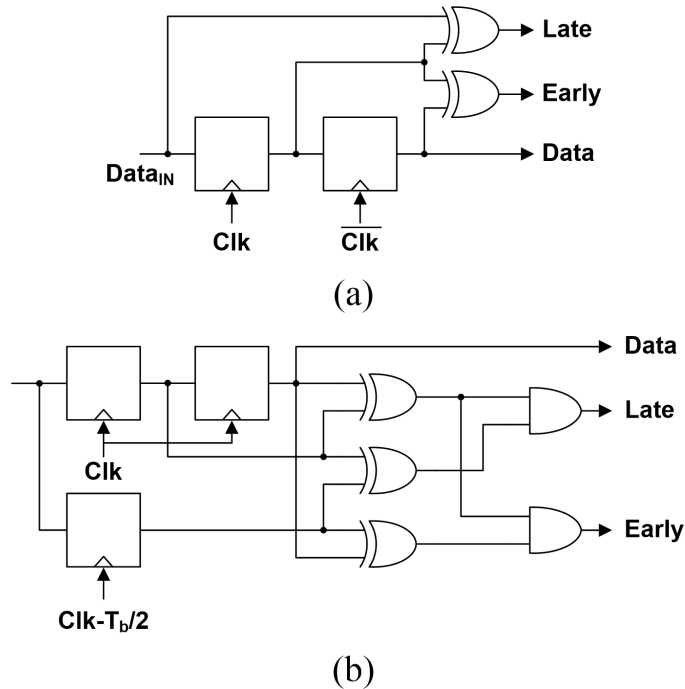


Figure 2.31: CDR phase detectors: (a) linear [51], (b) binary [52]

are not practical.

As shown in Figure 2.31, the phase detector can either be linear [51], which provides both sign and magnitude information of the phase error, or binary [52], which provides only phase error sign information. These phase detectors are also known as bangbang. While CDR systems with linear phase detectors are easier to analyze, generally they are harder to implement at high data rates due to the difficulty of generating narrow error pulse widths, resulting in effective dead-zones in the phase detector [54]. Bangbang phase detectors minimize this problem by providing equal delay for both data and phase information and only resolving the sign of the phase error [55].

## 2.2.5 Advanced Modulation Techniques

Modulation techniques which provide spectral efficiencies higher than simple binary signaling are an alternative solution to increase per channel data rate and achieve high aggregate bandwidth. These techniques have been implemented by link designers in order to increase




Figure 2.32: Simple binary (1bit/symbol) and PAM-4 (2bits/symbol) modulations.

data rates over band-limited channels. Multi-level pulse amplitude modulation (PAM), most commonly PAM-4, is a popular modulation scheme which has been employed in many designs [10, 56–58]. As shown in Figure 2.32, in PAM-4 modulation each symbol is comprised of two bits, which allows transmission of an equivalent amount of data in half the channel bandwidth. However, due to the transmitter’s peak-power limit, the voltage margin between symbols is 3X (9.5dB) lower with PAM-4 versus simple binary signaling. As a result, in order to benefit from the bandwidth efficiency of this modulation scheme, while not sacrificing the overall SNR the channel loss at the PAM-2 Nyquist frequency should be greater than 10dB relative to the previous octave [59]. However, this rule can be somewhat optimistic due to the differing ISI and jitter distribution present with PAM-4 signaling [24].

Another modulation format investigated by link designers is the use of multi-tone signaling. This type of signaling is commonly used in systems such as DSL modems [60], and has been recently employed for high-speed inter-chip communication applications [61]. In contrast with conventional baseband signaling, multi-tone signaling breaks the channel bandwidth into multiple frequency bands over which data is transmitted. This technique has the potential to greatly reduce equalization complexity relative to baseband signaling due to the

reduction in per-band loss and the ability to selectively avoid severe channel nulls. Typically, in systems such as modems where the data rate is significantly lower than the on-chip processing frequencies, the required frequency conversion is done in the digital domain and requires DAC transmit and ADC receive front-ends [60]. While it is possible to implement high-speed transmit DACs [61], the excessive digital processing and ADC speed and precision required for multi-Gb/s channel bands results in prohibitive receiver power and complexity. Thus, for power-efficient multi-tone receivers, researchers have proposed using analog mixing techniques combined with integration filters and multiple-input multiple-output (MIMO) DFEs to cancel out band-to-band interference [62].

Serious challenges exist in achieving increased inter-chip communication bandwidth over electrical channels while still satisfying I/O power and density constraints. As discussed, current equalization and advanced modulation techniques allow high data rates over severely band-limited channels. However, this additional circuitry comes with a power and complexity cost. The demand for higher data rates will result in increased equalization requirements and further degrade link energy efficiencies. These issues motivate investigation of novel low-power and compact equalization techniques to enable highly parallel data communication for chip-to-chip applications, which is discussed in the next section.

## 2.3 Summary

Due to the rapid scaling of the CMOS technology, computational capacity has increased drastically over the past decades. This has raised a great demand for efficient data communication links between computational cores. Electrical signaling is the most natural approach for interconnecting integrate circuits and achieves high level of integration; however, copper channels suffer from huge loss, reflection and crosstalk that limit the maximum possible data rate. The frequency dependent loss of the channel is mainly due to the skin effect and the dielectric loss, which create severe inter-symbol-interference and degrade signal integrity. As we discussed in this chapter, advanced timing recovery techniques as well as equalization techniques have allowed the maximum data rates to continue to scale. However these



techniques heavily add to the complexity, power consumption and area of the transceivers.

Parallelism is a natural way to achieve high aggregate data rates to meet the IO bandwidth requirements of the future computational and storage systems such as high performance processors and data centers. However, highly parallel links demand low power consumption per IO and additionally an efficient way to combat the excessive crosstalk induced from neighboring channels. Therefore, ultra low-power crosstalk cancellation techniques are necessary to allow for a high level of parallelism.

In the next chapter we introduce a transceiver system with a novel ultra low-power transmitter and receiver equalization and crosstalk cancellation scheme suitable for highly parallel communication links.

## Chapter 3

# High-Speed Low-Power Electrical Transceivers

As mentioned earlier, the rapid scaling of CMOS technology continues to increase the processing power of microprocessors and the storage volume of memories. This increases the need for high bandwidth interconnection between chips, which can be achieved by employing large numbers of input and outputs (IOs) per chip as well as high data rates per IO. Although CMOS scaling has increased the speed of transistors, the scaling of interconnect bandwidth has proven to be very difficult. The dielectric and resistive losses of printed circuit board (PCB) traces increase as the operation frequency increases. Such frequency dependent attenuation causes inter-symbol interference (ISI) and ultimately signal-to-noise-ratio (SNR) degradation. In addition, reflections from discontinuities in the signal path generate more ISI and further reduce the SNR. These problems are exacerbated as the data rate increases.

As explained in the previous chapter, a common approach in the design of high-speed serial links over bandwidth-limited channels is to employ equalization techniques, including DFE [64–67] and continuous time linear equalization [68–71] at the receiver and feed-forward equalization (FFE) at the transmitter [43–47]. These approaches can be used in parallel links with many IOs to increase the aggregate data rate; however the number of links will be eventually limited by area, noise, and power consumption of the IO. In addition, employing a large number of traces increases the electromagnetic coupling between adjacent traces and results in crosstalk, which degrades SNR at the receiver. A number of techniques

have been proposed to remove the effects of crosstalk. The design in [73] employs an FFE equalizer and [74] and [76] use crosstalk-induced jitter equalization at the transmitter and receiver, respectively. Other approaches to compensate for crosstalk noise include the use of staggered I/Os [77] or a finite-impulse response (FIR) filter at the transmitter [78–80]. All these schemes result in significant power consumption and are not suitable for parallel data links.

In this chapter, we introduce novel techniques to implement receiver and transmitter equalization as well as crosstalk cancellation in a low-power and compact manner to enable highly parallel electrical communication links that achieve high aggregate data rate as a solution to the ever increasing demand for IO bandwidth requirements. We will start with the low-power receiver architecture with far-end crosstalk cancellation and introduce a novel implementation of decision feedback equalization using a switched-capacitor (SC) technique to enable realization of many taps with better power efficiency compared to conventional techniques [83, 84]. We next elaborate on the crosstalk cancellation scheme and explain how the SC technique can be employed to implement this technique with minimal power overhead. Circuit implementation and experimental results will be discussed next.

In the second part of this chapter we introduce a transmitter equalization technique suitable for low-power high data rate application using passive linear equalization. Adjustable equalization is performed employing controllable peaking inductors realized in a 3-D structure to minimize area. We will compare this technique with the conventional FIR-based equalization and conclude this chapter with the CMOS implementation of the proposed transmitter and experimental results.

### 3.1 Receiver Architecture

A diagram of a typical parallel communication link is shown in Figure 3.1. Many ICs are connected via PCB traces. The loss in these traces, as well as the electromagnetic crosstalk due to adjacent traces, significantly degrades the signal integrity in such systems. A common approach to removing the ISI induced by the bandwidth-limited channel is to employ DFE

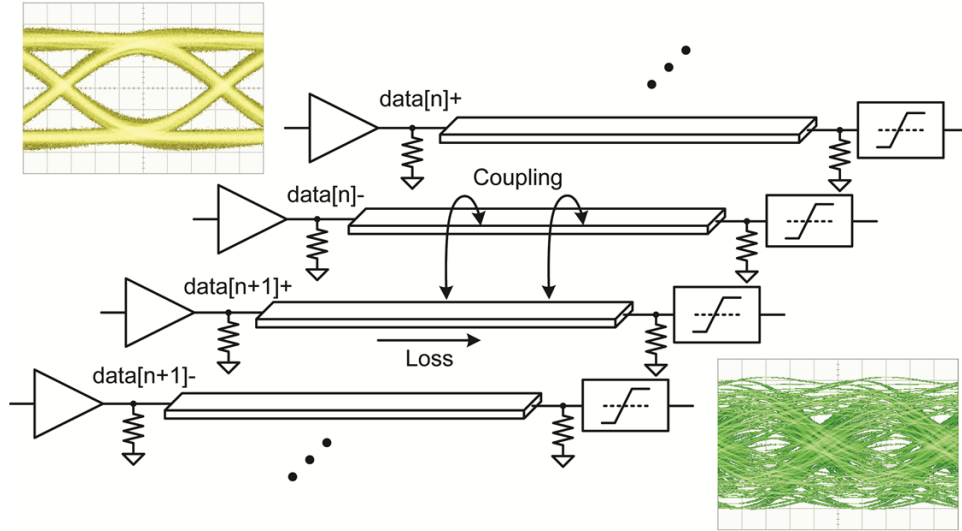


Figure 3.1: A typical parallel communication link. Signal integrity is degraded due to channel dispersion and crosstalk noise of the neighboring channels.

at the receiver. This technique helps to compensate for post-cursor ISI arising from the spread of a single pulse over time.

Figure 3.2 shows the top level architecture of the receiver. It employs a 2-tap DFE to compensate for the post-cursor ISI. It will be shown later that this architecture can be readily generalized to n-taps of equalization. As explained in previous chapter, in the DFE, weighted versions of previous bits are added to or subtracted from the main sample by a summer at the front-end. As shown in Figure 3.2, in the 2-tap DFE, the ISI from the previous two bits is compensated by adjusting taps  $H_1$  and  $H_2$ . To avoid timing issues in the first tap equalization path, a speculative DFE architecture is employed, in which the outcomes for both possible values (one and zero) of the previous bit are fully resolved to digital levels. A multiplexer then chooses the correct value based on the resolved previous bit. Therefore, two clocked latches are required to restore the digital levels. The design employed in this paper uses a combined analog multiplexer and latch to select one of the two analog voltages directly at the output of the summer and resolve the digital bit [64]. The power consumption associated with latches is reduced since a single latch is embedded into the multiplexer, as shown in Figure 3.2.

In this design a half rate clocking technique is adopted to further reduce power con-

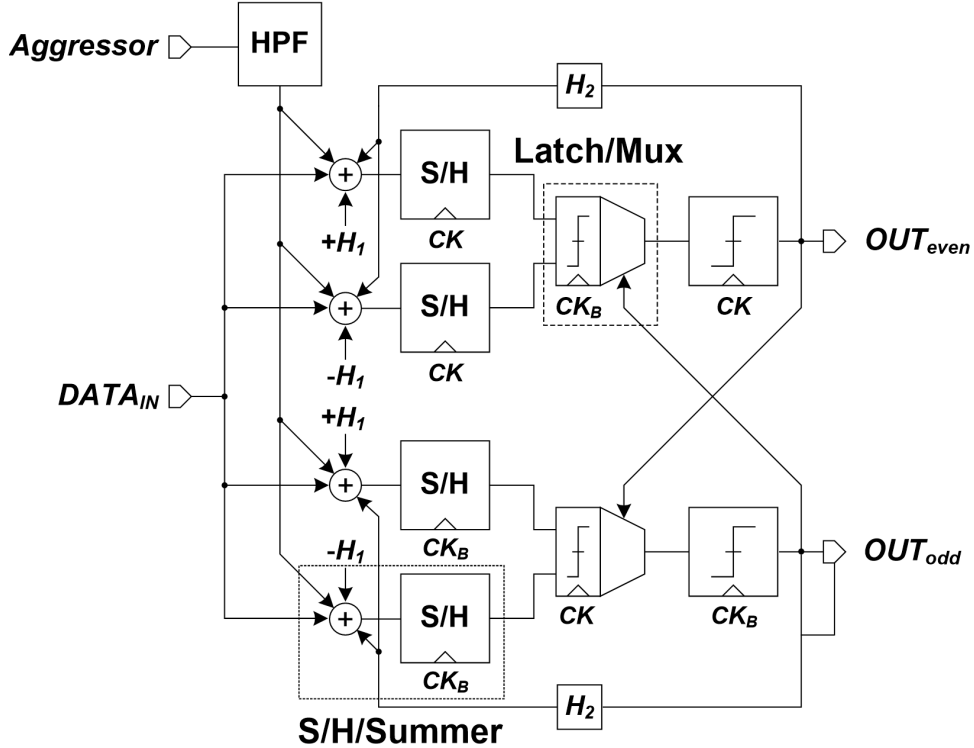


Figure 3.2: Top level architecture of the proposed receiver employing half-rate clocking and loop-unrolling.

sumption [82]. This approach helps to relax the DFE timing constraint and the design requirements of the clocking and the slicers as they will be operating in a fraction of the data rate. The inherent demultiplexing enabled by this technique helps to save power in the following digital stages. Crosstalk cancellation in the proposed design is performed via a novel, low-power technique in which the aggressor signal is used to reproduce the coupling noise, as shown in Figure 3.2. Further details of this technique will be provided later in this chapter.

## 3.2 Decision Feedback Equalization Implementation

As discussed in the previous section, a DFE operates based on adding weighted information from previous bits to the received signal. At high data rates, the design of a linear and precise summer that meets the DFE timing requirement is challenging. In particular, the power associated with the summer becomes critical in a half rate loop-unrolled architecture

where the number of summers is large.

### 3.2.1 Current-Mode Summers

Conventionally, taps of the equalizer are implemented using current-mode summers [11, 12, 48, 85]. Figure 3.3(a) illustrates the schematic of a current summer. The currents from each tap are set by the tap coefficient and added at the output nodes. The resulting output will be equal to  $R(I_1 + I_2 + \dots + I_n)$ . This configuration is quite power hungry due to its large static current. In addition, the power consumption of the summer increases proportionally with the number of taps. The extra capacitance due to equalization taps also introduces additional loading at the summer output and limits its bandwidth. For a resistive summer, the settling time is determined largely by the load resistor and load capacitance at the output of the summer. As approximately three RC time constants are needed to achieve 5% settling accuracy, the load time constant is set such that

$$T_b = 3R_D C_L, \quad (3.1)$$

where  $C_L$  is the total output capacitance. As a result, assuming complete switching of the input transistors, all the current steers into one side and the resulting voltage at the input will be equal to

$$V_{OUT} = R_D I_t = \frac{T_b I_t}{3C_L}, \quad (3.2)$$

where  $I$  is the tail current. As can be observed from this equation, in order to maintain the output voltage while adding more equalization taps, the tail current and hence overall power consumption has to be increased.

To mitigate the power consumption issue of the current-mode summer, designers have recently proposed current-integrating summers [66, 67], in which the resistor in the current-mode summer is replaced by a switch as shown in Figure 3.3(b). This summer operates in two phases. In the first phase the switches are turned on and pre-charge the output

nodes. In the second phase the switch is turned off and the current sum of the equalization tap currents are integrated over the parasitic capacitance at the output node, which will be resolved by the next stage slicer. The developed output differential voltage represents the incoming data with post-cursor ISI canceled based on previous bit decisions through equalization taps. The operation of the current-integrating summer is shown in Figure 3.3(c). This configuration mitigates the bandwidth problem of the current-mode summer through the integration process as no settling time requirement exists. To get a better sense of the power benefit of this configuration we will calculate the minimum tail current source required to achieve the same sensitivity as the current-mode summer. Assuming complete switching, the tail current is steered on one side and is integrated over  $C_L$ . This result in the output voltage equal to

$$V_{OUT} = \frac{T_b I_t}{C_L}. \quad (3.3)$$

Comparing Equations 3.2 and 3.3, we can observe that a power reduction factor of three is achieved through integrating summer. However, it should be noted that increasing the number of taps increases the total capacitance at the output node and hence degrades sensitivity for a certain current as the output voltage is inversely proportional to the output capacitance.

### 3.2.2 Charge-Mode Summers

In the previous chapter we showed that current-mode summers are becoming inadequate for implementing decision feedback equalization in terms of power efficiency. In particular, as we start to employ techniques such as loop-unrolling and fractional-rate clocking to enable the operation of DFE close to the theoretical limit of  $4 \times \text{FO-4}$  [18], the number of summers per receiver increases, which makes its power consumption an important performance metric. It was shown that the current-integrating summers can help reduce the power consumption by a factor of three compared to conventional current-mode summers. However, as they still need to consume static current to perform summation, this still leaves the power consumption

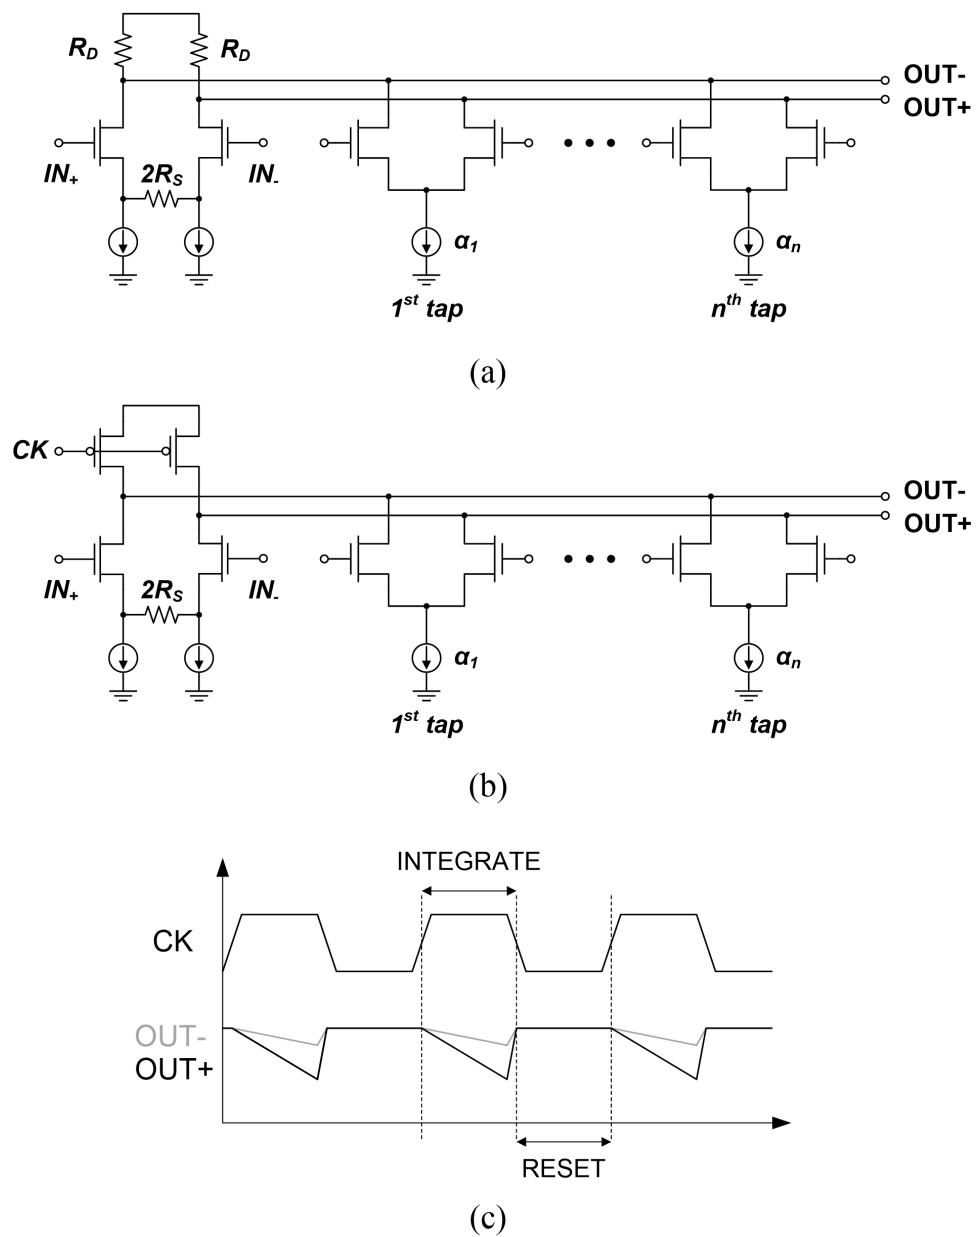


Figure 3.3: Current-mode summer employed in conventional DFE (a). Current-integrating summer (b). Operation of the current-integrating summer (c).



issue.

Operating in charge domain can greatly help in reducing power consumption [86–88]. In a charge-mode configuration, summation is performed via charge accumulation over capacitors through some switches, which does not involve any static current consumption. The only source of power consumption arises from turning the switches on and off. This introduces dynamic power dissipation, which can be much smaller than the static power associated with current-mode summers. In addition, dynamic power scales well with technology, which makes it a desirable solution for advance technologies. In the next section we will introduce a novel charge-mode summation technique which is capable of performing many taps of equalization with a performance better than its current-mode counterpart while offering lower power consumption.

### 3.2.2.1 4-tap Realization

In this section we first introduce a 4-tap realization of the proposed charge-mode switched capacitor summer, in which two taps are employed for equalization and two taps for crosstalk cancellation. Then we show how this technique can be generalized to implement  $n$  taps of equalization. In the proposed architecture, an SC front-end, which is denoted as S/H/summer, is employed to sample the input signal and combine it with the feedback coefficients as shown in Figure 3.2. Figure 3.4(a) shows the implementation of the S/H/summer. The single-ended version is shown for simplicity. The S/H/summer operates in two phases, shown in Figure 3.4(b), (c). In the first, sample/sum phase, the input is sampled into capacitor  $C_{S1}$  and the first tap coefficient,  $H_1 = \alpha_1 V_{REF}$ , is added to (or subtracted from) this sample. During this phase, as will be discussed later, the crosstalk canceling signal is stored into capacitor  $C_{S2}$ . As a result, at the end of this phase, the voltage across  $C_{S1}$  and  $C_{S2}$  will be respectively equal to

$$V_{CS1} = \alpha_1 V_{REF} - V_{IN}, \quad (3.4)$$

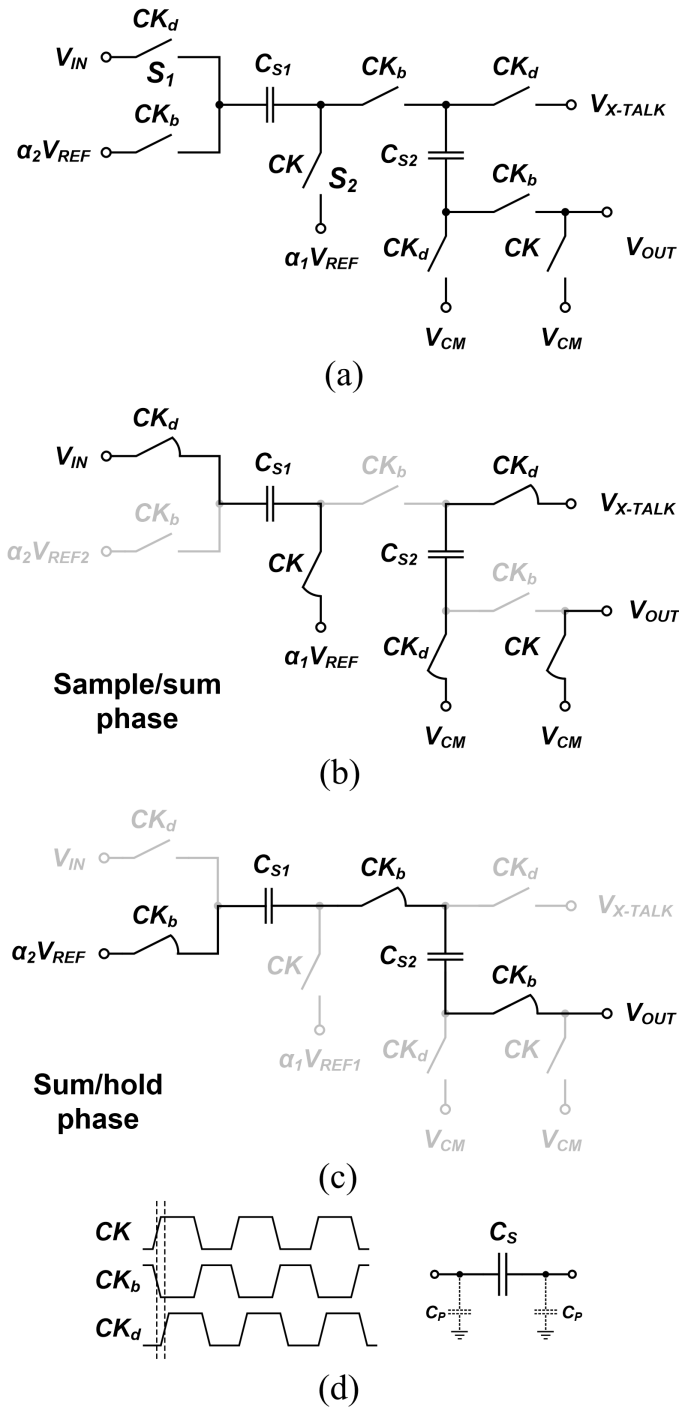


Figure 3.4: (a) Circuit-level implementation of the front-end S/H/summer operating in two phases: (b) sample/sum phase, (c) sum/hold phase (single-ended version is shown for simplicity).

$$V_{CS2} = V_{CM} - V_{X-TALK}, \quad (3.5)$$

where  $V_{X-TALK}$  is the mimicked crosstalk noise, explained in Section 3.3,  $V_{CM}$  is the common-mode voltage, and  $V_{IN}$  is the input signal. In the second, sum/hold phase, the result of the first phase is added to the second tap coefficient,  $H_2 = \alpha_2 V_{REF}$ , and applied to the next stage comparator. This is done by putting the two sampling capacitors in series which results in the output being equal to

$$V_{OUT} = \alpha_2 V_{REF} + V_{CS1} + V_{CS2} = \alpha_1 V_{REF} + \alpha_2 V_{REF} + V_{CM} - V_{IN} - V_{X-TALK}. \quad (3.6)$$

In order to minimize the input-dependent charge injection, the switch that connects the sampling capacitor to the input,  $S_1$ , is turned off slightly after  $S_2$ . This is possible through employing the delayed version of the clock,  $CK_d$ . A key aspect of the design of the SC DFE is the sizing of the sampling capacitors. Optimizing this sizing demands an understanding of how loss, SNR, and power change as a function of the sampling capacitor size, requiring an analysis of how it affects parasitic capacitance, receiver noise, and clocking power. When this technique is implemented with capacitors and switches, the parasitic capacitances, shown in Figure 3.4 as  $C_p$ , cause signal gain degradation due to charge sharing. Therefore, the resulting gain from the input to the S/H/Summer output is equal to

$$A_{sampler} = \frac{C_{S1}C_{S2}}{(C_{S1} + C_{S2} + 2C_P)(C_{S2} + C_P + C_L) - C_S^2}, \quad (3.7)$$

where  $A_{sampler}$  is the signal gain,  $C_p$  is the parasitic capacitance and  $C_L$  is the load capacitance from the next stage.  $C_p$  is due to the switch junction capacitance, sampling capacitors top and bottom plate parasitic, and the wiring capacitance. As a result, it almost linearly changes with the sampling capacitor size. Figure 3.5(a) shows how the signal gain changes with the sampling capacitors sizes for the proposed 4-tap S/H/summer. The loading from the next stage (slicer/MUX) is estimated to be about 2fF. The sensitivity of the receiver is

mainly determined by the performance of the front-end S/H/summer. In order to achieve reasonable SNR, the signal loss and the noise must be kept as low as possible. Contributing to the input-referred circuit noise are the slicers and samplers in the receiver. The slicer is modeled as a sampler with gain and has an input referred voltage noise variance of [89]

$$\sigma_{slicer}^2 = \frac{2kT}{A_{slicer}C_A}, \quad (3.8)$$

where  $C_A$  is the total capacitance at the internal slicer node ( $V_{int}$  in Figure 3.12), which is approximately 8fF.  $A_{slicer}$  is the gain from the input of the slicer to the internal node,  $V_{int}$ , which is estimated to be equal to two. This results in a slicer voltage noise sigma of  $0.5mV_{rms}$ . Sampler voltage noise variance is equal to

$$\sigma_s^2 = \frac{2kT}{C_{S1}} + \frac{2kT}{C_{S2}}. \quad (3.9)$$

Clock jitter also impacts the receiver sensitivity because any deviation from the ideal sampling time results in a reduced sampled voltage. This timing inaccuracy is mapped into an effective voltage noise on the sampled input signal with a variance of

$$\sigma_{clk}^2 = \sigma_j^2 \times R_b^2, \quad (3.10)$$


where  $\sigma_j$  is the clock jitter and  $R_b$  is the rate of the input voltage change around the sampling point. The timing noise is estimated to be equal to  $0.75mV_{rms}$ . The total noise is equal to

$$\sigma_n = \sqrt{\sigma_{slicer}^2 + \sigma_s^2 + \sigma_{clk}^2}, \quad (3.11)$$


which results in SNR equal to

$$SNR = \frac{V_b A_{sampler}}{\sqrt{\sigma_{slicer}^2 + \sigma_s^2 + \sigma_{clk}^2}}, \quad (3.12)$$


where  $V_b$  is the eye opening amplitude after equalization taps are applied. Simulation shows that if a  $400mV_{pp}$  differential input at 15 Gb/s is transmitted over a 5" FR-4, 2 mV of eye



(a)



(b)



(c)

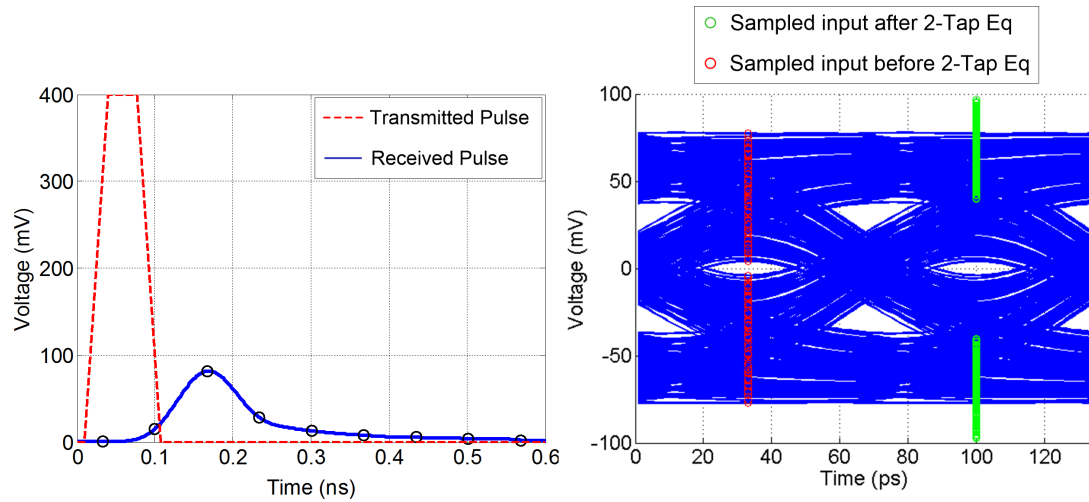
Figure 3.5: S/H/summer performance in the case of  $n=2$ . S/H/summer voltage loss,  $A_{\text{summer}}$  (a). SNR at the output of the S/H/summer (b). SNR normalized to clocking power consumption (c). To achieve high SNR (hence high BER) and power efficiency while maintaining the DFE speed,  $C_{S1}$  and  $C_{S2}$  are chosen to be equal to 19fF and 14fF, respectively.

opening amplitude is generated, as shown in Figure 3.6, which can be improved to 40mV using two taps of equalization. Replacing 40 mV in Equation 3.12 results in the S/H/summer SNR plot shown in Figure 3.5(b). The SNR increases with the sampling capacitor sizes. However, in order to maintain bandwidth, the corresponding switch size also has to be increased, which affects the sampler performance by adding more parasitic capacitance and hence increasing the signal loss. Furthermore, to drive a large switch, a large clock buffer is required that increases the overall power consumption. This power is proportional to the capacitance driven by the buffers. If we ignore the capacitance introduced by the buffer itself, the total clock buffer power consumption can be written as


$$P_{clk} = kC_{s-tot}V_{dd}^2f + kC_{slicer}V_{dd}^2f, \quad (3.13)$$

where  $k$  is the activity factor equal to 0.5 for clock buffers,  $C_{s-tot}$  is the total capacitance introduced by the S/H/summer switches,  $C_{slicer}$  is the capacitance in the slicer that is driven by clock buffers, and  $f$  is the operating frequency. Figure 3.5(c) shows the SNR normalized to the power consumption of the clock buffers. We employed this metric, along with the SNR, to choose the sampling capacitor sizes. In order to achieve  $BER < 10^{-12}$  ( $SNR > 7.1$ ) and low clocking power consumption in this design, the sampling capacitors ( $C_{S1}$ ,  $C_{S2}$ ) are optimized to be about 20fF, however, in order to meet the settling time requirements within a bit interval, according to simulations,  $C_{S1}$  and  $C_{S2}$  are chosen to be equal to 19fF and 14fF, respectively. As the  $\frac{SNR}{P_{clk}}$  is quite flat around the maximum point, the SNR and power penalty is small.

Two 4-bit current-steering digital-to-analog converters (DACs) generate the differential equalization coefficients ( $\alpha_1V_{REF}$ ,  $\alpha_2V_{REF}$ ) as shown in Figure 3.7. In this work, equalization tap coefficients are adjusted manually, however, on-chip adaptation algorithms such as SS-LMS [90, 91] and BER-based adaptation [92] can be utilized to optimize the tap values. The two DACs are designed to deliver enough current drive without considerable voltage drop while consuming less than  $900\mu A$ . A 1pF bypass capacitor at the DAC output filters noise and kickback charge from the switches. During each bit interval, the DAC should



(a) (b)



(c)

Figure 3.6: Simulated input eye diagram for a 5" FR-4 trace at 15Gb/s. Red and green circles show the sampled input before and after 2-tap decision feedback equalization, respectively (a). Histogram of the sampled input before and after equalization. 2-tap DFE improves eye opening from 2mV to 40mV (b).

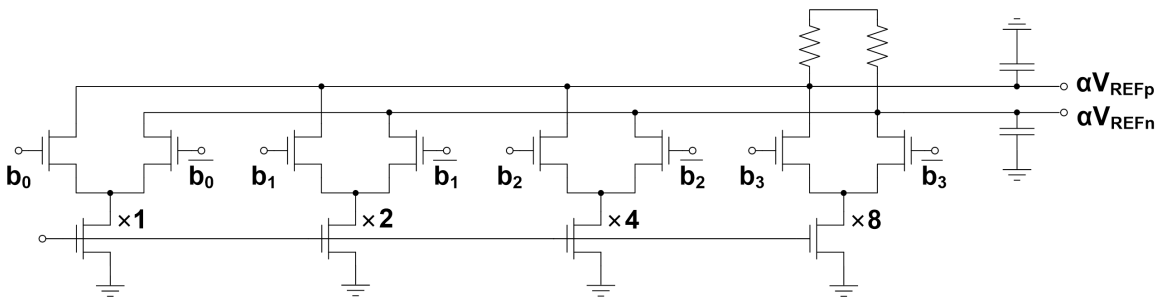


Figure 3.7: 4-bit current steering DAC that generates DFE tap coefficients.




Figure 3.8: Switched capacitor summer for 2n-tap, sample/sum phase (a), sum/hold phase (b).

charge/discharge one sampling capacitor,  $C_S \approx 20\text{fF}$ , from the S/H/summer, which causes a very small charge sharing with the bypass capacitor. In a differential implementation the voltage change due to charge sharing appears as a common-mode noise to the first order and hence does not affect the DFE performance significantly.

### 3.2.2.2 Generalization of Switched Capacitor Summers

The SC summation technique can be further extended to realize a large number of taps. Figure 3.8(a), (b) show the operation of a 2n-tap S/H/summer, where  $\alpha_k V_{REF}$  represents the weighted version of the  $k^{th}$  previous bit. In the first phase, the input and 2n-1 taps of the DFE ( $\alpha_1 V_{REF}, \dots, \alpha_{2n-1} V_{REF}$ ) are sampled and stored over n sampling capacitors ( $C_{S1}, \dots, C_{S(2n-1)}$ ). The capacitors are connected in series in the second phase, as shown in Figure 3.8(b) where one side is connected to the  $2n^{th}$  tap coefficient,  $\alpha_{2n} V_{REF}$ , and the other side is the output of the S/H/summer. The resulting output at the end of the second phase is equal to

$$V_{OUT} = V_{IN} + (\alpha_2 V_{REF} + \dots + \alpha_{2n} V_{REF}) - (\alpha_1 V_{REF} + \dots + \alpha_{2n-1} V_{REF}), \quad (3.14)$$

which is the sum of the sampled input and all feedback coefficients. Coefficients can be negated simply by swapping the differential signals in a differential implementation. As mentioned earlier, due to charge sharing between the sampling capacitors and the parasitics, the voltage gain in practice will be less than unity. In order to minimize the effect of parasitics and charge sharing on gain, the capacitor that samples the input signal is chosen to be the




last capacitor in the chain,  $C_{S1}$ . Figure 3.9(a) shows how the voltage gain decreases with increasing the number of taps.


In order to compensate for the voltage gain degradation for a larger number of taps, the input can be sampled using two sampling capacitors, which results in a factor of two increase in voltage gain, as shown by the dashed line in Figure 3.9(a). We denote this technique as dual-path sampling. The factor of two increase in gain comes as a result of sacrificing one of the DFE taps to sample the input. It also increases the input loading, but since the input is 50-ohm terminated, the additional loading due to the extra sampling capacitors has a negligible effect on the overall channel insertion loss. On the other hand, the sampling capacitors can be sized smaller to compensate for the loading effect. Scaling the sampling capacitors results in parasitic reduction and hence the charge sharing effect remains the same. As noise is proportional to the square root of the capacitor size, the same input loading can be achieved by reducing the sampling capacitors by a factor of 2 and still achieving  $\sqrt{2}$  improvement in SNR. This also helps in decreasing the clocking power consumption as smaller switches are required for a given data rate. For comparison purposes, Figure 3.3(a) shows the current-mode summer conventionally employed for tap summation in DFE designs. The summer voltage gain is equal to

$$A = \frac{R_D}{\frac{1}{g_m} + R_S}, \quad (3.15)$$


where  $g_m$  is the input transistor transconductance. To maintain the bandwidth while increasing the number of taps, load resistance should be reduced. This causes the voltage gain to decrease for a constant  $\frac{I}{g_m}$  in the main tap. The dotted line in Figure 3.9(a) shows how the voltage gain changes by increasing the number of taps while keeping the main tap current constant assuming each tap adds ten percent extra capacitance to the summing node [71]. To maintain the voltage gain and the bandwidth regardless of the number of taps, larger  $g_m$  and hence more power consumption is required in this technique. The reduction in voltage gain due to increasing the number of taps affects the receiver sensitivity by degrading the SNR. Figure 3.9(b) shows the normalized change in the front-end SNR with the number



(a)

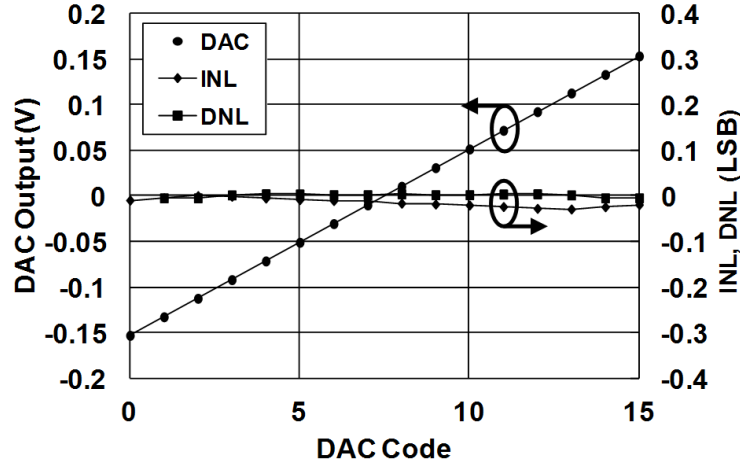


(b)




(c)

Figure 3.9: Signal gain comparison between SC and current-mode summer (a). Normalized SNR for the SC and current-mode summer (b). Tap coefficient gain for different number of taps (c).



(a)



(b)

Figure 3.10: Linearity of the post-cursor taps for an eight-tap (a) switched capacitor and (b) current-mode summer.

of DFE taps for the SC and current-mode summers. For the current-mode summer it is assumed that any additional tap contributes ten percent of the original noise [71].

Another important aspect of the post-cursor tap summer is its linearity. As the gain of the summer for each tap is determined by the ratio of capacitors, the only source of non-linearity is the parasitic capacitors introduced by transistor switches. For the tap voltage range that the SC summer operates, simulation shows that the parasitic capacitance (due to junction capacitance) changes about 10%. However, as mentioned earlier, the parasitic capacitance is comprised of sampling capacitor parasitics, interconnects and switches. As

a result, the overall parasitic capacitor change is less than 5%. Figure 3.10(a) shows the linearity simulation results for an 8-tap SC summer including the tap coefficient DAC. In the current-mode summer the main source of non-linearity is due to the tail current source that generates  $\alpha_n$ . As the tap differential pair transistors are sized small to add minimum parasitic to the output node, at high currents, the voltage across the tail current source reduces drastically. This causes reduction in the tail current due to channel length modulation. Figure 3.10(b) shows the linearity simulation results for an 8-tap current-mode summer where each tap transistor size is about 10% of the main tap to minimize loading [71] (Both in Figure 3.10(a), (b) transistor mismatch is not accounted for).

The tap coefficients gain in SC implementation is also a function of both number of taps and the tap number. Figure 3.9(c) shows how the tap gains change with the number of taps and also the tap number. The taps sampled by the first capacitor in the chain,  $C_{Sn}$ , experience larger loss compared to the taps sampled by the last capacitor in the chain,  $C_{S1}$ . As a result, the DAC that generates  $\alpha_{2n-1}$  should provide a larger voltage compared to that of  $\alpha_1$ . On the other hand, post-cursors generated in a realistic channel decay with the tap number, that is,  $\alpha_{2n-1}$  is inherently smaller than  $\alpha_1$ . This can compensate the lower gain of the tap coefficients sampled by the first capacitors in the chain.

### 3.2.3 Comparator Design

As explained before, even though the speculative architecture relaxes the timing requirement for the feedback path, the extra hardware increases the area and power. Figure 3.11 illustrates the top-level architecture of the DFE in which each path requires two comparators to perform data resolution on the incoming bit for a previously received zero and one. These slicers are followed by a digital multiplexer to choose the resolved data based on the previous bit. In order to reduce the DFE area and power the two slicer can be combined using an analog multiplexer as shown in Figure 3.12. Since the MUX is clocked with full-swing clock signals, the clock switch transistors are operating in the linear region, and the stacking of the transistors is easier given the limited voltage supply. In addition, as the tail transistor

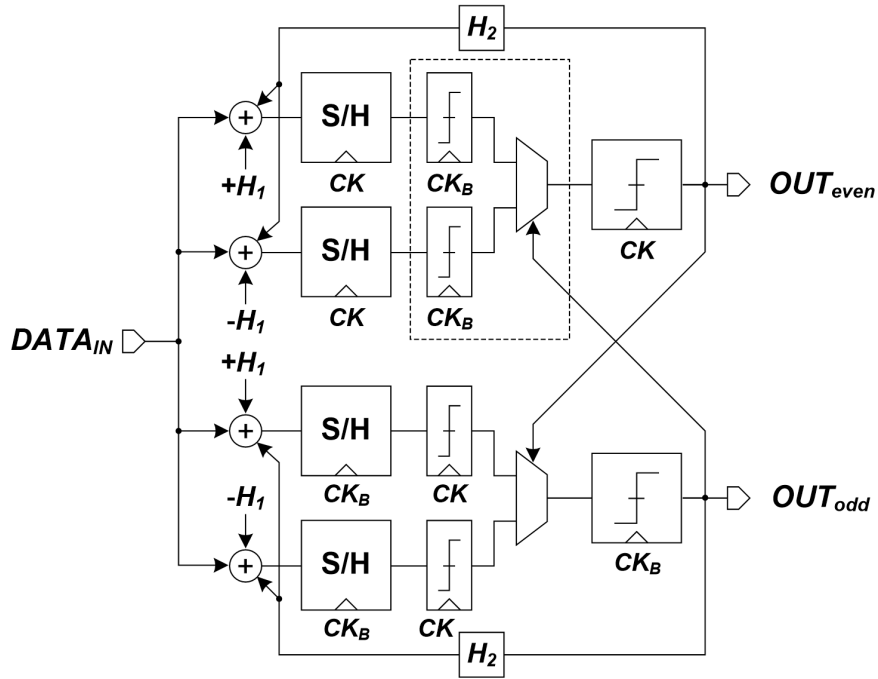


Figure 3.11: 2-tap DFE architecture with loop-unrolling.

and the MUX transistors are not in the DFE critical loop, they can be sized to be relatively large to allow for proper operation of this stage at low supply voltages. In order to cancel the kickback from the latch output to the sensitive sampling nodes, small metal capacitors cross-couple the output and the input. According to simulation, more than 80% reduction in the kickback noise can be achieved by proper sizing of these capacitors. These capacitors also reduce the loss of the S/H/summer due to the charge sharing between the sampling capacitors and the slicer/MUX input parasitic capacitor as they cancel the Miller capacitance associated to the gate-drain capacitance of the input differential pair.

### 3.3 Far-End Crosstalk Cancellation

Far-end crosstalk (FEXT) in transmission lines is a signal that is electromagnetically coupled from one line (aggressor line) to another line (victim line) and is received at the end of the victim line. It appears as an interfering signal at the victim channel receiver and degrades the horizontal and vertical eye-opening of the original victim signal. As shown in Chapter 2

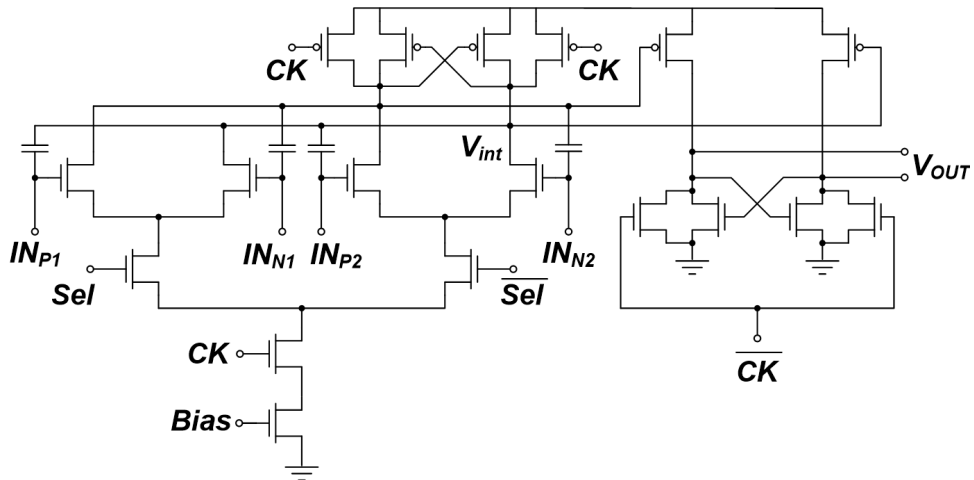


Figure 3.12: Combined analog MUX and latch with cross-coupled capacitors to reduce kickback.

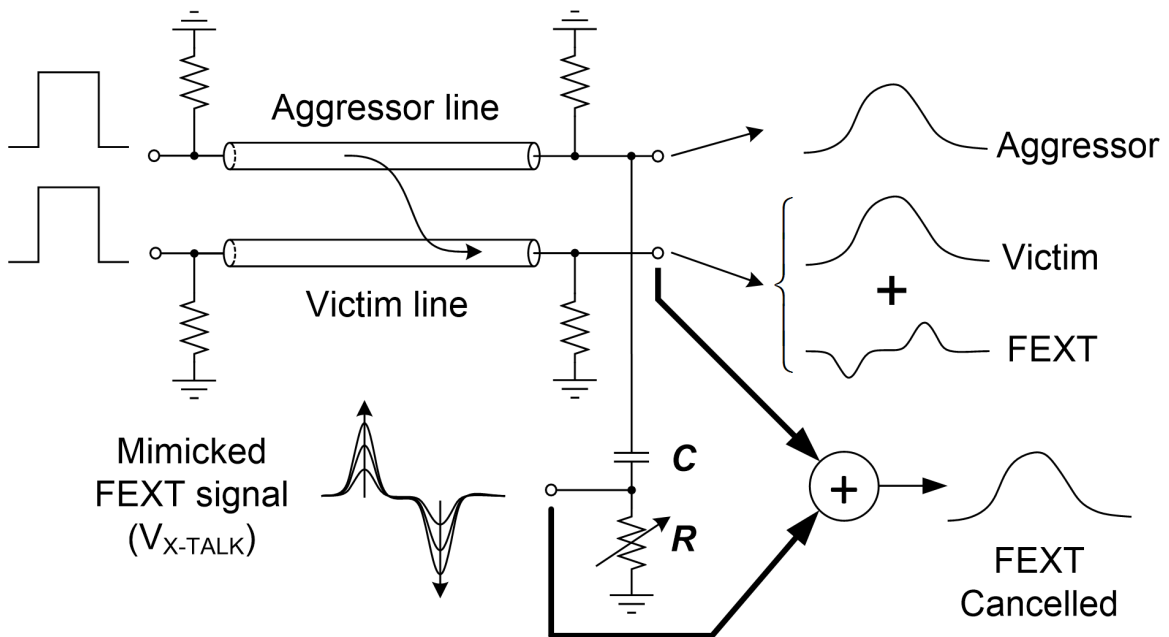


Figure 3.13: Crosstalk cancellation technique employing a high-pass filter as a differentiator to emulate FEXT.




Figure 3.14: Measured transfer characteristics of a 5" long, 32mil wide coupled trace with 40mil separation (a). Simulated FEXT noise due to a 15Gb/s pulse and the emulated FEXT employing the differentiator along with the residual FEXT noise (b).

the FEXT signal can be expressed as

$$V_{FEXT} = \frac{1}{2} \left( \frac{C_m}{C_t} - \frac{L_m}{L_S} \right) t_f \frac{dV(t - t_f)}{dt}. \quad (3.16)$$

The key observation in this equation is that FEXT noise appears at the front-end of the victim channel receiver as a voltage proportional to the derivative of the transmitted signal. This signal can have the same or opposite polarity as the aggressor signal if the link is capacitive or inductive, respectively.

Equation 3.16 shows that a differentiator with adjustable gain can be employed at the receiver to mimic the effect of FEXT. In the proposed design, the incoming aggressor signal is sent through an adjustable high-pass filter, as shown in Figure 3.13, to emulate FEXT for different levels of coupling. A simple  $RC$  circuit is employed as the high-pass filter. The transfer function of this filter is  $\frac{j\omega RC}{1+j\omega RC}$ . If the input frequency,  $\omega$ , is well below the cut-off frequency,  $(RC)^{-1}$ , the transfer function of the high-pass filter can be approximated by that of a differentiator,  $j\omega RC$ , with a gain equal to the time constant ( $RC$ ) of the filter.

The gain of the differentiator can be adjusted through a variable resistor, which is realized by an NMOS transistor operating in triode region. The output of the differentiator is finally




Figure 3.15: Simulated output of the S/H/summer before and after applying the aggressor as well as when the crosstalk cancellation is enabled.



subtracted from the received signal, which is comprised of the desired victim signal and the FEXT, as shown in Figure 3.13. Figure 3.14(a) shows the transfer characteristics of a coupled line on FR-4 with 40mil spacing and 32mil width. The simulated FEXT and the output of the high-pass filter for this channel, along with the residual FEXT noise are shown in Figure 3.14(b). The mimicked FEXT signal experiences a delay when it passes through the high-pass filter. As a result of this delay between the mimicked FEXT and the FEXT, the effect of the crosstalk noise is not completely eliminated. As the operating data rate is much smaller than the cut-off frequency of the filter, the amount of delay introduced is not considerable. Figure 3.15 illustrates the simulated output of the switched capacitor summer for the case where a low-loss and highly coupled channel is employed. The first eye diagram shows the output of the switched capacitor summer when the aggressor signal is disabled. Due to the low-loss of the channel, the eye diagram is quite open, however the aggressor signal causes eye closure. Employing the appropriate setting for the crosstalk cancellation filter, the eye opening is restored.

For channels with high coupling, the gain of the filter (RC) has to be increased which results in a large delay and hence degradation of the crosstalk cancellation. The dual-path sampling technique described in section 3.2.2.2 can be employed to resolve this problem as it can increase the filter gain while keeping the RC value constant.

In general, for dense parallel links, we may need to consider the effect of crosstalk from distant lines as well as the neighboring line. When differential signaling is employed, the crosstalk from one line to another diminishes approximately by a factor of  $D^{-3}$ , where  $D$  is the distance between the aggressor and the victim signal [77]. To investigate the functionality of the crosstalk cancellation technique in case of multiple channels, we employed three differential channels in parallel where each pair of the differential channels are 32mil wide and spaced 32mil, and 48mil apart from the adjacent channel pair, Figure 3.16. Figure 3.17(a), (b) show the simulated transfer characteristics for the channels with 10" and 20" length. It can be seen that the distant FEXT is considerably lower than the adjacent FEXT. The effectiveness of the crosstalk cancellation technique was determined by normalizing the root mean square (RMS) power of the residual FEXT noise to the FEXT signal RMS power.

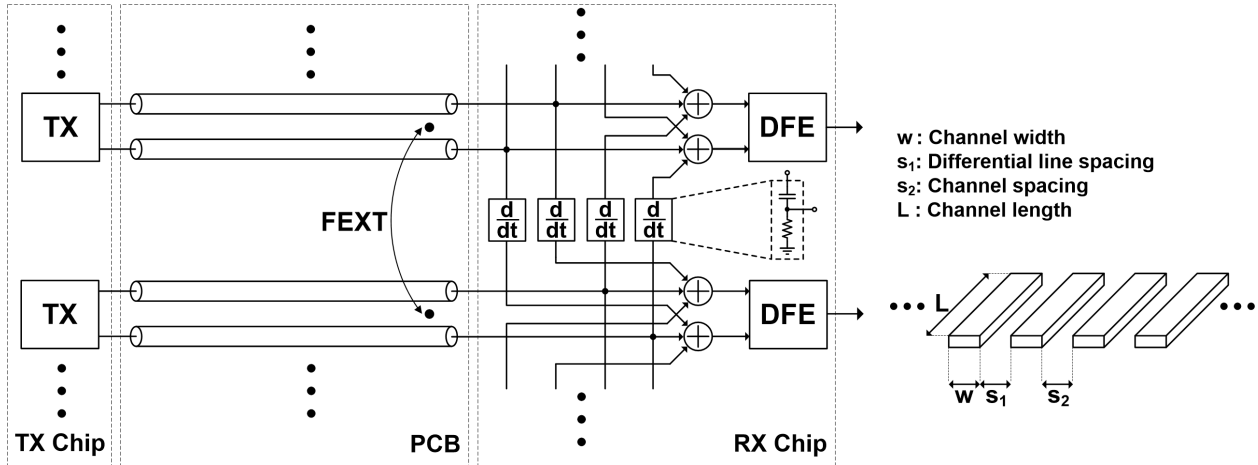


Figure 3.16: Crosstalk cancellation technique for multiple coupled channels.

With optimal selection of the differentiator gain more than 75% reduction in crosstalk noise can be achieved. Figure 3.17(c), (d) show the energy of the residual adjacent FEXT and the distant FEXT normalized to the FEXT energy for the 10" and 20" channel. The effect of the distant FEXT is less than the residual FEXT after the crosstalk cancellation is applied.

Figure 3.16 shows how the proposed crosstalk cancellation technique can be extended to multiple differential channels in parallel. As the effect of distant crosstalk is negligible, only the effect of crosstalk from the neighboring channel is canceled by taking the derivative of the signal received at the end of the adjacent line and subtracting it from the received signal in the victim line. As a result, there will be two differentiators connected at the end of each line. The loading effect of the differentiator is small compared to the parasitic capacitance of the bonding pad. In addition, as the aggressor line is 50-ohm terminated, the resulting insertion loss due to the crosstalk sensing is negligible. Figure 3.18 shows that the loading from the differentiator adds less than 0.2dB to the overall insertion loss.

The proposed crosstalk cancellation method does not involve resolving the aggressor signal to compensate for its effect on the victim signal. The effect of FEXT is removed by addition of the mimicked FEXT signal to the sampled input signal during the sum/hold phase. As addition and subtraction have minimal power overhead in this architecture, crosstalk cancellation adds very small power overhead, which is mainly due to the clock buffers. In addition, this scheme is not sensitive to the phase delay between the transmitted aggressor

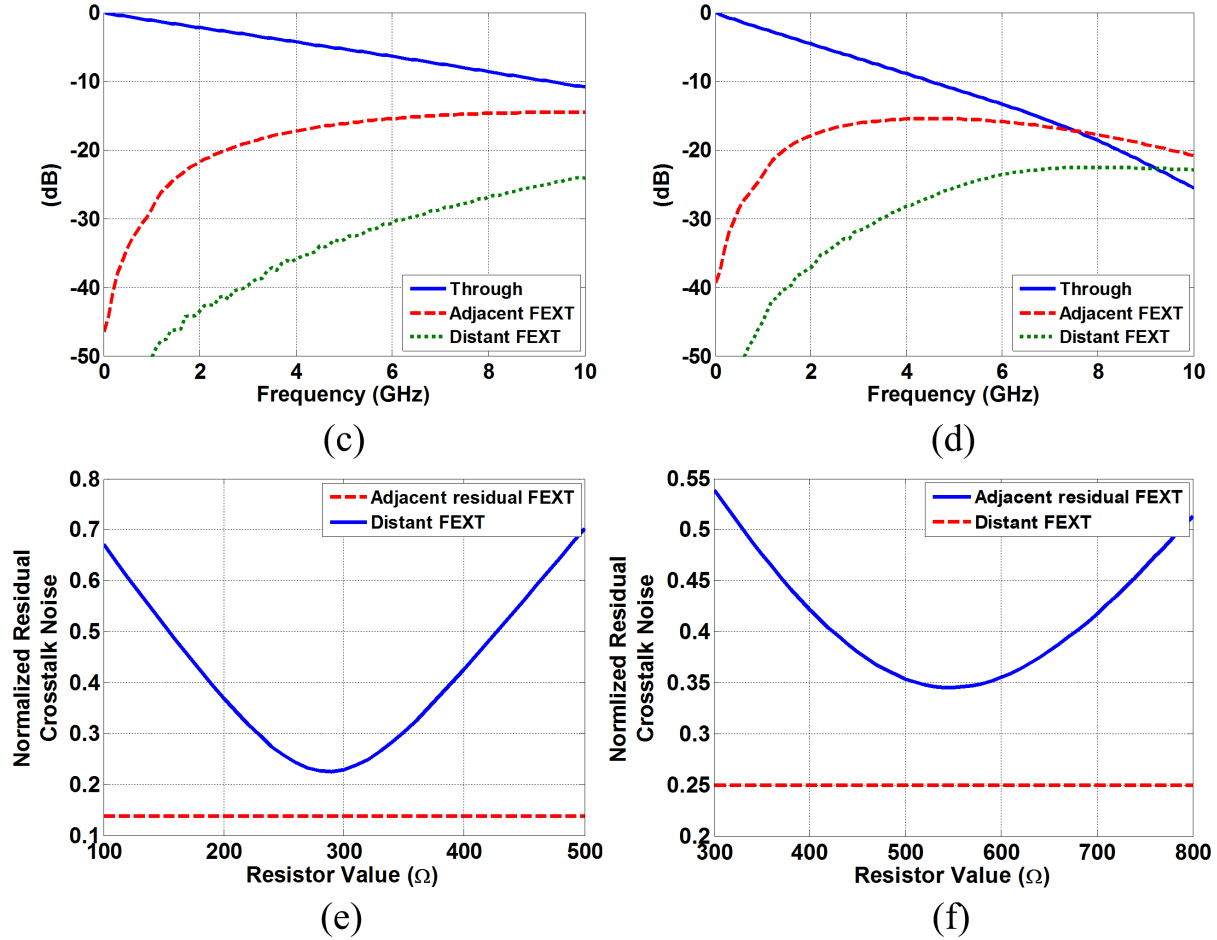




Figure 3.17: Simulated channel, adjacent FEXT, and distant FEXT response for width=32mil, spacing=48mil, and length=10", (d) for length=20". (e) Residual crosstalk from the adjacent aggressor after cancellation normalized to the FEXT energy along with the distant FEXT for 10" channel, and (f) 20" channel..

and the victim signals, as the effect of FEXT and the mimicked FEXT are sampled at the same time. As a result, synchronization between the aggressor and victim signals is not necessary, that is, the crosstalk cancellation operates whether the aggressor transition happens in the middle of the victim bit time (center of the eye) or when the victim signal transitions.

In order to adjust the crosstalk cancellation gain, RC, algorithms such as BER-based adaptation can be employed [92]. To apply this technique for a lossy coupled channel, the eye closure due to crosstalk should be decoupled from the eye closure due to ISI. As a result, during the adaptation process, the equalization tap coefficients should be first determined



(a)



(b)

Figure 3.18: Effect of loading due to the crosstalk cancellation circuitry on the overall channel insertion loss for a 10'' channel (a) and a 20'' channel (b).

by disabling the adjacent aggressor. When the tap coefficients for all channels are set by the adaptation algorithm, the adjacent aggressors are enabled and the same adaptation algorithm sets the crosstalk cancellation gain to achieve the minimum BER and hence maximum eye-opening.

## 3.4 Experimental Results

The prototype was fabricated in 45nm SOI CMOS technology. The die micrograph is shown in Figure 3.19 with all the blocks highlighted along with the receiver layout. The receiver, consisting of clock buffers, S/H/summer, slicer/MUX and DACs, occupies  $220\mu\text{m}\times 65\mu\text{m}$ . Figure 3.20 shows the measurement set-up. The prototype was tested with different channels. An Anritsu MP1800A signal quality analyzer was employed to generate a pseudo-random bit sequence (PRBS) signal which was transmitted over channels with different lengths as well as coupled channels to characterize the functionality of the receiver in channel loss compensation and crosstalk cancellation. The channel output was then carried to the receiver chip through high-speed SGS probes. The multiplexed output is monitored with a bit-error rate tester (BERT) and also through an oscilloscope.

The clock signal was also provided from off-chip and fed to the receiver through SGS probes. The CML input clock is converted into CMOS logic through a CML-to-CMOS converter, which is shown in Figure 3.21 [64]. An external differential half-rate clock signal is buffered by on-chip current-mode logic (CML) clock receivers and buffers. The duty cycle and noise performance of these CML-to-CMOS converters is of critical concern in the design of this block. The duty cycles of the CMOS clocks are corrected by using cross-coupled inverters, as shown in Figure 3.21.

Initially, the performance of the receiver was evaluated using input data with a low level of ISI. The PRBS7 data was transmitted to the receiver through low-loss cables and RF probes. Under this condition, the receiver operated error-free ( $\text{BER} < 10^{-12}$ ) up to 20Gb/s with an input sensitivity of 100mV, which reduces to 50mV at 15Gb/s. The input-referred offset was measured to be about 20mV at 15Gb/s. Note that offset compensation techniques are not incorporated into this design.

The equalization capability of the receiver was tested by transmitting data over 5", 10" and 18" FR-4 PCB traces. Figure 3.22 shows the characteristics of the channels employed to measure the functionality of the DFE, including the connecting SMA cables and connectors. With an  $800\text{mV}_{pp}$  differential PRBS7 data signal at 15Gb/s, the received eye is closed for

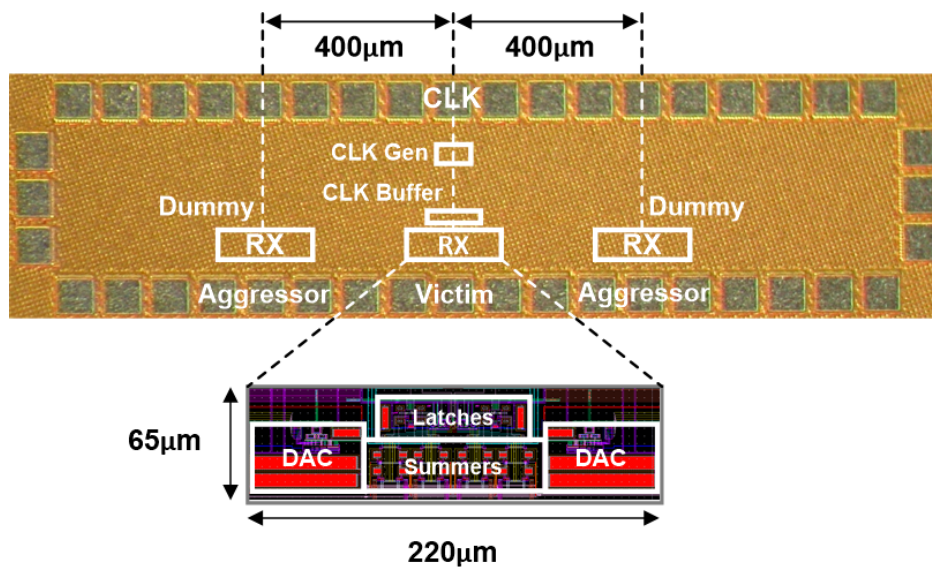


Figure 3.19: The die micrograph of the receiver with major blocks highlighted.

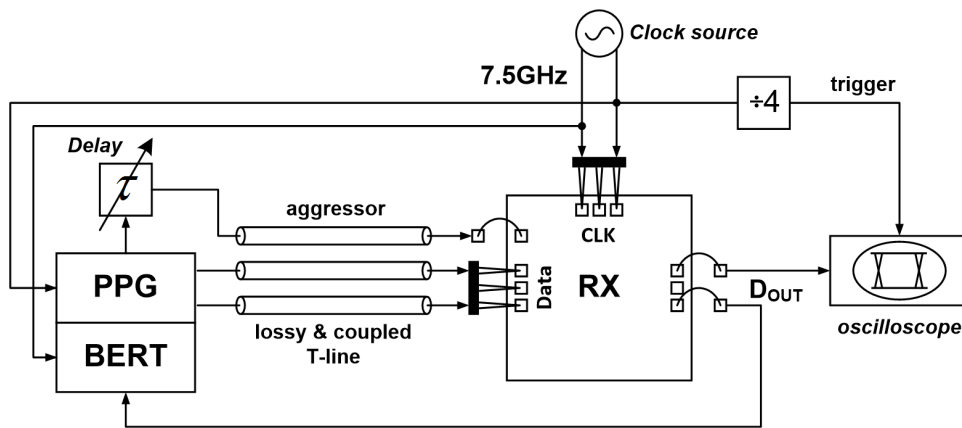


Figure 3.20: Receiver DFE and crosstalk cancellation test set-up.




Figure 3.21: CMOS clock generation through CML-to-CMOS conversion followed by duty cycle correction.




Figure 3.22: Channel transfer characteristics for 5", 10" and 18" PCB traces.

all these channels, shown as insets in Figure 3.23. The 5" channel exhibits a loss of 14.5dB at 7.5GHz. Employing the 2-tap DFE, while consuming 7.5mW from a 1.2V supply, a 24% horizontal eye opening ( $BER=10^{-12}$ ) is achieved. The DFE was also tested with 10" and 18" channels. Due to the increased loss of the 10" channel at 7.5GHz, the DFE failed to equalize pseudo-random data at 15Gb/s. The data rate was accordingly reduced to test the limits of the DFE, and 13Gb/s data was transmitted over the 10" channel with 17dB of loss at 6.5GHz. Under these conditions, the DFE achieved 35% horizontal eye opening while dissipating 6.1mW. In order for the DFE to equalize the 18" channel, the data rate was further reduced to 11Gb/s. This channel had about 21dB loss at 5.5GHz. The DFE achieved a 26.5% eye opening while consuming 5.5mW.

Table 3.1 summarizes the DFE performance. Compared to prior art, the proposed design offers the best figure of merit (FOM) [84]. It also provides one of the most compact DFE designs among recently published works. Figure 3.24 illustrates the power breakdown of the receiver. As can be observed, the major portion of the receiver power consumption is due to the clock buffers. In this design the clock distribution network was not optimized for minimum power consumption. An optimized clocking system can greatly improve the power efficiency of the receiver. Simulation shows that an optimized design can reduce clocking power consumption by a factor of two and achieve less than 0.35mW/Gbps power efficiency. The fact that the power consumption of the proposed receiver is mainly dominated by digital blocks means that the overall power consumption can be greatly reduced by technology scaling.

The crosstalk cancellation scheme was evaluated by transmitting random, uncorrelated victim and aggressor data over a 5" long, 32mil wide coupled trace with 40mil separation on an FR-4 PCB. To generate uncorrelated data sequences, differential outputs of the pulse pattern generator (PPG) were delayed with respect to one another through the delay element,  $\tau$ , as shown in Figure 3.20. The functionality of the crosstalk cancellation technique was tested at different data rates. In the first experiment, 8Gb/s PRBS data was transmitted over the coupled channel, while the aggressor was kept quiet. Figure 3.25(a) shows that without crosstalk noise, the DFE generates 59% of eye opening. Applying the aggressor






Figure 3.23: The PRBS7 eye diagram at the receiver input and the bathtub curve after equalization for, (a) 11Gb/s data over 18" trace, (b) 13Gb/s data over 10" trace, (c) 15Gbs data over 5" trace.

Table 3.1: DFE PERFORMANCE SUMMARY

Process	45nm SOI		
Architecture	2-tap speculative DFE		
Supply (v)	1.2		
Power (mW)	7.5	6.1	5.5
Data Rate (Gb/s)	15	13	11
Channel loss @ Nyquist (dB)	14.5	17	21
Horizontal eye opening (BER<10 <sup>-12</sup> )	24%	35%	26.5%
FOM (pJ/bit/dB)	0.034	0.028	0.024
Sensitivity	±50mV <sub>ppd</sub>		
Area (mm <sup>2</sup> )	0.014		




Figure 3.24: Receiver power breakdown.

Table 3.2: CROSSTALK CANCELLATION PERFORMANCE SUMMARY

Data Rate (Gb/s)	8	10	11	12.5
Horizontal eye opening (<math>10^{-12}</math>) no aggressor	59%	46%	31%	40%
Horizontal eye opening (<math>10^{-12}</math>) w/o cancellation	40%	11.3%	0%	0%
Horizontal eye opening (<math>10^{-12}</math>) with cancellation	53%	20%	16.5%	15%

signal degrades the eye opening to less than 40%, which is restored to above 53% when the crosstalk cancellation is activated. Next, 10Gb/s data was employed to test the crosstalk cancellation functionality. Figure 3.25(b) shows the bathtub curve of the receiver with the DFE activated to compensate for the channel loss. Applying the aggressor signal causes BER degradation to higher than  $10^{-10}$ . Crosstalk cancellation restores the eye opening.

The same experiment was repeated for 11Gb/s and 12.5Gb/s data rates. The channel has more than 11.5dB loss and about -15.6dB coupling at 5.5GHz. The DFE and crosstalk canceler provide more than 16.5% of horizontal eye opening while the input eye is completely closed due to channel loss and crosstalk noise. The values of FEXT and channel loss at 6.25GHz are -15dB and 12.5dB, respectively. The input eye is closed when no aggressor is applied. The DFE compensates for the loss and generates a 40% open eye at  $BER < 10^{-12}$ . Applying the aggressor closes the eye and degrades the BER to higher than  $10^{-5}$  at the center of the eye, as shown in Figure 3.25(d). The crosstalk canceler achieves more than 15% horizontal eye opening. Table 3.2 summarizes the performance of the crosstalk cancellation technique.

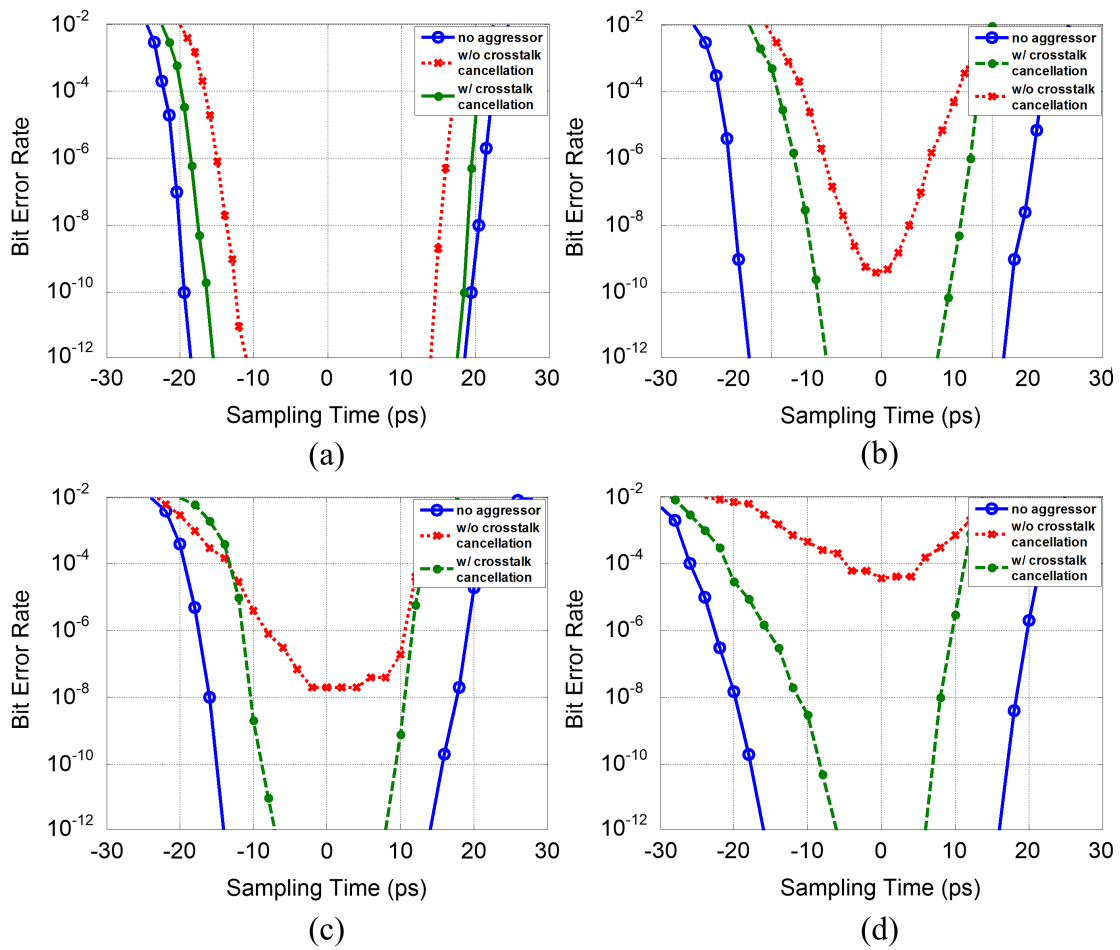


Figure 3.25: The receiver bathtub curve without and with crosstalk noise, and after crosstalk cancellation for, (a) 8Gb/s, (b) 10Gb/s, (c) 11Gb/s, and (d) 12.5Gb/s victim and aggressor data.

## 3.5 Transmitter Design

As mentioned earlier, one of the main challenges of modern wireline communication systems is the severe frequency-dependent loss in the channel. In the previous section we discussed the advantages of employing receiver equalization to combat the channel loss. However as the data rate keeps increasing, only utilizing receiver equalization becomes less effective. Transmitter equalization is another technique, which in conjunction with receiver equalization can enable high data rates over high attenuation channels [69, 70, 72]. In addition, low-power equalization at the transmitter can simplify the design of the receiver and help with the overall power consumption [94].

In this section we will introduce a novel transmitter equalization technique, which employs continuous-time linear equalization using compact programmable passive devices. This design offers a small form factor and adjustable equalization at the transmitter to accommodate channels with different characteristics [95].

### 3.5.1 Transmitter Architecture

As mentioned in Chapter 2, finite impulse response (FIR) and analog filters are the two most commonly used techniques for transmitter pre-emphasis. FIR-based transmitters require additional hardware for generating delayed versions of the data. At high data rates the extra hardware consumes a considerable amount of power. Additionally, FIR pre-emphasis reduces the output signal swing due to the low frequency attenuation of the transmitted data. As a result, the power penalty of employing this technique at high data rates can be very large. On the other hand, analog filtering does not require any additional active component and solely relies on passive devices and offers high power efficiency. Although passive devices occupy a large area, as the frequency of operation increases, the area drops. In addition, since high quality factor is not necessary, a 3D layout can help to reduce the area. The basic idea behind analog filtering technique is to boost the high frequency content of the transmitted signal to compensate the high frequency attenuation of the channel.

Shunt, series, shunt-series, and shunt-double-series techniques have been previously used




Figure 3.26: Shunt and double-series bandwidth enhancement technique, which requires three inductors. A T-coil can also perform the same functionality.




Figure 3.27: Segmented T-coil layout 3D and top views. The five top metals are employed.

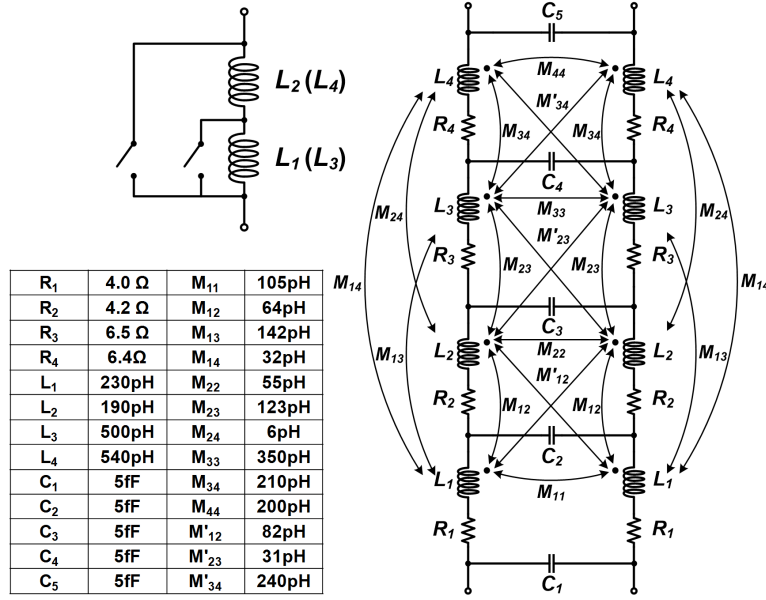


Figure 3.28: Segmented T-coil lumped model generated using IE3D electromagnetic simulator.

in order to enhance the bandwidth of amplifiers and compensate capacitive loads [38]. In many signaling applications, the bandwidth of the links is limited due to the lossy-capacitive (RC) behavior of the channel. In this work, modified peaking techniques are employed to compensate the effective capacitance of the trace and equalize the channel loss. The maximum bandwidth enhancement can be achieved through employing the shunt and double-series technique, which involves using six inductors in a differential implementation as shown in Figure 3.26. In this technique the line capacitance ( $C_L$ ) forms a series  $LC$  with  $L_3$ . As a result, the effect of  $C_L$  is canceled at the resonant frequency  $\frac{1}{2\pi\sqrt{L_3C_L}}$ . Around this frequency, the transfer function from the differential pair transistor current ( $I_{in}$ ) to the output can be simplified to Equation 3.17.  $L_1$  and  $C_P$  form a pair of complex conjugate poles that introduce a considerable amount of peaking in the transfer function.

$$\left| \frac{V_{out}}{I_{in}} \right|^2 = \frac{1}{(1 - L_1 C_P \omega^2) C_L \omega}. \quad (3.17)$$

As the quality factor of the inductor is not high, a fairly broad-band peaking is obtained. Another zero at  $f_z = \frac{R}{2\pi L_2}$  can be added to the transfer function through  $L_2$ , as in the shunt peaking technique. The main drawback of this technique is the excessive silicon area

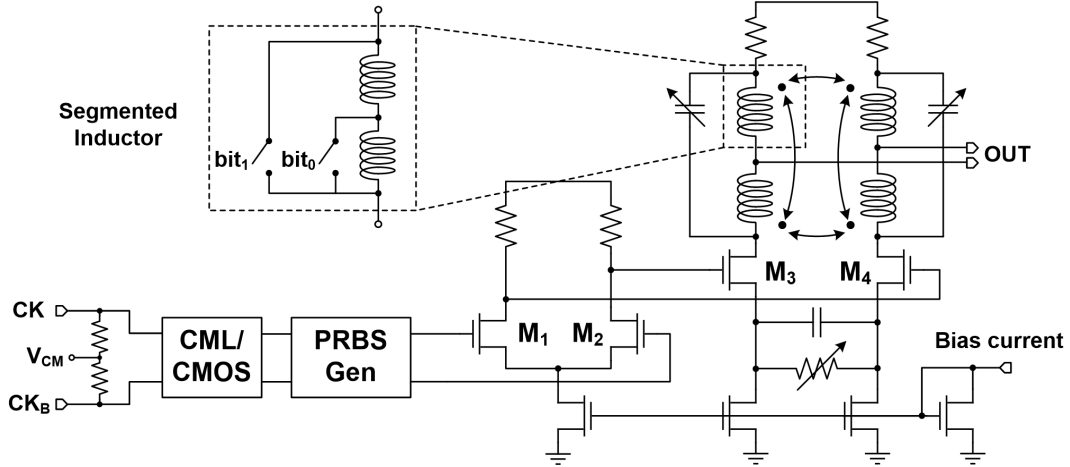


Figure 3.29: Transistor-level schematic of the transmitter employing segmented T-coils.

required. In order to alleviate the area penalty, the inductors are replaced by two T-coils that offer the same performance as shown in Figure 3.26. To further improve the area efficiency of this technique, we have utilized the fact that the two branches of the differential transmitter carry currents in opposite directions. By mutually coupling the coils in two branches, we obtain the same inductance in a smaller area

$$L_{eff} = L + M. \quad (3.18)$$

Moreover, the opposing currents in the two branches significantly reduce the magnetic field outside the T-coil area, allowing for dense arrays of transmitters. Figure 3.27 shows the layout of the two coupled T-coils. They are interwound using three thick metal layers. Although the utilized technology provides only two thick metal layers, a third is realized by stacking three thin metal layers. In order to be able to equalize different channels with different amounts of loss, the inductors in each T-coil are segmented using MOS switches with two bits of resolution as shown in Figure 3.28. This allows for adjusting the inductance to introduce different levels of peaking.

In addition to inductive peaking, a frequency-dependent source degeneration technique is employed by adding a variable parallel RC to the source of the transmitter to improve high frequency peaking, as shown in Figure 3.29. This introduces a zero at frequency  $\frac{1}{2\pi RC}$ .




Figure 3.30: Simulation results showing the programmable high frequency peaking achieved by RC source degeneration.

The variable resistor is implemented using an NMOS transistor in triode. According to simulation, about 10dB of peaking is achieved by using this technique, as illustrated in Figure 3.30.

The transmitter is comprised of two stages as shown in Figure 3.29. The first stage is the pre-amplifier that drives the output stage. The output stage utilizes T-coil and zero peaking. The segmented T-coils were simulated in an electromagnetic simulator (IE3D) and modeled, as shown in Figure 3.28. Figure 3.31(a), (b) shows the simulation results for the transmitter connected to two different channels. The channel response, along with different levels of pre-emphasis, is shown for 5" and 10" FR4 PCB traces. An improvement in bandwidth from 800MHz to 8GHz is achieved with optimal pre-emphasis.

### 3.5.2 Experimental Results

The prototype was fabricated in a 65nm CMOS process. The chip micrograph is shown in Figure 3.32. The transmitter consumes 10mW from a 1.2V supply and provides  $250mV_{pp}$  output swing. The transmitter was tested with an on-chip PRBS-7 generator. A high speed, 32bit shift register was also integrated to enable the application of arbitrary patterns to the transmitter for testing and debugging purposes. The output of the on-chip PRBS-7






Figure 3.31: Simulated transfer characteristics of the transmitter and channel for different levels of pre-emphasis, (a) 5" FR4 channel, (b) 10" FR4 channel.




Figure 3.32: Transmitter die micrograph along with the core layout.




Figure 3.33: Transmitter measurement setup for characterizing the performance over different channels




Figure 3.34: 20Gb/s on-chip PRBS-7 generator.

generator was sent through a lossy cable as well as 5” and 10” FR4 PCB traces at maximum rate of 20Gb/s. Figure 3.35 shows the transfer characteristics of these channels. A 2-tap FIR was simulated in MATLAB with the same channels to compare the performance of the analog filter and FIR-based pre-emphasis techniques in channel loss compensation.

Figure 3.36 illustrates the adjustability of the transmitter in equalizing 10Gb/s data over a 5” channel. By changing the equalization settings under over and optimal equalization is achieved. As shown in Figure 3.37, by employing the proposed pre-emphasis technique, at 15 Gb/s over the lossy cable (7dB at 7.5GHz), a completely open eye with less than 12ps peak-to-peak jitter was measured. Figure 3.38 illustrates the measured and simulated eye-diagrams at the output of 5” and 10” FR4 channels while a 15Gb/s data was sent by the proposed transmitter and a 2-tap FIR respectively. The 5” and 10” FR4 channels introduce about 15dB and 20dB loss at 7.5GHz. A 20Gb/s data was also sent through the lossy cable with 10dB loss at 10GHz. The output eye-diagram has less than 18ps peak-to-peak jitter as shown in Figure 3.38(c).

It is important to note that circuit non-idealities are not included in the simulations of the 2-tap FIR and the actual FIR circuit performance is expected to be worse. The high frequency peaking introduced in this technique may cause extra far-end crosstalk (FEXT) in highly parallel communication links. The resulting crosstalk can be canceled employing low-power crosstalk cancellation techniques [83]. Table 3.3 compares the performance of this design with FIR-based transmitters. To achieve the same output swing, the 2-tap FIR requires twice as much power as the proposed design. In many FIR-based pre-emphasis


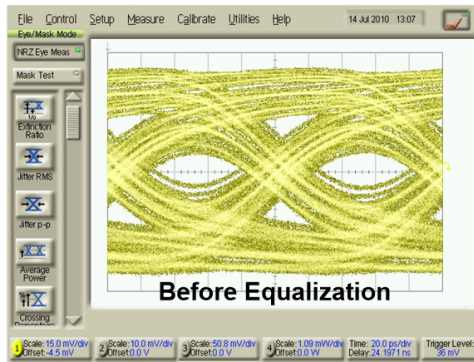
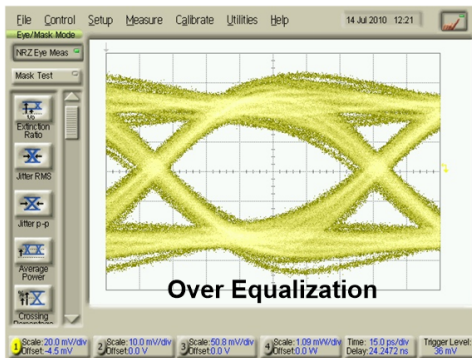


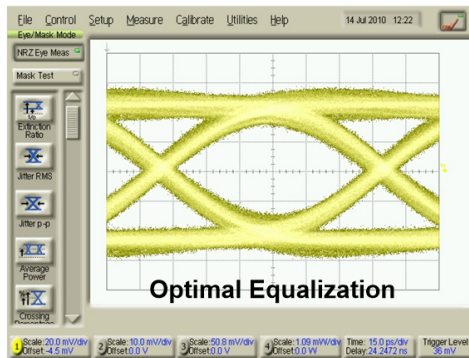
Figure 3.35: Channels' transfer characteristic.



(a)



(b)



(c)

Figure 3.36: The transmitted 10Gb/s PRBS7 data over 5" FR4 channel with about 10dB loss at Nyquist, before equalization (a), over equalized (b) and optimally equalized (c).




Figure 3.37: Transmitter output at 15Gb/s over lossy channel with 7dB loss.

Table 3.3: TRANSMITTER PERFORMANCE SUMMARY

Rate	Loss	Power	Output Swing	process	Supply	Area	Architecture
Rate	20dB	10mW	250mV <sub>pp</sub>	65nm	1.2V	0.01mm <sup>2</sup>	Continuous-time Equalization

designs [85, 96, 97] shunt peaking inductors are required to meet the bandwidth constraint of the CML delay elements. These inductors impose a large area penalty as shown in table 3.3. The comparison table and the measured/simulation results in Figure 3.38 show superior performance for the analog filter-based equalization technique both in terms of channel loss compensation and power and area efficiencies.

### 3.6 Summary

The power consumption and area of the receiver is a critical design aspect of parallel electrical interconnects, where many IOs are placed in a single chip. To address these challenges we have proposed a low-power DFE receiver with crosstalk cancellation capability. The equalization is implemented through a novel switched capacitor technique that allows many taps of equalization with small power overhead. A 4-tap realization is implemented in a 45nm SOI CMOS technology as a proof of concept. It has been shown that the proposed SC DFE can be extended to implement a large number of taps. The receiver is suitable for channels with considerable amount of ISI. The simple, low-power DFE can significantly enhance

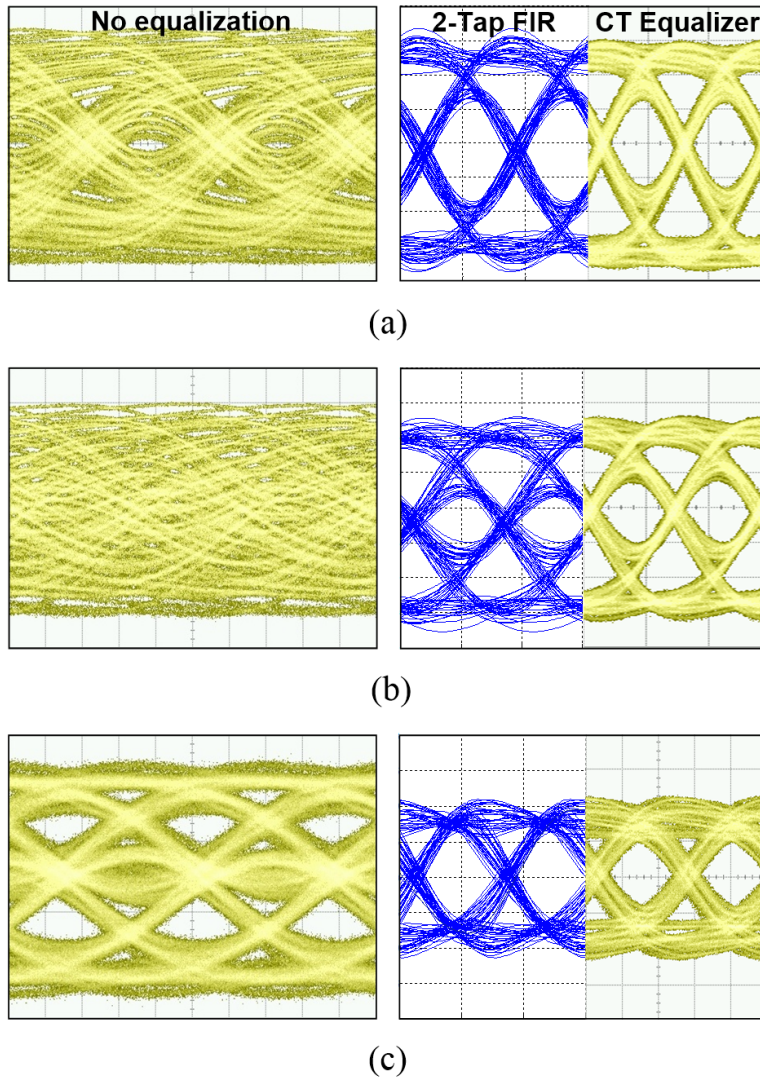


Figure 3.38: Output of the channel before and after equalization with the 2-tap FIR equalizer and continuous-time equalizer for 15Gb/s data over 5" channel (a), 15Gb/s data over 10" channel (b), and 20Gb/s data over lossy channel (c).

the data rate over lossy channels. In this design, high power efficiency (0.5mW/Gb/s) is achieved by using SC summation technique, analog multiplexers, and half-rate clocking. Unlike current-mode summers, the SC summer does not require a high bias current. It can also benefit from technology scaling due to switch performance improvement and reduced power consumption of the clock distribution network. The major portion of the receiver power consumption is due to the clock buffers. As a result, the overall power consumption can be greatly reduced by technology scaling. A novel crosstalk cancellation is incorporated in the receiver, which removes more than 75% of crosstalk noise. As addition and subtraction have minimal power overhead in the SC summer architecture, the extra hardware required for crosstalk cancellation results in only 5% ( $33\mu\text{W}/\text{Gbps}$ ) extra power dissipation. By multiplexing analog instead of digital signals, the number of digital latches is minimized. The half-rate clocking allows the use of CMOS clock buffers instead of CML buffers and relieves the speed requirements of the front-end circuits. Experimental results validate the feasibility of the DFE receiver for ultra-low-power, high-data rate and highly parallel I/O links.

A 20Gb/s transmitter utilizing efficient continuous-time equalization is presented. The proposed pre-emphasis technique supports data transmission over PCB channels with loss levels in excess of 20dB at  $\text{BER} < 10^{-12}$ . The proposed architecture consumes significantly less power compared to an FIR transmitter that has similar performance, while occupying very small silicon area. The receiver and transmitter together can enable 15Gb/s data communication over channels with more than 30dB loss with a power efficiency less than 1mW/Gb/s. The low-power consumption and the crosstalk capability offered by this transceiver makes it well-suited for densely parallel chip-to-chip communication.

# Chapter 4

## Overview of High-Speed Optical Links

We showed in the previous two chapters that equalization can significantly help in increasing the overall communication bandwidth over copper channels with severe frequency-dependent loss. As the communication distance grows while bandwidth requirement keeps scaling, equalized channels exceed the power envelope and become inadequate in delivering the required data in a power efficient manner. The power consumption and area of the optical transmitter and receiver electronics can be the limiting factors for the number of IOs possible on-chip. An example of such situation is the board-to-board communication in data centers and high-performance computers. A promising solution to this IO bandwidth requirement is the use of optical signaling.

The primary motivation for an I/O architecture modification as radical as optical signaling is the magnitude of potential bandwidth offered with an optical channel. In conventional optical data transmission, data is transmitted by modulating the optical intensity or amplitude of the high-frequency optical carrier signal. In order to achieve high fidelity over the most common optical channels, optical fiber, high-speed optical communication systems typically use infrared light from source lasers with wavelengths ranging from 850-1550nm, or equivalently frequencies ranging from 200-350THz, which provide a high potential data bandwidth. Moreover, because the loss of typical optical channels at short distances only changes by fractions of dB over wide wavelength ranges (tens of nanometers) [100], there is the potential for data transmission of several Tb/s without the requirement of channel equalization through wavelength division multiplexing (WDM). This simplifies design of op-




Figure 4.1: Optical signal transmission over fiber.

tical links in a manner similar to electrical links with high bandwidth channels. However, optical links require additional circuits that interface to the optical sources and detectors. As a result, in order to achieve the potential link performance advantages, emphasis is placed on using efficient optical devices and low power and area interface circuits. Using optics is particularly promising because of the emergence of silicon photonics technology which enables integration of high-performance optical devices such as photodiodes, modulators and waveguides on the same platform [101]. As a result, high data rate and power efficiency can be achieved through efficient and cost-effective hybrid integration of advanced CMOS and silicon photonics technologies using techniques such as flip-chip bonding [102] and copper pillars [103].

We start this chapter by briefly reviewing the basics of optical links and optical devices used in them. We will introduce optical channels, transmitters and receivers. In the first part of two sections of this chapter, we briefly discuss different options for optical channels and optical light generation and modulation. However, the main emphasis of this chapter is to go over the design of low-power optical receivers. Since the receiver circuitry directly interfaces with the photodetectors, understanding the operation and characteristics of these devices is essential for an optimum design. In the second part of this chapter, we focus on the design of the receiver electronics. We examine the prior art in front-end design for optical communication. Investigating the existing designs provides a motivation for the next chapter, which describes the proposed double-sampling RC front-end.






Figure 4.2: An optical fiber cross-section with the core and cladding having refractive index of  $n_1$  and  $n_2$ , respectively to allow for total internal reflection.

## 4.1 Optical Channels

The data transmission scheme in optical links is very similar to electrical links, Figure 4.1. A clock signal at the transmitter is used to define uniform time periods for sending the data signals successively one after another. A modulated optical signal is generated at the transmitter by directly modulating a vertical cavity surface emitting laser diode or by externally modulating a continuous wave (cw) laser. At the receiver side a photodiode converts the optical signal to an electrical signal proportional to its input optical power. The receiver circuitry following the photodetector is responsible for resolving the data from the incoming electrical signal. Similar to electrical receivers, a clock signal, synchronized with the data, is used for sampling and data decision.

The two optical channels relevant for short distance chip-to-chip communication applications are optical fibers [104–106] and on-board polymer optical waveguides [107–109]. These optical channels offer potential performance advantages over electrical channels in terms of loss, cross-talk, and both physical interconnect and information density.

Optical fiber-based systems provide alignment and routing flexibility for chip-to-chip interconnect applications. As shown in Figure 4.2, an optical fiber confines light between a higher index core and a lower index cladding via total internal reflection. Fibers are classified based on their ability to support multiple or single modes.

Multi-mode fibers with large core diameters (typically 50 or  $62.5\mu\text{m}$ ) allow several propagating modes, and therefore provide good coupling characteristics. These fibers are used in short and medium distance applications such as parallel computing systems and campus-scale interconnection. They suffer from a relatively large loss ( $\sim 3\text{dB/km}$  for 850nm light),

which might not be an issue for aforementioned applications. The major performance limitation of multi-mode fibers is modal dispersion caused by the different light modes propagating at different velocities. This effect can cause major ISI, particularly at high data rates.

Single-mode fibers with smaller core diameters (typically 8-10 $\mu\text{m}$ ) only allow one propagating mode (with two orthogonal polarizations), and thus require careful alignment in order to avoid coupling loss. These fibers are optimized for long distance applications such as links between Internet routers spaced up to and exceeding 100km. Fiber loss typically dominates the link budgets of such systems, and thus they often use source lasers with wavelengths near 1550nm which match the loss minima ( $\sim 0.2\text{dB/km}$ ) of conventional single-mode fibers. While modal dispersion is absent from single-mode fibers, chromatic (CD) and polarization-mode dispersion (PMD) exist. However, these dispersion components are generally negligible for distances less than 10km, and are not issues for short distance inter-chip communication applications.

Another means of chip-to-chip optical communication is to employ an on-board polymer optical waveguide. Similar to optical fibers, polymer waveguides can either support multiple or single optical modes. Usually, to facilitate coupling and reducing assembly costs, multi-mode waveguides are favorable. Figure 4.3 illustrates the cross-section of a polymer waveguide. The waveguide core is surrounded by a cladding layer with smaller refractive index to enable total internal reflection. Due to their large core area, they provide negligible coupling loss. In addition, the modal dispersion is fairly small as they are intended for short range board-level interconnection.

In summary, both fiber-based and polymer waveguide systems are applicable for chip-to-chip optical interconnects. For both optical channels, loss is the primary advantage over electrical channels. This is highlighted by comparing the highest optical channel loss, present in multi-mode fiber systems ( $\sim 3\text{dB/km}$ ), to typical electrical backplane channels at distances approaching only one meter ( $> 20\text{dB}$  at 5GHz). In addition, since pulse dispersion is small in optical channels for distances appropriate for chip-to-chip applications ( $< 10\text{m}$ ), no channel equalization is required. This provides another advantage over electrical interconnects with complex equalization required to compensate for the channel frequency dependent loss.

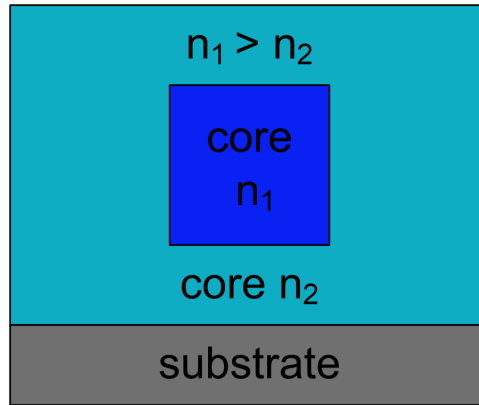


Figure 4.3: Cross section of a polymer waveguide.

## 4.2 Optical Receivers

Optical receivers generally determine the overall optical link performance, as their sensitivity sets the maximum data rate and amount of tolerable channel loss. Typical optical receivers use a photodiode to sense the high-speed optical power and produce an input current. This photocurrent is then converted to a voltage and amplified sufficiently for data resolution. In order to achieve increasing data rates, sensitive high-bandwidth photodiodes and receiver circuits are necessary.

### 4.2.1 Photodiodes

The most commonly used devices for optical to electrical conversion are p-i-n diodes. In this type of diode, an electrical field in a semiconductor material drives the electrons and holes generated by the incident photons in the intrinsic region to the n and p terminals, respectively. The result is a current proportional to the number of photons absorbed per second, which is called photocurrent. In the p-i-n diode, a reverse-bias across the diode ensures a strong field in the intrinsic region and a very small current in the absence of light, which is referred to as dark current. Figure 4.4 shows a simple electrical model for a photodiode. The optically generated current  $I_{opt}$  is proportional to the input optical power  $P_{opt}$  with the proportionality factor  $\rho$ , which is also known as the photodiode responsivity.

The capacitance of the photodiode is usually the dominant load for a receiver, impacting




Figure 4.4: Electrical model of a photodiode.

the sensitivity, electrical power consumption, and the bandwidth of the front-end. Photodiodes can be hybrid-integrated to CMOS chips via a number of techniques, such as wire bonding, flip-chip bonding, and copper pillars. The advantage of hybrid integration is that the material and design of the photodiode is independent of the transistor technology. Flip-chip bonding and nano pillar are manufacturing technologies that enable the integration of large device arrays while introducing low parasitic. Wire bonding has larger parasitic inductance and capacitance, reducing performance at high bit-rates.

### 4.2.2 Receiver Front-End

The task of an optical receiver front-end is to convert the current from the photodiode into a voltage and resolve the data. As shown in Figure 4.5, a simple resistor can perform the current to voltage conversion. In case of a high extinction ratio, a zero generates zero voltage and a “one” generates  $R \times I_1$ . A voltage amplifier then amplifies this signal for the next stage slicer for data resolution. It should be noted that in this architecture, sensitivity has a direct relation with the resistor size. On the other hand, the resistor and the photodiode capacitance form an RC time constant, which can limit the maximum data rate of the receiver. In order to avoid ISI, the bit-time has to be greater than the input time constant by almost a factor of four. This results in a strong trade-off between the sensitivity and the data rate as they both depend on  $R$ . This trade-off can be resolved by employing transimpedance amplifiers (TIA). Integrating front-end offers another solution to this trade-off. In the next two sections we will discuss these techniques in detail and introduce their advantages and disadvantages.




Figure 4.5: Simple resistive receiver front-end performing current to voltage conversion.

#### 4.2.2.1 Transimpedance Amplifiers

The strong trade-off between the bandwidth and SNR of a front-end with a simple resistor makes it impractical for many applications. The effective input resistance of the front-end can be reduced significantly by adding an active component to the design, resulting in a transimpedance architecture. A transimpedance amplifier (TIA) is an analog front-end with reduced input impedance and a relatively high current-to-voltage gain. The addition of active components, like transistors, will add to the noise. However, with a careful design, very high SNRs are possible at the output of an optimized transimpedance amplifier. Detailed analysis of TIAs are covered in numerous publications [110–112]. In this section, we briefly discuss the performance and trade-offs of a number of different TIAs. For TIAs, like any other receiver, the most important specs are bandwidth, sensitivity, power consumption, and area.

**Common-Gate and Regulated Cascode TIA.** In order to achieve a low input impedance and at the same time a high gain, common-gate (CG) topology can be employed. The common-gate TIA acts as a current buffer with a gain close to unity. As shown in Figure 4.6(a) it creates isolation between the diode capacitance  $C_p$  and the gain resistor  $R_D$  and therefore provides a wide bandwidth. The effective input impedance is inversely proportional to the input transistor transconductance ( $g_m^{-1}$ ), while the transimpedance is equal to  $R_D$ . The main disadvantage of this architecture is that the noise from the input transistor and the bias current directly adds to the input photocurrent and degrades SNR. In addition, to

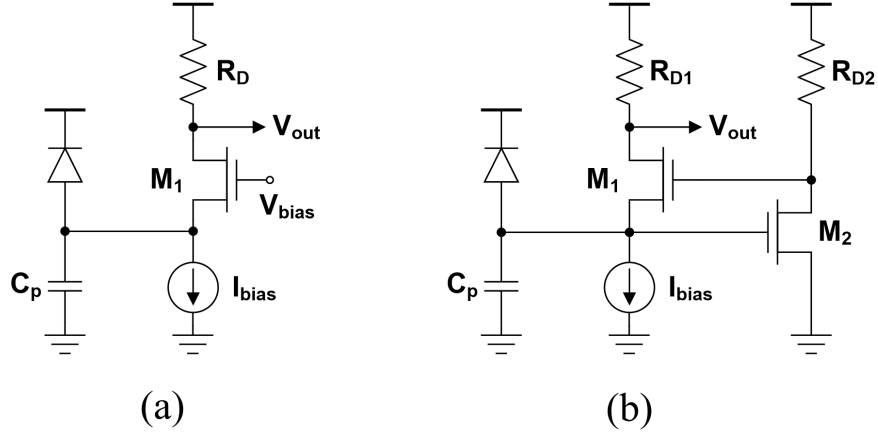


Figure 4.6: Schematic of a common-gate TIA (a) and a regulated cascode TIA (b).

reduce input resistance, the transconductance of the input transistor has to be maximized. This necessitates a large transistor which adds large capacitance at the input, as the  $C_{gs}$  of the transistor is directly loading the input. Moreover, the noise power of this transistor scales directly with its  $g_m$ , which results in a strong trade-off between bandwidth and SNR.

A regulated cascode (RGC) configuration addresses these issues [113, 114]. Due to the negative feedback at the input, the input resistance effectively reduces, while the transconductance remains almost the same as the CG TIA. The reduction factor in the input resistance is equal to the feedback loop gain. As a result, higher bandwidths are feasible. Figure 4.6(b) shows the schematic diagram of the RGC circuit. With a simple small-signal analysis, the input resistance of the RGC circuit can be approximated by

$$R_{in} = \frac{1}{g_{m1}(1 + g_{m2}R_2)}. \quad (4.1)$$

The local feedback employed in this configuration creates a second pole which can cause instability. As a result, careful design is necessary to guarantee proper operation of this TIA.

**Common-Source and Shunt-Shunt Feedback TIA.** An alternative design with better noise performance is a shunt-shunt feedback TIA [110]. A TIA with resistive feedback, followed by a chain of post amplifiers is the most common type of receiver design, which is shown in Figure 4.7(a). A detailed analysis of the performance of this configuration

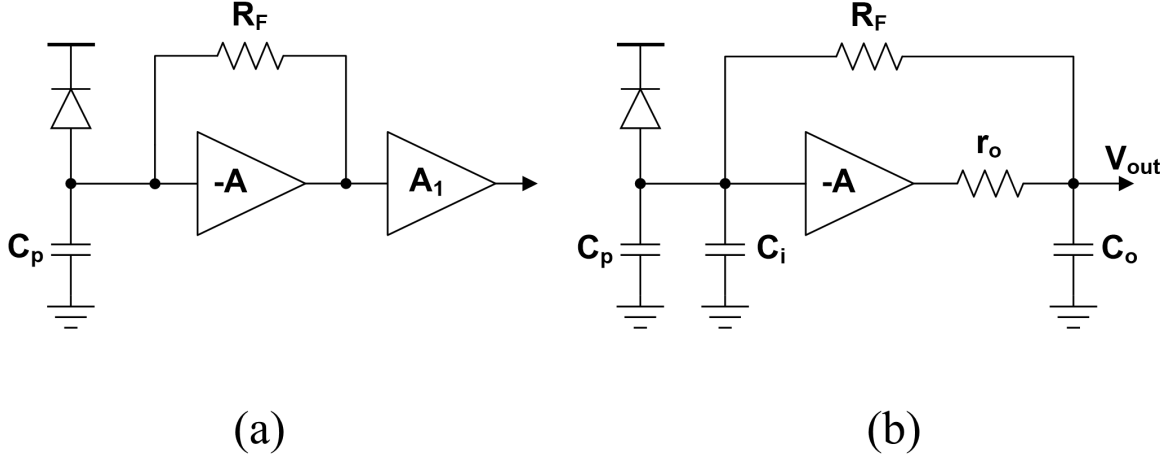


Figure 4.7: Typical shunt-shunt feedback TIA (a). Simplified small signal model of the TIA (b).

is presented next, using a simple model shown in Figure 4.7(b). If we ignore the output resistance of the amplifier, the performance of this system is dominated by a single pole at  $(A + 1)/(R_F C_T)$ , where  $C_T$  is the total capacitance at the input node. In addition, the DC transimpedance is equal to  $R_F/(1 + A^{-1})$ , while the input impedance at DC is reduced to about  $R_F(1 + A)$  due to the Miller effect. In most designs the overall TIA bandwidth is limited by the pole at the input node. However, real amplifiers have a finite open-loop pole that limits their gain bandwidth product. Accounting for this pole by including  $r_o$  and  $C_o$  leads directly to a more complete transfer function. Here we see that for small (low-power) amplifiers,  $r_o$  directly reduces gain and lowers the input pole, shown as [155]

$$Z_t = \left(\frac{AR_F - r_o}{A + 1}\right) \left(1 + s \frac{r_o C_o + (R_F + r_o)C_T}{A + 1} + s^2 \frac{R_F C_T r_o C_o}{A + 1}\right)^{-1}. \quad (4.2)$$

Assuming that the two poles are well apart, we can approximate the bandwidth as

$$BW \approx \frac{A + 1}{2\pi(r_o C_o + (R_F + r_o)C_T)}. \quad (4.3)$$

The sensitivity of the TIA depends on the total input-referred noise of the TIA and the noise of the diode itself. The TIA noise is mostly due to the thermal noise of the feedback resistor  $R_F$  and the input-referred noise of the amplifier. A detailed derivation of the output noise due to the amplifier and feedback resistor can be found in [155]. Assuming that the

amplifier noise is due to a transistor with noise contribution of  $4\gamma kTg_m$ , the total output noise can be written as

$$\overline{V_n^2} = \overline{V_{na}^2} + \overline{V_{nf}^2}, \quad (4.4)$$

where  $\overline{V_{na}^2}$  and  $\overline{V_{nf}^2}$  are the amplifier and  $R_F$  noise contributions, respectively, and can be expressed as

$$\overline{V_{nf}^2} = \frac{A}{A+1} \frac{\gamma kT}{C_o} \frac{r_o C_o + (A+1)R_F C_T}{r_o C_o + (R_F + r_o)C_T}, \quad (4.5)$$

$$\overline{V_{na}^2} = \frac{1}{A+1} \frac{kT}{C_o} \frac{A^2 R_F C_o + (A+1)r_o C_T}{r_o C_o + (R_F + r_o)C_T}. \quad (4.6)$$

Considering a fanout of 1 ( $C_o = C_T$ ), for large  $A$ , the ratio of the amplifier and  $R_F$  noise is equal to  $1 + C_p/C_i$ . For a low-power TIA with  $C_p/C_i \gg 1$  the noise is mainly dominated by the amplifier noise. On the other hand, if  $C_p/C_i \ll 1$ , the two noise sources will have almost the same contribution.

We will consider two configuration of the shunt-shunt feedback TIA in this section. In the first design shown in Figure 4.8, with enough open-loop gain the noise contributions from resistance  $R_D$  and transistor  $M_2$  can be very small and the noise of the TIA is mainly dominated by the noise contribution of transistor  $M_1$  and  $R_F$ . The main disadvantage of this design is its limited voltage headroom, which makes it difficult to implement in advanced technology nodes with low supply voltage. The minimum possible supply for this configuration is mandated by the output signal swing and the gate-source voltage ( $V_{gs}$ ) of transistors  $M_1$  and  $M_2$ , as well as the voltage drop on  $R_D$ .

Figure 4.9(a) shows another shunt-shunt feedback TIA employing an inverter as the gain stage. There are several advantages associated with this TIA. First, the use of both NMOS and PMOS transistors enhances the effective transconductance ( $G_m = g_{mn} + g_{mp}$ ) for a certain power, however this comes at the expense of added input capacitance. This configuration can be readily ported to other technology nodes due to the use of inverters as






Figure 4.8: Schematic of a common source TIA.

the building block. In addition, it can operate with small supply voltages and offer a large output swing. For an inverter gain stage we have  $A = G_m/G_{ds}$ , where  $G_{ds} = g_{dsn} + g_{dsp}$  and  $r_o = (G_m + G_{ds})^{-1}$ . Therefore, the bandwidth of this configuration can be written as

$$BW = \frac{1 + \frac{G_m}{G_{ds}}}{\frac{C_o}{G_m + G_{ds}} + (R_F + \frac{1}{G_m + G_{ds}})C_T}. \quad (4.7)$$

For small inverters, the bandwidth is mainly limited by the photodiode capacitance. Therefore increasing the inverter size helps improve the bandwidth due to the increase in  $G_m$ . For large inverters, the bandwidth is dominated by the inverter parasitic capacitance. Increasing the inverter size in this situation increases both  $G_m$  and  $C_i$ , however as  $C_i$  scales proportionally to the inverter size while  $G_m$  scales with a square root relation, the net result is the degradation of the bandwidth. The low-frequency transimpedance,  $R_t$ , of this configuration is equal to  $R_F - G_m^{-1}$ , therefore, it is desired to increase  $G_m$  to maximize  $R_t$ . This creates a strong trade-off between the bandwidth and transimpedance. Figure 4.10 (a) shows the maximum BW versus transimpedance for different photodiode capacitance values in 28nm CMOS technology. The associated power is also illustrated in Figure 4.10(b). In order to achieve a high data rate we have to sacrifice transimpedance and hence sensitivity, as shown in Figure 4.11. As one can see, it is difficult to achieve better than  $1K\Omega$  of transimpedance from one stage. As a result, a larger number of stages have been used by designers [155] to achieve higher transimpedance as shown in Figure 4.9(b).

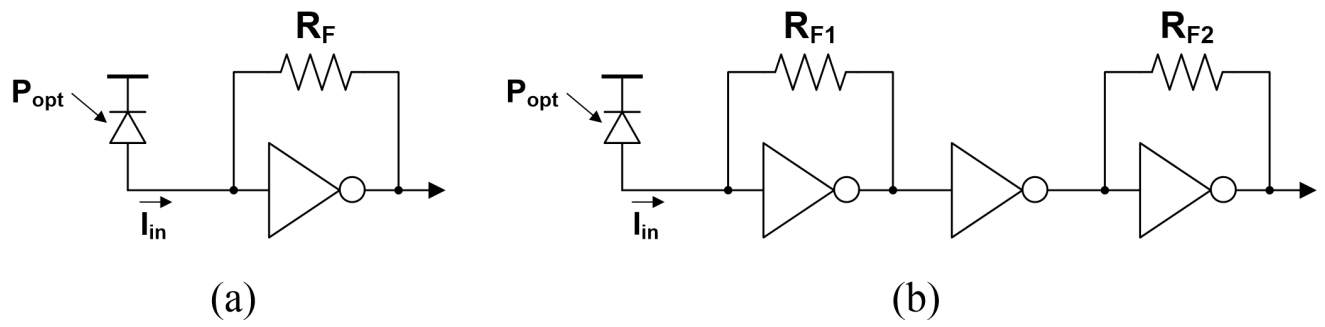




Figure 4.9: Inverter-based TIA with one (a) and three (b) stages.



(a)



(b)

Figure 4.10: TIA data rate (a) and power consumption (b) as a function of transimpedance.




Figure 4.11: Sensitivity degradation as a result of data rate scaling.

#### 4.2.2.2 Integrating Front-Ends

Integrating front-ends have been used to reduce the power consumption and area of the front-end by eliminating the need for TIAs. In this type of front-end, the input impedance of the receiver is designed to be purely capacitive within the frequency range of the input data. The optically generated current from the photodiode is then integrated onto the capacitor seen at the input node,  $C_{tot}$ , which is the sum of the diode, bonding and front-end circuitry capacitances. If the average input current during a bit zero is  $I_0$  and during a bit one is  $I_1$ , the voltage swing at the input node will be  $\Delta V_0 = I_0 T_b / C_{tot}$  for a zero and  $\Delta V_1 = I_1 T_b / C_{tot}$  for a one. To implement this technique, two photodiodes are required, which receive complementary data, as shown in Figure 4.12, to avoid saturation. A logic one charges the inverter input and a logic zero discharges it. The sensitivity directly depends on the size of the capacitor  $C_{tot}$ . If the input optical power is high enough to charge and discharge the input node close to  $V_{dd}$  and  $Gnd$  in less than a bit-time, a simple inverter can recover a full swing voltage and resolve the received data [115]. The required input modulation optical power for a full voltage swing is  $P_{opt} = C_{tot} V_{dd} \rho T_b$  where  $\rho$  is the photodiode responsivity. The minimum optical power is proportional to  $C_{tot}$ , requiring very small photodiode capacitances as well as the small voltage buffer that follows it.

With the existing photodiodes, the receiver-less front-end needs a relatively high input




Figure 4.12: Integrating optical front-end employing balanced photodiode and an inverter.

optical power to generate a full swing voltage. The required voltage swing at the input node can be reduced by replacing the inverter stage with a sense-amplifier [116]. The sense-amplifier-based front-ends either need two photodiodes and a pair of complementary optical beams or a precise reference current, in order to resolve the data. For each bit-time, the receiver has two phases, integration phase and reset phase. The input optical power is used only during the integration phase, the data is evaluated and then both integrating nodes are reset to the initial voltage. Figure 4.13 shows a sense-amplifier-based front-end with a precise reference current. As seen in this figure, a DC current equal to the average of a zero and a one current,  $I_0$  and  $I_1$ , is added to the integrating node to balance the integrating voltage around zero and eliminate the need for a reference voltage. The reset operation eliminates the saturation problem that the inverter-based front-end faces. This technique requires return-to-zero (RZ) data transmission scheme. As a result, only half of the bandwidth is employed for data communication. A synchronous internal clock is used to set the correct integration, evaluation and reset phases. It should be also noted that a clock with the same rate as the data is required.

The sense-amplifier based front-end improves the sensitivity compared to the receiver-less topology, and has lower power consumption compared to the TIA. However, it requires a reset phase (RZ data stream), which reduces the effective data rate of the system. Double-sampled integrating front-end is a technique proposed in [117] that solves some of the problems associated with the integrate-and-reset scheme.

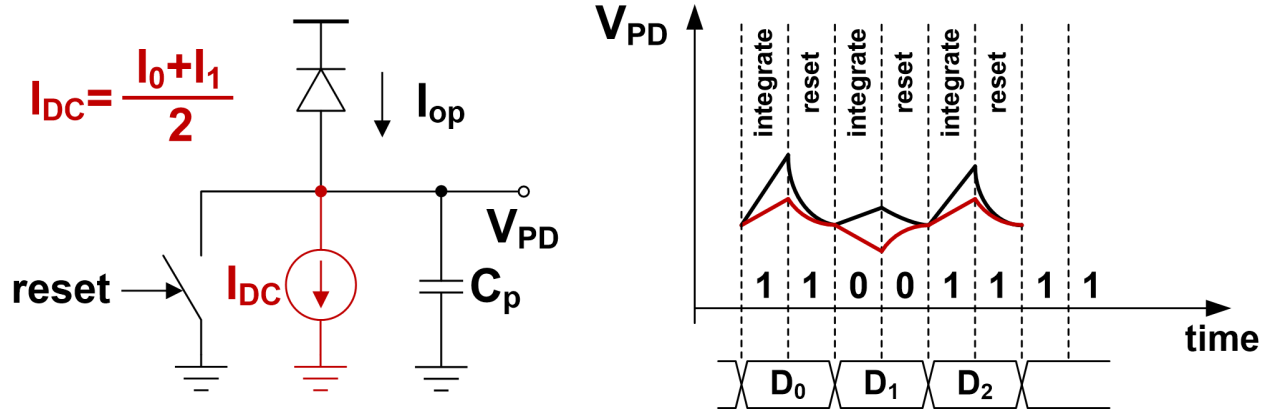


Figure 4.13: Integrate and reset front-end.

As mentioned earlier, in an integrating front-end, the photodiode current is integrated over the parasitic capacitance, which results in distinct voltage changes at the integrating node due to a received zero or one. The basic idea behind the double-sampling technique is to sample and compare the voltage at the integrating node to distinguish between a zero and a one. This technique can be applied to the integrate-and-reset front-end to resolve the incoming data. To better exploit the bandwidth, the reset process can be performed after receiving a certain number of bits. For instance, if after every nine bits we reset the integrating node, a factor of 1.8 increase in the data rate can be achieved. By limiting the number of integration bits before the reset phase, the saturation problem is also resolved. However, this technique requires data encoding to insert null bits during the reset phase and precise synchronization between the receiver and the transmitter to locate the null bits in the incoming data stream.

In an integrating front-end, the voltage of the input node at the end of each bit-time,  $V[n]$ , is always the sum of the voltage change due to the incoming signal and the voltage of the input node just before that bit;

$$V[n] = V[n - 1] + \Delta V[n] = V[n - 1] + \frac{I_{in} T_b}{C_T}, \quad (4.8)$$

where  $I_{in}$  is the input current to the receiver,  $I_1$  in case of a received one and  $I_0$  in case of a received zero. Similar to the sense-amplifier-based front-end, a DC current equal to the

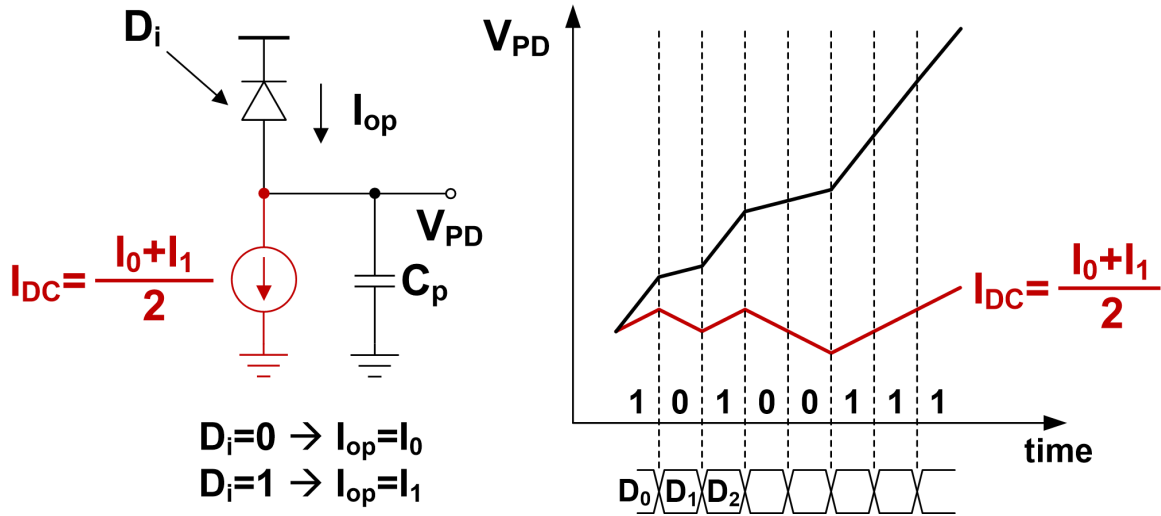


Figure 4.14: Double-sampling front-end with a DC current to provide bipolar voltage difference for a one and a zero.

average of  $I_1$  and  $I_0$  can be added at the integrating node to make the effective  $I_{in}$  balanced around zero. Therefore, if we have the voltage samples at the end of each bit-time,  $V[n]$  and  $V[n - 1]$ , we have enough information about the input signal at time  $t_n$  to determine whether it was a one or a zero. Figure 4.14 illustrates the operation of the double-sampling front-end.

As seen in this figure, the integrating node voltage monotonically increases upon receiving both zeros and ones, which causes saturation within a short time. To partly resolve this problem, a DC current is injected to the integrating node, similar to the integrate-and-reset configuration, which also provides a bipolar voltage change,  $\Delta V[n]$ . The bipolar voltage change at the input allows us to decide the input value by comparing the two adjacent samples of the input voltage. If the new sample is higher ( $\Delta V[n] = V[n] - V[n - 1] > 0$ ), the input signal is one, otherwise ( $\Delta V[n] = V[n] - V[n - 1] < 0$ ), it is zero, as shown in Figure 4.14.

Figure 4.15 illustrates the top level block diagram of this receiver. The input signal from the photo detector is single-ended, with a positive current. The DC current,  $I_{DC}$ , needs to be adjusted by a feedback loop looking at the DC value of the voltage of the input node. The feedback loop not only adjusts the DC current but also sets the average voltage of the

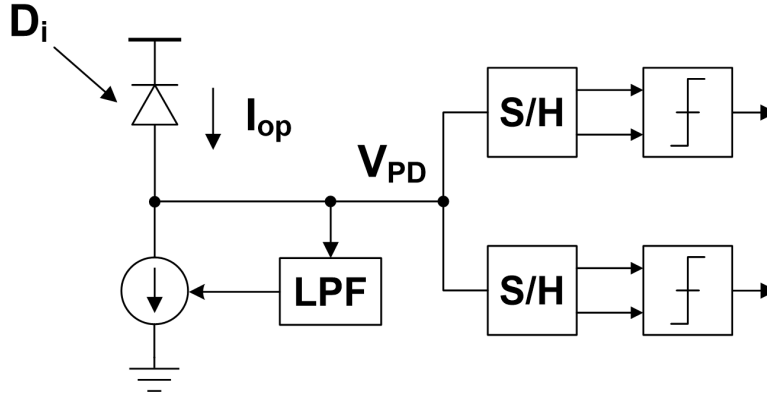


Figure 4.15: Double-sampling integrating front-end.




Figure 4.16: Input demultiplexing receiver using multiple sampler clock phases.

input node.

Figure 4.16 illustrates a possible implementation of this front-end. Two non-overlapping clock phases perform the double sampling, which results in a bit rate twice the on-chip clock frequency. This design has two samplers and each sampled value is used twice by the two comparators. The two slicers operate with complementary clock phases and are triggered once every clock cycle, after every other sampling period. An important aspect of this technique is the capability of performing demultiplexing by a factor of  $N$  immediately at the front-end by employing  $N$  phases of clock as shown in Figure 4.16. The samplers and slicers will operate at a rate which is  $N$  times lower than the data rate at the expense of more hardware requirement.

Even though the low-pass filter and the DC current injected at the integrating node provide bipolar voltage, they do not help prevent the saturation of the integrating node voltage

after receiving long sequences of ones or zeros. In [117] and [89] it has been suggested to employ encoding schemes to avoid many consecutive ones or zeros, such as 8B/10B encoding.

## 4.3 Optical Transmitters

With the recent advances in optical technology, optical devices that can handle tens of Gb/s data rates are available. These high performance optical devices facilitate very high data rates in optical interconnects, if high bandwidth transceiver circuitries, optimized for the characteristics of optical devices are designed. Optical signal transmission involves electrical to optical conversion. This is typically performed using a laser diode. The laser output then has to be modulated to embed the digital data intended to be communicated. The light from the laser diode can be directly modulated using a vertical cavity surface emitting laser while being generated or externally modulated through a Mach-Zehnder or ring resonator after being generated by a cw laser diode. Even though in this work we do not provide a solution for optical transmission, a brief introduction to optical transmitters is provided in the next section to give the reader a complete picture of a chip-to-chip optical communication link.

### 4.3.1 Vertical Cavity Surface Emitting Laser

A VCSEL is a semiconductor laser diode which emits light perpendicular to its top surface. The most common VCSELs are GaAs-based operating at 850nm [118, 119], with 1310nm GaInNAs-based VCSELs in recent production [120], and research-level devices near 1550nm [121]. While VCSELs appear to be the ideal source due to their ability to both generate and modulate light, they suffer from serious inherent bandwidth limitations and reliability concerns.

The output optical power of a VCSEL is a linear function of the forward current. It should be noted that there exists a threshold current after which this relation is held. Once the VCSEL begins lasing, the optical output power is related to the input current by the slope efficiency  $\eta$  (typically 0.3-0.5mW/mA) and a high extinction ratio between a logic one






Figure 4.17: Typical current-mode VCSEL driver

signal and a logic zero signal can be achieved by placing the zero current value near the threshold. Low level current for a zero provides a high extinction ratio, however, due to the low bias current of the laser diode, the speed is sacrificed. As a result the transition from zero to one occurs more slowly than the one to zero transition. As a result, equalization techniques are usually employed in VCSEL drivers to avoid ISI [89, 122–125].

Current-mode drivers are typically used to modulate VCSELs due to the direct relationship between drive current and optical output power [89]. A typical VCSEL output driver is shown in Figure 4.17, employing a differential stage to steer current into the VCSEL. Usually an additional static current source is used to bias the VCSEL sufficiently above the threshold current in order to ensure adequate bandwidth. In order to accommodate the large forward bias voltage of the VCSEL diode, the output stage usually uses a separate higher voltage supply [89, 122], which incurs an increase in the overall power consumption of the transmitter.

### 4.3.2 Mach-Zehnder Modulator

These devices generally use a pn diode structure which is positioned with the junction in or around an optical waveguide such that the depletion region, whose width changes with applied reverse bias, interacts with the light propagating along the waveguide. A change




Figure 4.18: A Mach-Zehnder modulator comprising of two arms which introduce different phase shifts to the optical signal in order to perform amplitude modulation.

in the phase of the light exiting the waveguide then occurs with changing depletion width due to the resultant change in effective refractive index. A beam splitter divides the light from a continuous wave laser into two paths. One of these paths includes a pn junction to introduce phase modulation. The beams are then recombined. If enough phase shift is introduced through the phase modulating path, the two beams can combine destructively and cancel each other, which results in zero net output optical power. On the other hand, if the two beams experience the same phase shift, they will add constructively and the same optical power that was inputted to the device appears at the output. A simple schematic of a Mach-Zehnder Modulator (MZM) is shown in Figure 4.18. For better efficiency, the two arms of the MZM could be driven differentially by the data signal.

Due to the poor electro-optical properties of silicon, the change in the refractive index is quite small in silicon-based devices. As a result, to achieve a 180 degree phase shift using silicon phase modulators, a fairly long waveguide is required. The p and n regions are usually driven through a transmission line in the phase-shifting arm, which makes the electrical driver power hungry. In addition to excessive power consumption associated with the MZM driver, they suffer from large area. The main advantages of these structures is their insensitivity to temperature, providing a high extinction ratio (the ratio between optical power for a one and a zero) and supporting a wide range of optical wavelengths.

### 4.3.3 Ring Resonator

Optical ring resonators operate based on the optical signal coupling between waveguides. When a beam of light passes through a waveguide as shown in Figure 4.19, part of the light will be coupled into the optical ring resonator. The reason for this phenomenon is the wave property of the light. In other words, if the ring and the waveguide are close enough, the light in the waveguide will be transmitted into the ring.

There are three aspects affecting the optical coupling: the distance, the coupling length and the refractive indices between the waveguide and the optical ring resonator. In order to optimize the coupling, the usual practice is to narrow the distance between the ring resonator and the waveguide. The closer the distance, the more easily the optical coupling occurs. In addition, the coupling length affects the coupling as well. The coupling length represents the effective curve length of the ring resonator for the coupling phenomenon to occur with the waveguide. Furthermore, the refractive indexes of the waveguide material, the ring resonator material and the medium material in between the waveguide and the ring resonator also affect the optical coupling. The medium material is usually the important one been studied since it has a large effect on the transmission of the light wave. The refractive index of the medium can be either large or small according to various applications and purposes.

The key property of the ring resonator waveguide is that only lights with certain wavelength are coupled into the ring. The necessary condition for coupling to occur is that the optical path length (OPD) of the ring is a multiple of the wavelength  $\lambda$ . OPD can be expressed as

$$OPD = 2\pi r n_{eff}, \quad (4.9)$$

where  $r$  is the radius of the ring resonator and  $n_{eff}$  is the effective refractive index of the waveguide material. It should be noted that  $n_{eff}$  must be larger than the surrounding material to satisfy total internal reflection requirement.

In order to exploit the ring resonator to perform amplitude modulation a mechanism




Figure 4.19: A ring resonator with pn structure for performing resonant wavelength shift enabling optical amplitude modulation.

is needed to change the OPD for a certain incoming wavelength. This can be done using a carrier-depletion waveguide. As shown in Figure 4.19, a pn junction around the ring can change the effective refractive index [126, 127]. The larger the difference between the refractive index of the waveguide and its surrounding material is, the sharper the bends can be implemented. Therefore ring resonators can potentially offer very compact and low power electro-optical modulators.

The main challenge of using ring resonators is the sensitivity of the resonant wavelength to the temperature and the fabrication process. To address this problem, thermal resonant wavelength tuning is usually employed which adds directly to the overall power dissipation [128–130].

## 4.4 Summary

In this chapter we went over the basics of optical links as a solution to the bandwidth problem of electrical communication links. The negligible frequency dependent loss of optical channels provides the potential for optical link designs to fully utilize increased data rates provided through CMOS technology scaling without excessive equalization complexity. Optics also allow very high information density through wavelength division multiplexing (WDM). Hybrid integration of optical devices with electronics has been demonstrated to achieve high performance. These approaches pave the way for massively parallel optical communications. In order for optical interconnects to become viable alternatives to established electrical links,

they must be low cost and have competitive energy and area efficiency metrics. Dense arrays of optical detectors require very low-power, sensitive, and compact optical receiver circuits. Particularly, the power efficiency of the optical receiver is of great importance. Existing designs for the input receiver, such as TIA, require large power consumption to achieve high bandwidth and low noise, and can occupy large area due to bandwidth enhancement inductors. Moreover, these analog circuits require extensive engineering efforts to migrate and scale to future technologies. In this chapter, we discussed different TIA configurations and introduced their advantages and disadvantages.

While TIAs have a relatively high sensitivity, the power consumption and stability issues of TIAs make them less than optimal for densely parallel applications. We showed that integrating front-ends can reduce the power consumption by avoiding an analog amplifier that runs at the bit-rate. The integrate-and-reset configuration is a configuration that provides low-power consumption but requires a full rate clock and a non-return to zero (NRZ) data modulation. We also introduced the double-sampling technique, which helps avoid the reset phase in the integrate-and-reset configuration and enables return to zero (RZ) modulation. The main problem with this technique is the limited number of consecutive ones or zeros that can be resolved before entering saturation. As a result, data encoding such as 8B/10B scheme is necessary to avoid long sequences of identical bits. In the next chapter we introduce a double-sampling RC front-end which benefits from the advantages of the integrating front-end while resolving the headroom problem.

# Chapter 5

## Low-Power Optical Receiver Design

As mentioned in previous chapter, integrated circuit scaling has enabled a huge growth in processing capability, which necessitates a corresponding increase in inter-chip communication bandwidth. This trend is expected to continue, requiring both an increase in the per-pin data rate and the I/O number. Unfortunately the bandwidth of the electrical channels and the number of pins per chip are not following the same trend. As data rates scale to meet increasing bandwidth requirements, the shortcomings of copper channels are becoming more severe. While I/O circuit performance benefits from technology scaling, the bandwidth of electrical channels does not scale with the same trend. Especially as the data rate increases, they exhibit excessive frequency-dependent loss, which results in significant inter-symbol interference (ISI). In order to continue scaling data rates, equalization techniques can be employed to compensate for the ISI. However, the power and area overhead associated with equalization make it difficult to achieve target bandwidth with a realistic power budget. As a result, rather than being technology limited, current high-speed I/O link designs are becoming channel and power limited.

A promising solution to the I/O bandwidth problem is the use of optical inter-chip communication links. The negligible frequency dependent loss of optical channels provides the potential for optical link designs to fully utilize increased data rates provided through CMOS technology scaling without excessive equalization complexity. Optics also allow very high information density through wavelength division multiplexing (WDM). Hybrid integration of optical devices with electronics has been demonstrated to achieve high perfor-

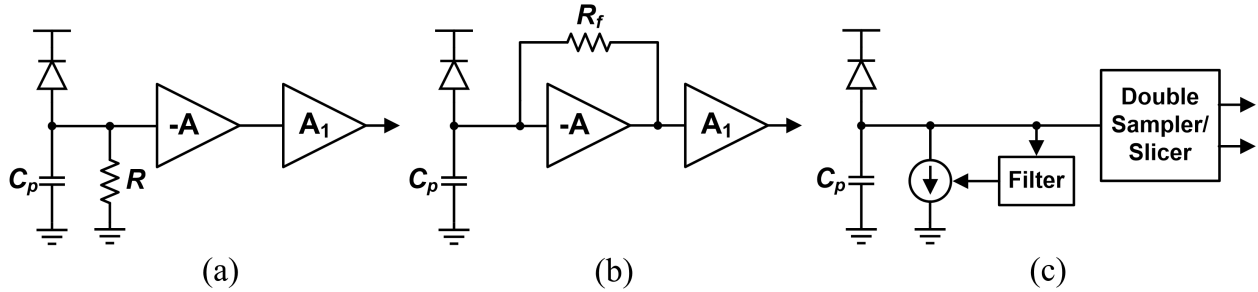


Figure 5.1: Different optical receiver architectures. (a) simple resistive front-end, (b) transimpedance front-end with limiting amplifiers, (c) integrating double-sampling receiver.

mance [131–133], and recent advances in silicon photonics have led to fully integrated optical signaling [135, 136]. These approaches pave the way for massively parallel optical communications. In order for optical interconnects to become viable alternatives to established electrical links, they must be low cost and have competitive energy and area efficiency metrics. Dense arrays of optical detectors require very low-power, sensitive, and compact optical receiver circuits. Existing designs for the input receiver, such as TIA, require large power consumption to achieve high bandwidth and low noise, and can occupy large area due to bandwidth enhancement inductors.

In this chapter we introduce a compact low-power optical receiver that scales well with technology to explore the potential of optical signaling for future chip-to-chip and on-chip communication. In the next section, we present the overall architecture of the proposed receiver. Next, the detailed circuit level implementation of the proposed receiver along with the sensitivity analysis is presented. System-level design considerations such as clocking and adaptation for the proposed receiver are discussed, and finally, we present experimental results from the evaluation of a 65nm bulk CMOS implementation of the optical receiver.

## 5.1 Low-Power Double-Sampling RC Front-End

The task of the optical receiver is to resolve the value of the incoming signal by sensing the changes in the magnitude of photodiode current. To minimize the transmit optical power, the receiver has to be able to resolve small optically generated currents from the photodiode.

In order to achieve a robust data resolution with low BER, the total input-referred noise current from the circuitry and the diode itself should be well below the optically generated current. In general, design of a low-noise front-end with a very high bandwidth is difficult and requires high electrical power consumption. In most optical receivers, the photodiode current is converted to a voltage signal. As discussed in previous chapter, a simple resistor, Figure 5.1(a), can perform the I-V conversion if the resulting  $RC$  time constant is on the order of the bit interval ( $T_b$ ) [135]. A voltage amplifier then amplifies the voltage swing for the following data resolution slicer block. Assuming that the voltage amplifier has a high bandwidth, the bit rate of such a front-end is limited by the input node time constant,  $RC_{in}$ , where  $C_{in}$  is the sum of the diode capacitance and other parasitic capacitors at the input node. The time constant of the input node sets a maximum limit on the resistor  $R$ . On the other hand, the maximum possible voltage swing at this node is equal to  $\Delta V = RI_{op}$  where  $I_{op}$  is the input photocurrent. It is clear that lower  $R$  values degrade the signal-to-noise-ratio (SNR) at the input. This results in a strong trade-off between the sensitivity and the bandwidth as they both depend on  $R$ . This trade-off between sensitivity and data-rate can be resolved by employing TIAs, as shown in Figure 5.1(b). TIA provides low impedance at the input node while introducing a high transconductance to convert the optical current from the photodiode into voltage. As shown in Equation 5.1, the maximum bandwidth, and hence the data rate, supported by TIA is proportional to its gain,  $A$

$$BW = \frac{1 + A}{2\pi R_F C_{in}}. \quad (5.1)$$

As a result, to achieve a high data-rate, a TIA with large gain-bandwidth product is required, which can result in high power consumption. Passive components such as inductors can be employed to enhance the bandwidth of TIA [112], but impose a significant area overhead. As discussed in Chapter 5, there is a strong trade-off between sensitivity and bandwidth of TIAs. In addition, their poor scalability makes them inadequate for dense arrays of optical links.

An alternative to TIA is the integrating front-end [117], shown in Figure 5.1(c). The




input signal from the photodetector is a single-ended, positive current. The injected charge is higher if the bit value is one but it is not necessarily zero when the bit value is zero. Therefore, in order to have a bipolar voltage change at the input of the receiver a constant charge is subtracted from the input capacitor for every bit. This is done by subtracting an adjustable current from the input through a feedback loop. By sampling the input voltage at the end of each bit period the received bit is resolved. The double-sampling technique allows for immediate demultiplexing at the front-end by employing multiple clock phases and samplers. It also eliminates the need for high gain stages, such as TIA, that operate at the input data rate. Another advantage associated with this technique is the inherent single-ended to differential conversion that happens at the front-end and reduces receiver sensitivity to common-mode interferences. The main advantage of this technique is that it mainly employs digital circuitry that allows for achieving considerable power saving by scaling to advance technology nodes. However, this technique suffers from voltage headroom limitations and requires short-length DC-balanced inputs such as 8B/10B encoded data.


In this chapter, we propose an RC front-end that employs a double sampling technique to break the trade-off between data rate and sensitivity without the described headroom problem, as shown in Figure 5.2(a). This technique allows for an input time-constant much larger than  $T_b$  ( $RC_{in} \gg T_b$ ) as opposed to TIA in which the input time-constant should be smaller than the bit time. The additional resistor,  $R$ , in the front-end automatically limits the input voltage and prevents out of range input voltages due to long sequences of ones or zeros. The input voltage can be expressed as

$$V_{PD} = V_{DD} - RI_1 e^{\frac{-t}{RC_{in}}} \quad (5.2)$$

for a long sequence of zeros following a long sequence of ones, where  $V_{PD}$  denotes the input voltage,  $R$  is the front-end resistance,  $C_{in}$  is the total capacitance at the input and  $I_1$  is the current due to a one input. Double-sampling can be applied to sample the input voltage at the end of two consecutive bit times,  $V[n-1]$ ,  $V[n]$ , Figure 5.2(b), and these samples are compared to resolve each bit ( $\Delta V[n] = V[n] - V[n-1] > 0$  results in one and



(a)



(b)

Figure 5.2: The proposed RC double-sampling front-end architecture (a). The exponential input voltage and the corresponding double-sampled voltage for a long sequence of successive ones (b).

$\Delta V[n] < 0$  results in zero). However, the resistor causes the double-sampled voltage,  $\Delta V[n]$ , to be input-dependent as expressed in Equations 5.3 and 5.4

$$\Delta V[n] = RI_1 e^{\frac{-nT_b}{RC_{in}}}, \quad (5.3)$$

$$\Delta V[n] = RI_1 (1 - e^{\frac{-T_b}{RC_{in}}}) e^{\frac{-(n-1)T_b}{RC_{in}}} = \Delta V[1] e^{\frac{-(n-1)T_b}{RC_{in}}}, \quad (5.4)$$

where  $T_b$  denotes the bit time. For instance, a one after a long sequence of zero generates larger  $\Delta V[n]$  than a one after a long sequence of ones. The dependency of the voltage difference on the input signal can be resolved by introducing a dynamic offset to the sense amplifier, Figure 5.3(a). This offset effectively increases the voltage difference  $\Delta V[n]$  for weak ones/zeros, and decreases it for strong ones/zeros as shown in Figure 5.3(b). We call this technique dynamic offset modulation (DOM). The idea behind this technique is to introduce an offset to the double-sampled voltage based on the value of the voltage at the input. As an example, a long sequence of ones followed by a long sequence of zeros is considered, Figure 5.3(b). The first one after zeros generates a large voltage at  $V_{PD}$ . As the number of successive ones increases, this voltage decays exponentially due to  $R$  and  $C_{in}$ . If the maximum double-sampled voltage is equal to  $\Delta V_{max}$ , DOM introduces an offset so that the sense amplifier differential input is  $\frac{\Delta V_{max}}{2}$ , regardless of the previous bits. For instance, an offset equal to  $-\frac{\Delta V_{max}}{2}$  is applied when  $\Delta V[n] = \Delta V_{max}$ , no offset is applied when  $\Delta V[n] = \frac{\Delta V_{max}}{2}$ , and an offset equal to  $\frac{\Delta V_{max}}{2}$  is applied if  $\Delta V[n] = 0$ .

Figure 5.4(a) shows a simple model of the double sampler where  $\Delta V[n]$  can be expressed in z-domain as

$$\Delta V(z) = (1 - z^{-1})V(z). \quad (5.5)$$

After subtracting the previous sample,  $V[n - 1]$ , the resulting voltage difference,  $\Delta V'[n]$  can be written in the z-domain as

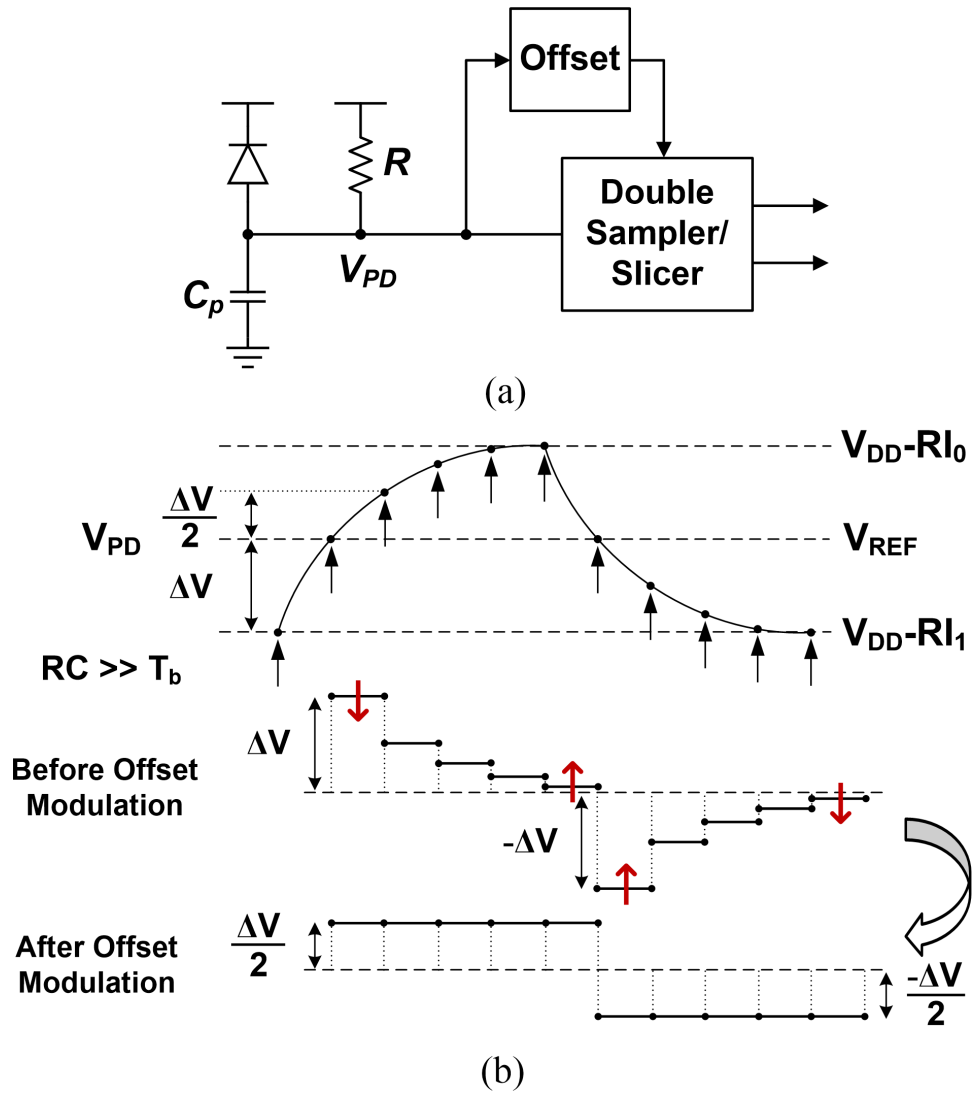


Figure 5.3: Modified RC front-end with DOM to resolve input dependent double-sampled voltage (a). The basic operation of DOM technique (b).


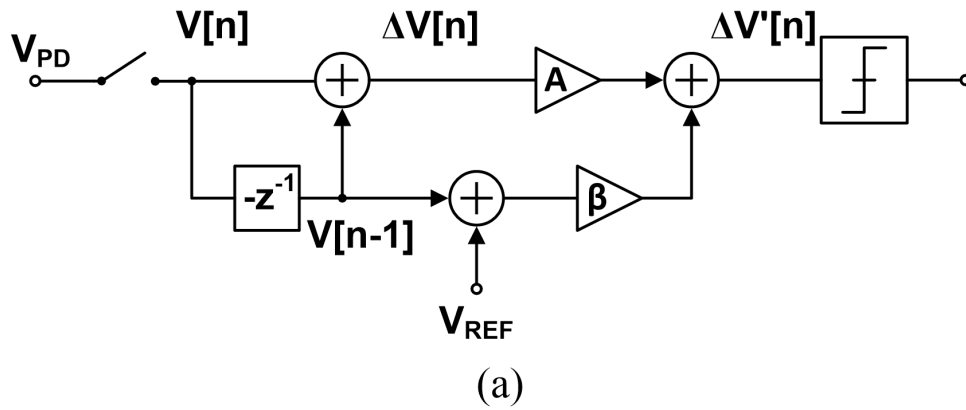


Figure 5.4: Block diagram of the offset modulation technique (a). The first sample is subtracted from the double-sampled voltage,  $\Delta V[n]$ , to make it constant regardless of the input sequence. Simulated operation of the DOM for a long sequence of ones showing  $\Delta V[n]$  before and after DOM (b).

$$\Delta V'(z) = (1 - z^{-1})V(z) + \beta z^{-1}V(z), \quad (5.6)$$

where  $\beta$  is the DOM coefficient and  $V(z)$  is equal to

$$V(z) = \frac{RI_1}{1 - e^{\frac{-T_b}{RC_{in}}}} z^{-1}. \quad (5.7)$$

In order to have a constant  $\Delta V'[n]$  regardless of the received input sequence, we should find  $\beta$  for which  $\Delta V'(z)$  is independent of  $z$ . By substituting 5.6 in 5.7, it can be shown that for

$$\beta = 1 - e^{\frac{-T_b}{RC_{in}}}. \quad (5.8)$$

$\Delta V'(z)$  will be independent of  $z$  and equal to

$$\Delta V'[n] = \frac{1}{2} \Delta V_{max} \forall n, \quad (5.9)$$

where

$$\Delta V_{max} = RI_1(1 - e^{\frac{-T_b}{RC_{in}}}). \quad (5.10)$$

$\Delta V_{max}$  is the double-sampled voltage due to a one (zero) following a long sequence of zeros (ones). Figure 5.4(b) shows the simulation results showing the double-sampled voltage before and after DOM. The target value of  $\beta$  can be determined using adaptive algorithms as described in section 5.2.1.

### 5.1.1 Front-End Sensitivity Analysis and Implementation

Figure 5.5 shows the top-level architecture of the receiver. The input current from the photodiode is integrated over the parasitic capacitor, while the shunt resistor ( $R$ ) limits the voltage.  $R$  can be designed to be adjustable to prevent saturation at high optical powers, allowing for a wide range of input optical power. As mentioned earlier, the employed double-

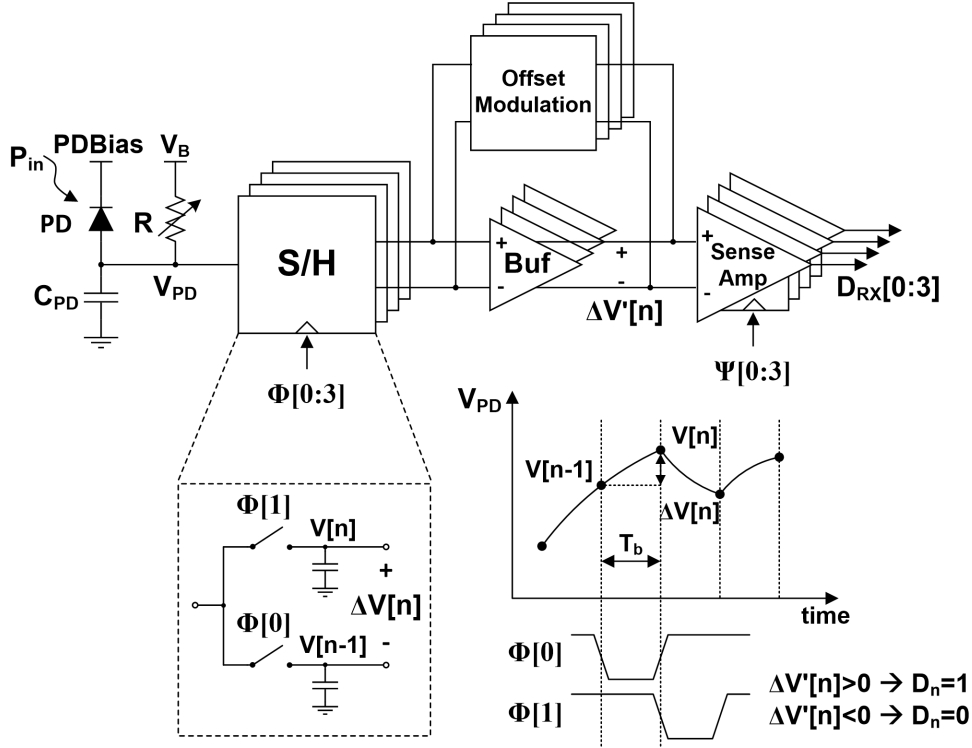


Figure 5.5: Top level architecture of the RC double-sampling front-end.

sampling technique allows demultiplexing by use of multiple clock phases and samplers. In this design a demultiplexing factor of four is chosen as the minimum possible demux factor to allow for proper operation of the double sampler and the following comparator stage. The front-end S/H is composed of a PMOS switch and the parasitic capacitor ( $C_S$ ) from the following stage. The optimum size of  $C_S$  is chosen considering the noise performance of the front-end and S/H speed as will be explained later. An amplifier with about 6dB of gain is inserted between the S/H and the comparator to provide isolation between the sensitive sampling node and the comparator and minimize kick-back noise. This also creates a constant common-mode voltage at the comparator input and improves its speed and offset performance. A StrongARM sense amplifier [137] is employed to achieve high sampling rate and low power. Figure 5.6 shows the transistor level schematic of the sense amplifier. Banks of digitally adjustable NMOS capacitors are employed to compensate the offset due to mismatch. DOM is implemented using a differential pair at the input of the sense amplifier [138]. This differential pair along with the resistors of the buffer stage form an amplifier with

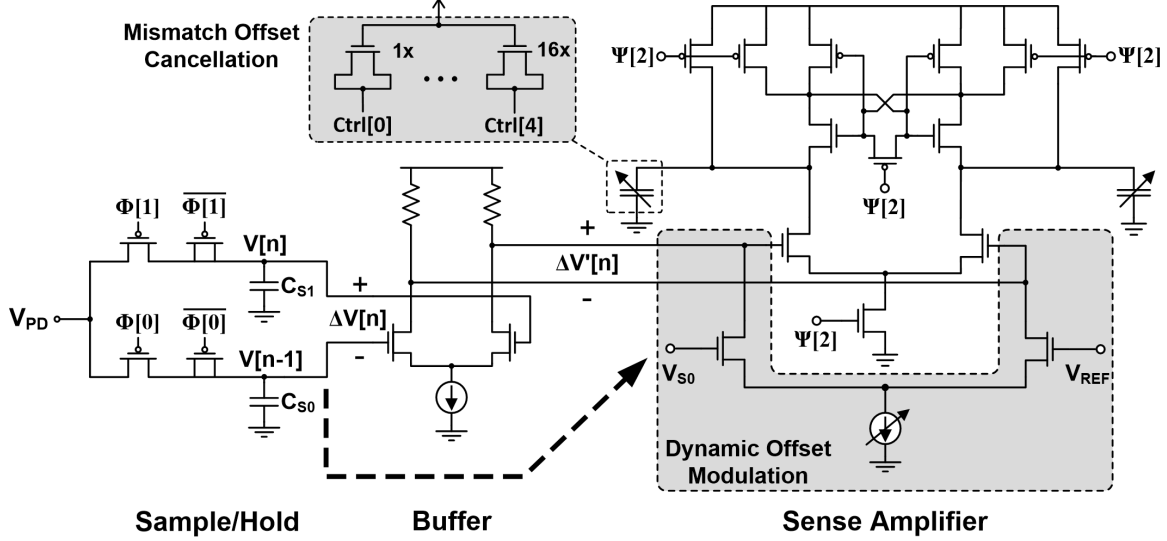


Figure 5.6: Detailed schematic of the RC double-sampling front-end.

variable gain  $\beta$ , which is adjusted through the variable tail current source. As the bandwidth of this amplifier and the buffer stage are equal,  $V[n-1]$  and  $\Delta V[n]$  experience the same delay in reaching the input of the sense amplifier. This eliminates any timing issue in the DOM operation. The dynamic offset is proportional to the difference between the sampled voltage ( $V[n-1]$ ) and a reference voltage ( $V_{REF}$ ).  $V_{REF}$  is defined as the average of the maximum ( $V_{DD} - RI_1$ ) and minimum ( $V_{DD} - RI_0$ ) voltages at  $V_{PD}$ , however, it should be noted that the resulting double-sampled voltage is constant regardless of the  $V_{REF}$  value, as discussed in the previous section. Here  $V_{REF}$  only sets the DC value of the double-sampled voltage, that is, with this value for  $V_{REF}$ , the resulting double-sampled voltage changes around zero. As shown in the previous section the double-sampled voltage is equal to

$$\Delta V_b = \frac{1}{2}RI_1(1 - e^{-\frac{T_b}{RC_{in}}}). \quad (5.11)$$

For  $T_b \ll RC_{in}$  Equation 5.11 can be approximated by

$$\Delta V_b \approx \frac{I_1}{2C_{in}} = \frac{\rho P_{avg} T_b}{2C_{in}}. \quad (5.12)$$

As a result the receiver sensitivity is a strong function of the bit period ( $T_b$ ), total input



capacitance ( $C_{in}$ ), photodiode responsivity ( $\rho$ ), and the total input referred noise:

$$P_{avg} = \frac{I_1}{\rho} = \frac{2C_{in}\Delta V_b}{\rho T_b}. \quad (5.13)$$

The receiver input capacitance is comprised of

$$C_{in} = C_{PD} + C_{pad} + C_{WB} + C_{int} + 2C_S, \quad (5.14)$$

where  $C_{PD}$  is the photodiode capacitance,  $C_{pad}$  denotes the bonding pad capacitance,  $C_{WB}$  is the wirebond capacitance,  $C_{int}$  is the input interconnect capacitance and  $C_S$  is the total sampling capacitance of each sampler. The required  $\Delta V_b$  is set by the minimum signal-to-noise ratio (SNR) for the target BER and the residual input-referred offset of the sense amplifier after correction,  $V_{offset}$ . As a result, the minimum required  $\Delta V_b$  is equal to

$$\Delta V_b = SNR \times \sigma_n + \frac{V_{offset}}{A}, \quad (5.15)$$

where  $\sigma_n$  is the total input voltage noise variance, which is computed by input referring the receiver circuit noise and the effective clock jitter noise.

The main sources of noise in the RC front-end are the sampler noise, buffer noise, sense-amp noise and finally clock jitter noise, as shown in Figure 5.7(a). The single-ended version is shown for simplicity. The sense amplifier is modeled as a sampler with gain and has an input referred voltage noise variance of [117]

$$\overline{\sigma_{SA}^2} = \frac{2kT}{A_{vsa}^2 C_A}, \quad (5.16)$$

where  $C_A$  is the internal sense amplifier node capacitance, which is set to approximately 15fF in order to obtain sufficient offset correction range. The sense amplifier gain,  $A_{vsa}$ , is estimated to be equal to near unity for the 0.8V common-mode input level set by the buffer output, resulting in a sense amplifier voltage noise sigma of 0.75mV. The buffer noise can be written as

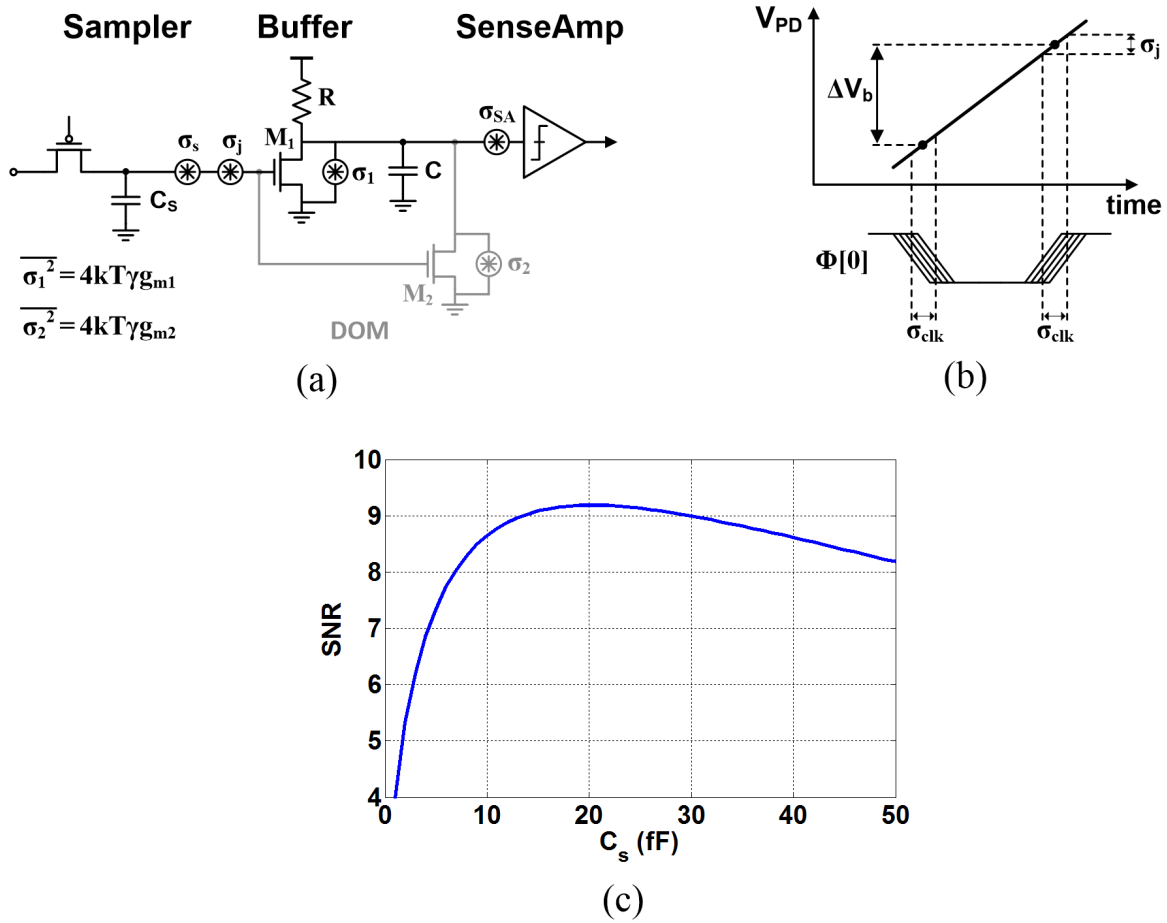


Figure 5.7: Schematic showing the noise sources in the front-end (a). This plot shows how the clock jitter is translated into the double-sampled voltage noise (b). There is an optimum range, 15-25fF, for the sampling capacitor to achieve maximum SNR (c).

$$\overline{\sigma_A^2} = \frac{2kT\gamma}{C}, \quad (5.17)$$

where  $\gamma$  is the transistor noise coefficient. According to simulation, the input referred voltage noise variance of the buffer stage is equal to  $0.6mV_{rms}$  while it provides about 6dB gain. Sampler voltage noise variance is equal to

$$\overline{\sigma_S^2} = \frac{2kT}{C_S}, \quad (5.18)$$

where the factor of two is due to the two sampling capacitors employed in the sampler block, which generate the differential input voltage to the buffer.

Clock jitter also has an impact on the receiver sensitivity because any deviations from the ideal sampling time results in a reduced double-sampled differential voltage as shown in Figure 5.7(b). This timing inaccuracy is mapped into an effective voltage noise on the input signal with a variance of

$$\overline{\sigma_j^2} = \left(\frac{\sigma_{CLK}}{T_b}\right)^2 \Delta V_b^2. \quad (5.19)$$

Using the measured clock jitter of about  $1ps_{rms}$ , it is estimated to be about  $0.5mV_{rms}$ . As shown in Figure 5.7(b), DOM also contributes some noise to the system. The total noise due to DOM can be written as

$$\overline{\sigma_{DOM}^2} = \frac{\beta^2}{2A^2} (\overline{\sigma_j^2} + \overline{\sigma_S^2} + 2\overline{\sigma_A^2}), \quad (5.20)$$

where  $A$  is the buffer DC gain. As  $\frac{\beta}{A} \ll 1$ , the noise contribution of the DOM is negligible. Combining the input referred circuit noise and effective clock jitter noise and ignoring  $\sigma_{DOM}$ , results in a total input noise power equal to

$$\sigma_n = \sqrt{\overline{\sigma_S^2} + \overline{\sigma_A^2} + \frac{\overline{\sigma_{SA}^2}}{A^2} + \overline{\sigma_{CLK}^2}}. \quad (5.21)$$

Figure 5.7(c) shows how the input SNR changes with the sampling capacitor for an

estimated total input capacitance of about 250fF. The maximum SNR is achieved for  $C_S$  equal to 20fF. However, a large sampling capacitor requires a large switch size in order to maintain performance at high data rate. This creates a trade-off between power consumption and data rate as clock buffer power consumption increases with the switch size. As a result in this design we chose  $C_S$  to be about 15fF to both achieve high SNR and remain in the flat part of the SNR curve to minimize sensitivity to process variation and at the same time reduce the power consumption due to clock buffers. A dummy transistor with half the size of the sampling transistor was also added to the sampler to minimize clock feed-through.

In order for the receiver to achieve adequate sensitivity, it is essential to minimize the sense amplifier input-referred offset caused by device and capacitive mismatches. While the input-referred offset can be compensated by increasing the total area of the sense amplifier [139], this reduces the buffer bandwidth by increasing input capacitance and also results in higher power consumption. Thus, in order to minimize the input-referred offset while still using relatively small devices, a capacitive trimming offset correction technique is used [140]. In this technique the capacitance is digitally adjusted to unbalance the amplifier and cancel the offset voltage. The residual offset is limited by the minimum offset canceling capacitance possible. As shown in Figure 5.6, digitally adjustable nMOS capacitors are attached to internal nodes and cause the two nodes to discharge at different rates and modify the effective input voltage to the positive-feedback stage. Using this technique, an offset correction range of 60 mV with a residual of 0.9 mV is achieved. The fixed input common-mode voltage provided by the buffers eliminates variability in the offset correction magnitude as the input signal integrates over the input voltage range.

The maximum input optical power is set by the requirements of the sampling switches and the transistor oxide break-down voltage. In order to accurately sample the input voltage in a bit time, on-resistance of the switch must be low enough. Given the 15fF sampling capacitor, for a 20Gb/s input, the on-resistance of the transistor must be less than  $1\text{K}\Omega$  in order for the resulting time constant to be small enough for the sampled voltage to settle to its final value within a bit interval,  $T_b=50\text{ps}$ . For the employed 65nm CMOS technology, this translates to 0.4V minimum possible voltage at the integrating node,  $V_{min}$ . On the other hand, the

maximum possible  $V_{PD}$  is equal to the oxide break-down voltage

$$V_{PDmax} - V_{PDmin} = RI_{1max} = R\rho P_{max}. \quad (5.22)$$

In this design the variable resistor ( $R$ ) at the input changes from about  $0.8\text{K}\Omega$  to  $4\text{K}\Omega$ , which allows the receiver to operate for up to  $0\text{dBm}$  input optical power with a photodiode responsivity of about  $1\text{A/W}$ . According to simulation, the receiver operates at higher input optical powers as the double-sampled voltage is quite large, however the excessive voltage at  $V_{PD}$  will stress the transistors connected to this node. The minimum input optical power is determined by the noise performance of the front-end as explained earlier in this section.

In the fabricated prototype, the DOM coefficient,  $\beta$ , is adjusted manually. In the next section an adaptive algorithm is introduced, which can automatically set  $\beta$  for optimum operation. In addition, the required clock signals are provided from off-chip. In a complete system the clock is generated on-chip using a CDR. In the next section we explain how a bang-bang CDR technique can be applied to the proposed receiver.

## 5.2 System-Level Design Considerations

In this section we will discuss a number of additional design considerations such as adaptation techniques for DOM, scaling behavior of the receiver, and suitable clocking techniques. The feasibility of these techniques is validated through circuit- and system-level simulations.

### 5.2.1 Adaptation of Dynamic Offset Modulation

As shown in section 5.1 the DOM coefficient depends on the front-end time constant ( $RC$ ). As a result, at the beginning of the operation and in order to maintain the operation of the receiver over slow dynamic variations such as temperature or supply drifts, an adaptation technique should be employed. We first consider the  $RC$  front-end without the DOM. As previously discussed, and as shown in Figure 5.8, consecutive ones or zeros generate double-sampled voltages that are not equal. In fact it is clear that the first bit generates a larger




Figure 5.8: Basic operation of the DOM gain,  $\beta$ , adaptation algorithm. The error signal is generated for a certain pattern depending on the difference between  $\Delta V[n]$  and  $\Delta V[n - 1]$ .

$|\Delta V[n]|$  than the second bit. This difference can be employed as an error signal to adjust the DOM coefficient,  $\beta$ . In this study we employ a bang-bang-controlled loop in which the sign of the error signal is used to correct the coefficient with constant steps. The corresponding UP/DN commands can be generated by simply duplicating the comparator part of the RC front-end. However, the new comparator should compare the two voltage differences. The modified sense amplifier for implementing this task is shown in Figure 5.9(b).

Figure 5.9(a) shows the input voltage waveform of the RC front-end upon receiving the data. As an example we choose a “11” data pattern. For this particular pattern, if  $\beta$  is equal to the optimal value, the two double-sampled voltages  $\Delta V[n - 1]$  and  $\Delta V[n]$  will be equal. Any error in  $\beta$  would lead to non-equal  $\Delta V[n - 1]$  and  $\Delta V[n]$ . The error direction, low or high, determines the sign of the error difference between the two double-sampled voltages for the “11” pattern. Therefore, if each double-sampled voltage is compared with its previous one, the result can be used for  $\beta$  adjustment. The operation is similar to normal data resolution where we compare each sample  $V[n]$  with a one-bit older sample  $V[n - 1]$ . Modified comparators, which generate  $P$  signals in Figure 5.9(c), are added to the front-end

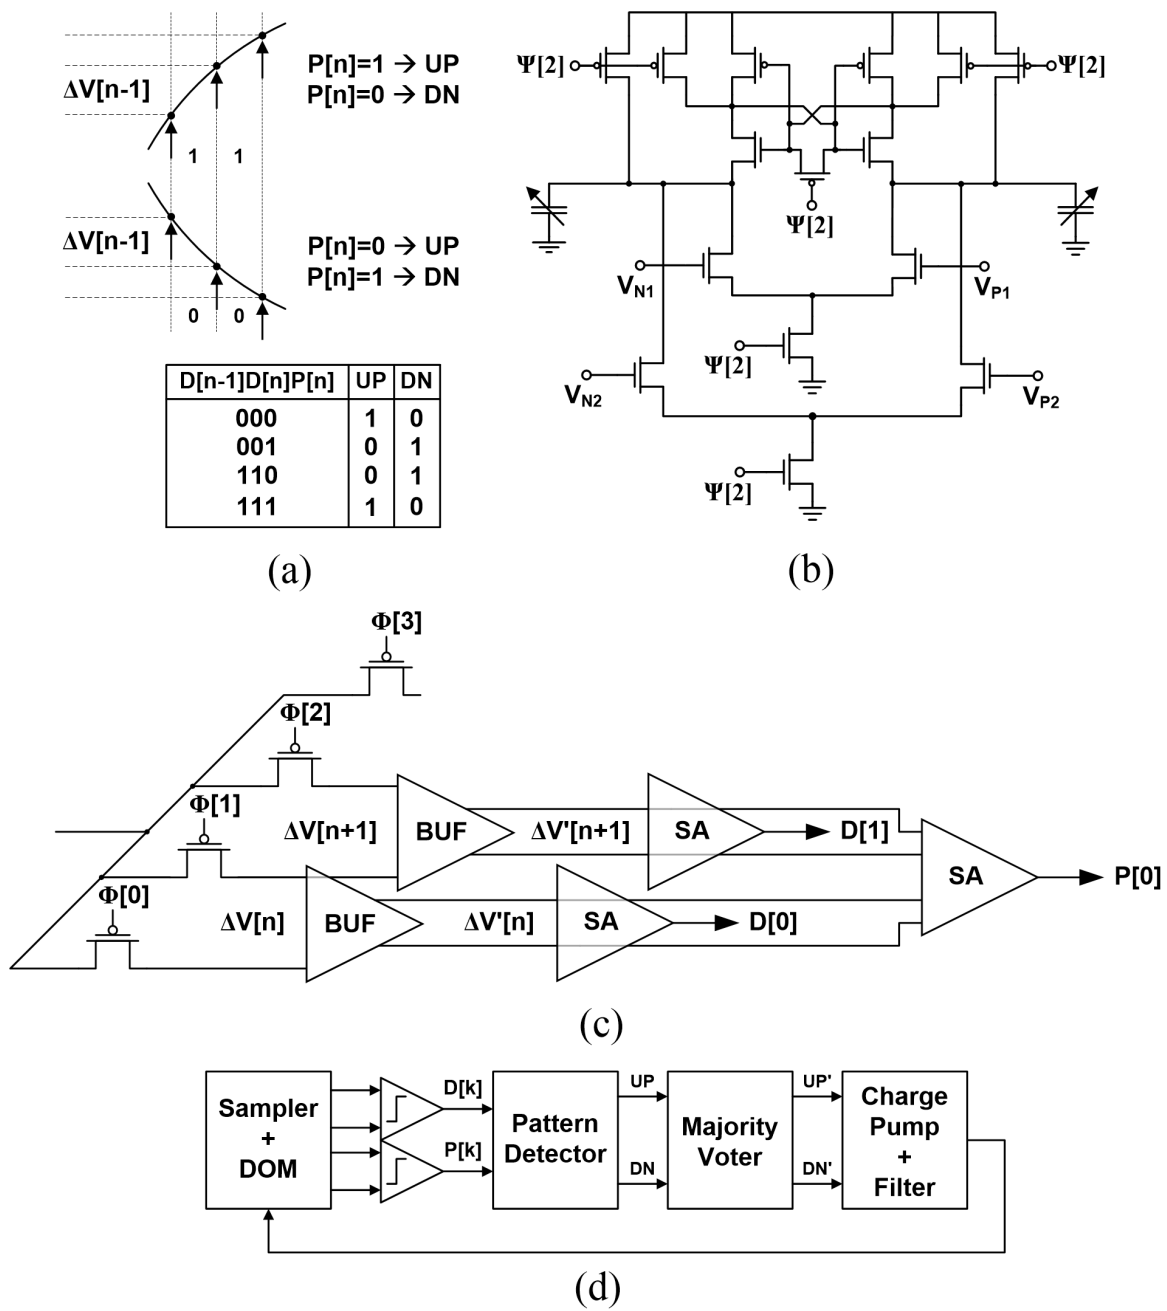


Figure 5.9: The input waveform and  $\beta$  error detection (a), modified sense-amplifier as the difference comparator (b), samplers and comparators for error detection (c), bang-bang  $\beta$  adaptation loop (d).

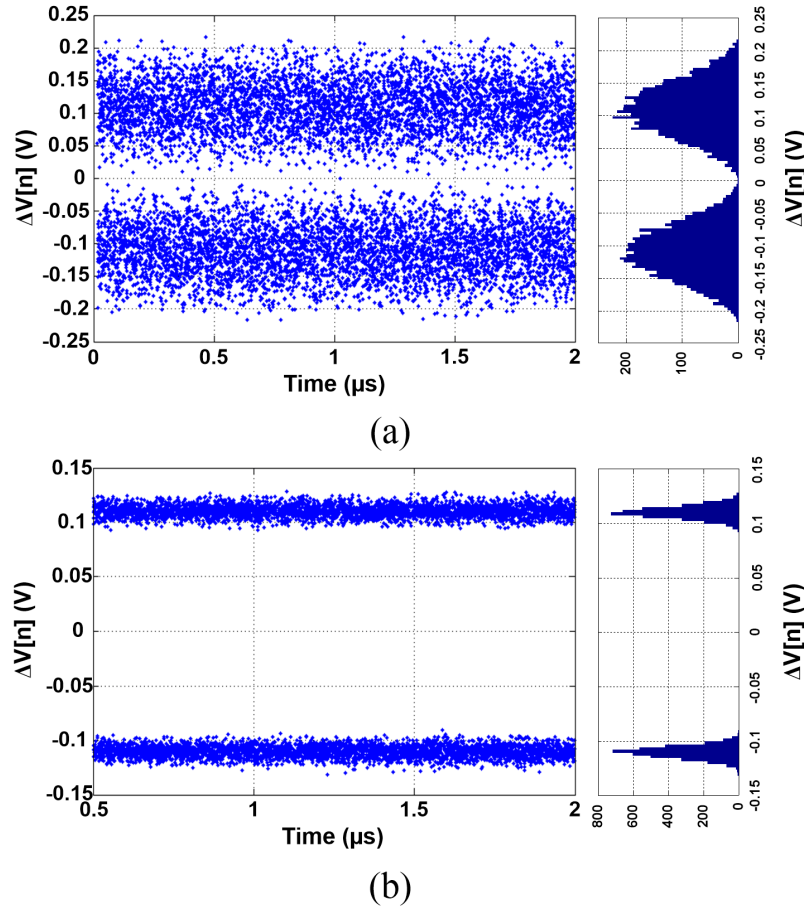


Figure 5.10: Simulated performance of the front-end before (a) and after (b) DOM adaptation. Gaussian noise with  $\sigma=10\text{mV}$  is applied at the sampler.

for this purpose. The error information for the adaptation loop is now the difference in the two double-sampled voltages and the 2-bit pattern that corresponds to samples  $V[n-1]$  to  $V[n+1]$ . Not all 2-bit patterns provide error information for the adaptation loop. The valid patterns are those that give equal  $\Delta V[n-1]$  and  $\Delta V[n]$  when  $\beta$  is adjusted to its optimal value. “11” and “00” are patterns that have such error information. The table in Figure 5.9(a) lists valid patterns with the corresponding condition for a meaningful result. Out of four possible 2-bit patterns, two give information for the adaptation loop. The effective probability of getting phase information from a random input is close to 0.5 when long sequences of ones or zeros do not happen often. Long sequences of ones or zeros result in near zero  $\Delta V[n-1]$  and  $\Delta V[n]$ , which does not provide error information.



A block diagram of a bang-bang-controlled gain adjustment loop is shown in Figure 5.9(d). This is a first-order loop and hence it is unconditionally stable. From the incoming data, two sets of samplers and comparators resolve the data,  $D$  signals, and raw error information,  $P$  signals. A pattern detector then generates the UP/DN correction commands for the loop. The UP/DN commands are filtered and used to adjust  $\beta$ . For instance, in the case where the  $RC$  time constant was equal to 200ps and  $T_b=50$ ps, optimal  $\beta$  is about 0.23 according to Equation 5.8. As a result when the loop is closed,  $\beta$  will converge to this quantity, also confirmed by closed-loop circuit-level simulations. Figure 5.10(a) shows the output of the sampler,  $\Delta V[n]$ , when DOM and adaptation circuits are disabled. Applying DOM with the adaptation loop creates a constant double-sampled voltage difference, as shown in Figure 5.10(b). The variation in the double-sampled voltage is due to the sampler noise of  $\sigma=10$ mV being incorporated in the simulations. The adaptation loop can be designed to operate only occasionally to correct for slow variations, and the same hardware can be reused for clock recovery as will be explained in section 5.2.4.

## 5.2.2 Double-Sampling Front-End Scaling

CMOS has been the leading technology for building integrated circuits for many years [1]. Owing to the scaling of this technology more functionality and faster devices become available with every new generation. The shrinking feature sizes and lower supply voltages offered by scaled CMOS technologies reduce the power consumption of most processing blocks while enabling faster operation. This reduction in power and increase in speed usually happens to digital systems. If the same paradigm can be applied to IOs, a huge reduction in power consumption and area of transceivers could be achieved, which would allow for higher numbers of IOs with higher data rates per chip. The combination of higher number of IOs and higher data rates per IO leads to a huge improvement in the overall chip-to-chip bandwidth.

This section examines how the performance of the proposed optical receivers will scale with advances in CMOS technology. We look into the challenges and problems introduced to the double sampling front-end design in deep sub-micron CMOS.




Figure 5.11: Optimum sampling capacitor size,  $C_S$  versus the photodiode capacitance,  $C_{PD}$ .

The removal of the high-speed gain stage in the proposed double sampling integrating front-end makes it a promising candidate for CMOS technologies that commonly have a poor gain-bandwidth product. While the digital behavior of the front-end prompts a good scalability, there are many subtle design issues as we move to more advanced technologies. In this section, we examine the scaling behavior of the double sampling/integrating front-end with respect to sensitivity, bandwidth, power consumption, and dynamic range.

### 5.2.2.1 Sensitivity

Our discussion in Section 5.1, and Equation 5.8 show that the required optical power for this receiver, strongly depends on the diode capacitance,  $C_p$  and the sampling capacitance,  $C_s$ . If these capacitances stay constant, the input optical power increases linearly with the data rate as the technology scales. This becomes particularly problematic as the number of transceivers in an array of parallel IOs increases. As a result of this effect, in order to accommodate a larger number of IOs, larger optical power has to be provided by optical transmitter. In order to alleviate this problem, both the photodiode capacitance and the sampling capacitance have to be scaled. Figure 5.11 shows that the optimum value for  $C_S$  reduces only if the diode capacitance is reduced.

Recent advances in design of photodiodes promise diode capacitances lower than 10fF [141, 142]. In Section 5.2.3 we will discuss the effects of the photodiode capacitance scaling on the performance of the RC front-end double-sampling receiver.

### 5.2.2.2 Data Rate

Scaling the feature sizes in CMOS technology with rate  $\alpha$  results in an increase in the switching speed of transistors linearly with almost reverse relation,  $\alpha^{-1}$ . The fundamental limit on the bandwidth of the double-sampling receiver is the aperture time of the samplers. For higher data rate the  $R_S C_S$  time constant should scale down, where  $R_S$  is the ON resistance of the NMOS switch. For the optimum sensitivity we need to keep capacitance  $C_S$  constant. Thus, for higher data rates the resistance  $R_S$  should scale down. This is possible by keeping the width of the pass-transistor (in microns) constant, while the length is equal to the minimum channel size of the technology. As technology scales, the transistor channel length scales down and hence  $R_S$  also decreases. Similarly the on-chip clock frequency for generation of multi-phases increases and allows higher data rates.

In addition to the sampler aperture time, the photodiode capacitance has a significant impact on the maximum possible data rate of this receiver which is mandated through the trade-off between sensitivity and data rate. This will be discussed in more detail in Section 5.2.3.

### 5.2.2.3 Power Consumption

The power consumption in high-speed IOs is the most important issue. Increasing the number of IOs per chip is possible only if the power consumption per IO reduces. In digital systems the power consumption is dominated by the dynamic power for switching internal capacitances,  $P = CV^2f$ , as well as the leakage dissipations. As technology scales, the dynamic power consumption of a similar block reduces with a factor of almost  $\alpha^2$ . This is because the capacitances reduce with  $\alpha$ , the power supply reduces with  $\alpha$ , and the frequency increases with  $\alpha^{-1}$ . For the proposed front-end, the capacitances of the samplers/comparators are

dictated by the sensitivity requirements and hence the photodiode capacitance size. Therefore, the power consumption of the front-end scales only as  $\alpha$ . It should be noted that the power consumption in the following stages (sense-amplifier and SR latch), wires, and clock generation circuits scale as  $\alpha^2$ . For instance, scaling to 28nm technology with 1V power supply and constant diode capacitance, results in a power reduction of almost a factor of two for the same data rate, which agrees well with simulations. It is important to mention that the supply scaling has slowed down. As a result, the front-end power will decrease at a slower rate; however as long as scaling results in smaller feature size and hence parasitic capacitance, power reduction can be achieved by employing advanced technologies.

#### 5.2.2.4 Dynamic Range

The dynamic range of the proposed double-sampling receiver relies on the size of the resistance in the front-end. As mentioned earlier, this resistor at the front-end limits the voltage to  $V_{dd} - RI_1$ . The interesting point about this technique is that as long as the input time constant (RC) is much larger than the bit time ( $T_b$ ) the resulting double-sampled voltage is constant and equal to


$$V_{PDmax} - V_{PDmin} = RI_{1max} = R\rho P_{max}. \quad (5.23)$$

Therefore, by adjusting the input resistor, different optical powers can be accommodated. The maximum optical power that can be received without sacrificing the sensitivity occurs when the resistor size becomes so small that the input time constant is comparable to the bit time. On the other hand, the minimum input voltage is limited by the fact that the PMOS sampling switches will introduce large on-resistance at low input voltages. As the time constant of the sampler exceeds a fraction of the bit time, a larger error will be generated. Due to the scaling of  $V_{dd}$ , the input voltage range becomes smaller and smaller. For a 28nm technology with  $V_{dd} = 1$  this range is about 400mV for a total capacitance of 250fF and 20Gb/s data rate. This translates to a minimum of 1K $\Omega$  of resistance and about -3dBm maximum optical power. According to simulation, the receiver operates at higher input


optical powers. This is due to the fact that, even though the sampled voltage is smaller than the actual voltage, the double-sampled voltage is quite large. The measurement results agree with this observation and the receiver operates with 0dBm of received optical power.

### 5.2.3 Photodiode Capacitance Scaling

Silicon photonics has offered high-performance optical components, such as Germanium photodiodes, waveguides and modulators. This integration allows for very small photodiode parasitic capacitance. In this section we investigate the effect of photodiode capacitance scaling on the performance of the proposed receiver. According to Equation 5.10 the double-sampling voltage is inversely proportional to the photodiode capacitance. As a result, larger double-sampling voltage can be achieved using a smaller parasitic capacitance. This allows for scaling of the receiver sensitivity,  $P_{avg}$ , for a fixed data rate. This argument is valid under the assumption that no charge sharing happens between the photodiode capacitance and the receiver sampling capacitors. In order to minimize this charge sharing, a certain ratio between the photodiode capacitance and the sampling capacitance has to be kept. In this design this ratio is chosen to be about 10. Therefore, while we scale the photodiode capacitance, the sampling capacitor should also scale with the same rate. This in turn increases the  $\frac{kT}{C}$  noise of the sampler and degrades the front-end SNR. However as the noise is inversely proportional to the square-root of the capacitor size, the overall SNR and hence the sensitivity of the receiver increases proportionally to the square-root of the photodiode scaling factor as shown in Figure 5.12(a). For instance, in the case of a photodiode capacitance of 50fF, the double-sampling receiver achieves about  $34\mu\text{A}$  of current sensitivity at 20Gb/s, which improves to  $17\mu\text{A}$  at 10Gb/s as the integration time is doubled. For an extinction ratio of 10dB this translates to about -20dBm sensitivity. Therefore, the proposed receiver can greatly benefit in terms of sensitivity from advanced photodiode technologies with small parasitic capacitances and efficient integration techniques such as flip-chip bonding or copper pillars, as explained in the previous chapter. On the other hand, as the photodiode capacitance scales to 10-20fF, in the case of monolithically integrated photodiodes, the charge



(a)



(b)

Figure 5.12: Receiver current sensitivity versus photodiode capacitance, with and without scaling sampling capacitor (a). Receiver data rate versus photodiode capacitance for  $100\mu A$  sensitivity (b).

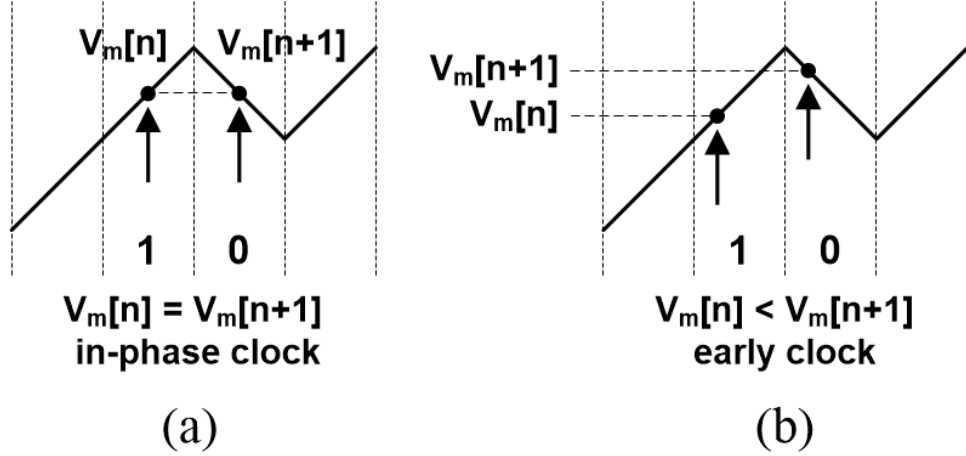


Figure 5.13: 2x-oversampled phase detection for the proposed receiver.

sharing between the photodiode capacitance and the sampling capacitance limits the sensitivity of the receiver.

The receiver maximum data rate is also a function of the photodiode capacitance. According to (25), for a given sensitivity, the data rate ( $R_b$ ) can be increased by scaling the photodiode capacitance.

$$R_b = \frac{1}{T_b} = \frac{\rho P_{avg}}{2C_{in}\Delta V_b} = \frac{\rho P_{avg}}{2C_{in}(SNR \times \sigma_n + V_{offset})}. \quad (5.24)$$

As mentioned earlier, in order to minimize charge sharing, the sampling capacitor scales with the same rate as the photodiode capacitance. As a result the input referred-noise,  $\sigma_n$ , changes accordingly. For the target RX sensitivity of  $100\mu\text{A}$ , Figure 5.12(b) shows how the data rate changes as the photodiode capacitance scales. The maximum achievable data rate is ultimately limited by the speed of transistors.

## 5.2.4 Clocking

An interesting problem in a clocked integrating front-end is to recover the clock from the incoming data. As mentioned in section 5.1.1 the clock jitter could be one of the limiting factors in the receiver sensitivity. As a result, an efficient low-jitter clocking technique is crucial. For highly parallel links, a dual-loop CDR [143] can be employed with one loop for




Figure 5.14: The input waveform and baud-rate phase detection, for in-phase (a), and out of phase clock (b).

the frequency synthesis, which can be shared between all the channels, and the other for phase correction in each channel (alternatively in a source-synchronous clocking scheme the frequency synthesis loop can be eliminated and a phase correction loop will be sufficient). An alternative technique is to employ a forwarded-clock scheme in a WDM link using one of the channels (wavelengths), which allows for simple phase correction loops to set the optimal sampling time.

The most common phase detection technique employed in electrical signaling is the 2x-oversampled phase detector known as Alexander phase detector [52]. A similar technique can be applied to the proposed double-sampled front-end. Figure 5.13 shows the DOM output voltage upon receiving a one-zero transition. The front-end samples the signal in the middle of each bit-period,  $V_m[n]$  and  $V_m[n - 1]$ . At any transition, if the clock is in-phase with the data, the two samples taken at the middle of these consecutive non-equal bits are expected to be equal. Any phase error would cause these two voltages to be different. This difference can be used as an error signal to adjust the phase of the sampling clocks. In order to implement the clock recovery loop, we can duplicate the samplers/comparator part of the front-end. This set of samplers/comparators needs to be clocked with an extra clock phase, shifted by half a bit-period.

Removing the extra phases for oversampled phase recovery can help to reduce the power






Figure 5.15: Electrical measurement setup.

consumption in the oscillator and clock buffers and relax the difficulties of phase spacing control. The RC front-end allows us to create an efficient baud-rate phase recovery scheme similar to [89], [144], based only on data samples as shown in Figure 5.14. The difference between this method and the one proposed in [144] is that instead of extracting the phase error data from the sampled input, the double-sampled voltage difference at the output of the DOM,  $\Delta V[n-1]$  and  $\Delta V[n]$  in Figure 5.14, are employed. This is similar to a  $\beta$  adaptation loop except that instead of looking at a 2-bit pattern, 4-bit patterns are investigated. As an example we choose a “0110” data pattern in Figure 5.14 to explain the operation of this technique. It is clear from the figure that for this particular pattern, if the sampling clock is in-phase with the incoming data,  $\Delta V[n-1]$  and  $\Delta V[n]$  will be equal. Any error in the sampling clock phase would lead to non-equal  $\Delta V[n-1]$  and  $\Delta V[n]$ . The phase error direction, early or late clock, determines the sign of the error difference between the two samples for this pattern. Therefore, if each two consecutive double-sampled voltages are compared, the resulting information can be used for phase recovery. The valid patterns for phase corrections are those that give equal  $\Delta V[n-1]$  and  $\Delta V[n]$  when the clock is synchronized with the incoming data. “1001” and “0110” are patterns that have complete early/late phase information. Most other patterns have conditional phase information, e.g.




Figure 5.16: Photodiode current emulator.

“1110” only gives valid results when the clock leads the input. Due to the lower update density in the baud-rate phase detection technique, the overall loop gain is smaller compared to the conventional 2x-oversampling by almost a factor of 2.67 [144]. As a result, the 2x-oversampling phase correction loop provides higher bandwidth, for identical loop filter and charge-pump, and hence superior jitter tolerance. On the other hand, the baud-rate phase detector has the additional advantage of being less sensitive to clock phase errors, as the same clocks are used for both the data and phase samples, whereas the 2x-oversampling detector relies on quadrature phase matching.

Another important aspect of the phase correction loop is its effect on the operation of the  $\beta$  correction loop. As explained earlier, these two loops operate based on the same correction signal, P, to minimize the difference between the two consecutive double-sampled voltages,  $\Delta V'[n-1]$  and  $\Delta V'[n]$ . As a result, they can operate concurrently to adjust  $\beta$  and the clock phase. This has been validated in simulation for a PRBS-7 pattern when the initial phase is about half UI apart from the optimal point. The bandwidth of the  $\beta$  and phase correction loop in this simulation was about 2MHz. This experiment was repeated for the case where the clock phase was leading and lagging with respect to the optimal clock phase as well as under- and over-compensated  $\beta$ .

As mentioned earlier in this section the only difference between the  $\beta$  adjustment loop and the CDR loop is the length of the pattern that should be monitored. As a result,




Figure 5.17: Receiver sensitivity characteristics for different data rates.

the same hardware ( $P$  comparators) employed in the  $\beta$  adaptation loop can be reused to perform clock recovery, except for the pattern detection logic. This allows for saving in power consumption and area.

### 5.3 Experimental Results

The prototype was fabricated in a 65nm CMOS technology with the receiver occupying less than  $0.0028mm^2$ . It is composed of two receivers, one with a photodiode emulator and one for optical testing with a photodiode. In the first version, an emulator mimics the photodiode current with an on-chip switchable current source and a bank of capacitors ( $C_{PD}$ ) is integrated to emulate the parasitic capacitances due to photodiode and bonding (PAD and wirebond), Figure 5.16. The four phases of clock are provided by an off-chip signal generator as shown in Figure 5.15. An on-chip CML-to-CMOS converter generates the full-swing clocks for the receiver. The on-chip clock was measured to have about 9ps peak-to-peak jitter.




Figure 5.18: Current and voltage sensitivity versus data rate.




Figure 5.19: Power consumption and efficiency at different data rates.




Figure 5.20: The receiver power breakdown.

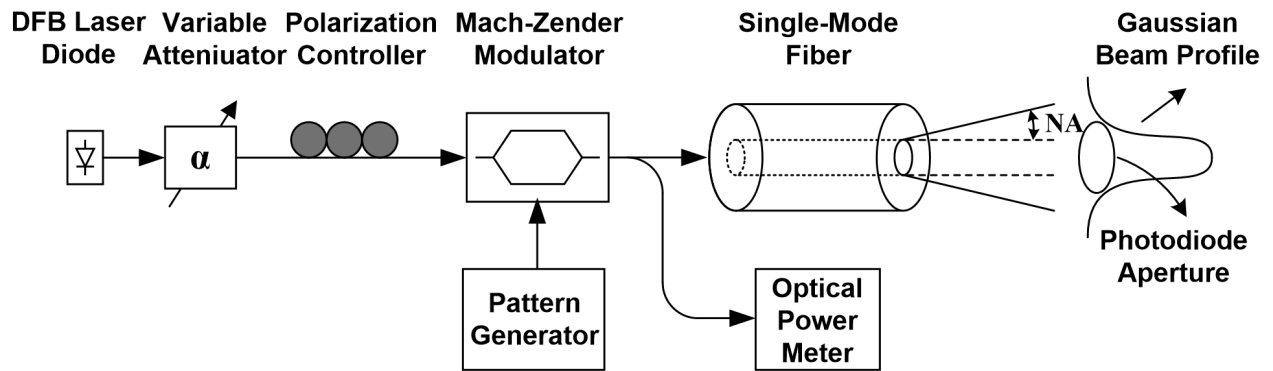


Figure 5.21: Optical test set-up.

The functionality of the receiver was first validated using the on-chip emulator and PRBS-7, 9, 15 sequences.  $R$  and  $C_{in}$  were chosen to be 2.2K $\Omega$  and 250fF ( $RC_{in} > 550$ ps). Figure 5.17 shows how the bit error rate changes with the input current at 14.2Gb/s, 16.7Gb/s, 20Gb/s, and 24Gb/s. For all these data rates the condition  $T_b \ll RC_{in}$  is valid. The receiver achieves about 75 $\mu$ A of sensitivity at 14.2Gb/s, which reduces to 160 $\mu$ A at 24Gb/s. Due to the integrating nature of the receiver the current sensitivity almost linearly increases with data rate, as shown in Figure 5.18. The voltage sensitivity of the receiver is measured to be about 13mV up to 20Gb/s and increases to 17mV at 24Gb/s, which is believed to be partly due to degradation of the eye opening at the emulator input. The receiver power consumption (including all clock buffers) at different data rates is shown in Figure 5.19. The power increases linearly with the data rate as the receiver employs mostly digital blocks. The receiver offers a peak power efficiency of 0.36pJ/b at 20Gb/s data rate. The power breakdown of the receiver is also shown in Figure 5.20, which confirms that the power is dominated by digital blocks. In order to validate the functionality of the DOM for long sequences of ones or zeros, a 200MHz square-wave current was applied as the input to the receiver while the front-end sampled the input at 20Gb/s. In this case, 50 consecutive zeros were followed by 50 ones. For an input time constant of about 0.55ns, this number of zeros or ones pushes the input to the flat region, where close to zero double-sampling voltage,  $\Delta V[n]$ , is obtained. Enabling the DOM resulted in error-free detection of the received pattern.

In the second set of measurements, the receiver was wire-bonded to a high-speed photodiode and tested at different data rates. The photodiode, bonding pad, wire-bond, and the receiver front-end are estimated to introduce more than 200fF capacitance. Figure 5.21 shows the optical test setup. The optical beam from a 1550nm DFB laser diode is modulated by a high-speed Mach-Zender modulator and coupled to the photodiode through a single-mode fiber. The optical fiber is placed close to the photodiode aperture using a micro-positioner. The responsivity of the photodiode at this wavelength is about 1A/W. As the beam has a Gaussian profile, the gap between the fiber tip and the photodetector causes optical intensity loss. This, combined with the optical connectors and misalignment introduces some loss, which can be characterized by comparing the sensitivity in the two experiments.




Figure 5.22: Micrograph of the receiver with bonded photodiode (a). Coupling laser through fiber to the photodiode (b).

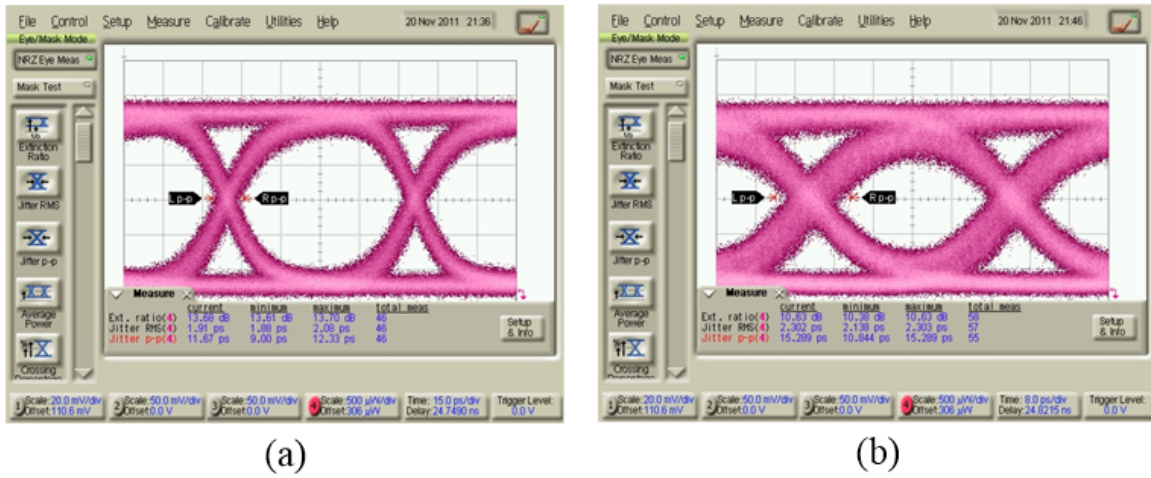


Figure 5.23: Optical input eye-diagram to the photodiode at 14Gb/s (a) and 24Gb/s (b).

Current and optical power sensitivity are related according to Equation 5.23

$$P_S = \frac{\rho I_S}{2} \frac{1 + 10^{-\frac{ER}{10}}}{1 - 10^{-\frac{ER}{10}}}, \quad (5.25)$$

where  $P_S$  is the optical power sensitivity,  $I_S = I_1 - I_0$  is the current sensitivity and  $ER$  is the extinction ratio. The measured extinction ratio at 14.2Gb/s is about 13dB using the external modulator. As a result, the nominal optical sensitivity according to the current sensitivity of  $75\mu A$  will be equal to -14dBm. The difference between the nominal and measured optical sensitivities is about 5dB, which is believed to be due to the coupling loss. This difference




Figure 5.24: Optical sensitivity at different data rates.

grows as the data rate increases due to the limited bandwidth of the external modulator. Therefore, the sensitivity can be improved by employing advance optical packaging technologies. Figure 5.24 shows how the sensitivity of the receiver changes with data rate. Note that the coupling loss is not considered in this plot. The receiver achieves more than -12.5dBm of sensitivity at 10Gb/s, which reduces to -7.3dBm at 18.6Gb/s and -4.6dBm at 24Gb/s. The maximum optical power at which the receiver was tested is 0dBm. This is the maximum power available from the measurement setup.

As mentioned in previous section, the variable resistor at the front-end allows for a wide range of optical input power. For large input optical power the variable input resistor can be reduced to avoid saturation. Figure 5.25 compares the calculated voltage sensitivity ( $BER < 10^{-12}$ ) achieved for the electrical and optical input experiments. In both cases the voltage is calculated by using Equation 5.12 and the measured current sensitivity. As expected the sensitivity of the receiver degrades with data rate. In the electrical test the receiver achieves almost constant voltage sensitivity regardless of the data rate. However for the optical experiment, the calculated voltage sensitivity degrades as data rate increases. The excessive sensitivity degradation in the optical test is partly due to the wire-bonded photodiode and limited bandwidth of the optical modulator, which causes reduced vertical






Figure 5.25: Comparison between voltage sensitivity for electrical and optical measurement.

Technology	65nm Bulk CMOS
Supply	1.2V
Data Rate	24Gb/s
Power Consumption	0.36mW/Gb/s
Sensitivity	-4.6dBm
RX $C_{in}$	>200fF
Area	0.0028mm <sup>2</sup>

Table 5.1: OPTICAL RECEIVER PERFORMANCE SUMMARY.

and horizontal eye opening as shown in Figure 5.23. Table 5.1 summarizes the performance of the proposed optical receiver and compares it with prior art.

## 5.4 Summary

The power consumption and area of receiver front-end circuitry is a critical design aspect of parallel optical interconnects for future arrays. While transimpedance amplifiers offer a relatively high sensitivity, the power consumption and area overhead of the TIAs make them inadequate for highly parallel IOs. Integrating front-ends can reduce the power consumption by avoiding an analog amplifier that runs at the bit-rate. However, the limited headroom, specially for highly scaled technologies, and the required encoding scheme to avoid long sequences of ones or zeros are major drawbacks attributed to this technique. This chapter described a double sampling  $RC$  front-end that offers low power consumption, high data

rate and relatively good sensitivity. The high data-rate is possible through a de-multiplexing scheme where multiple samplers and comparators run in parallel, while using a single photodiode. The resistive front-end provides an automatic mechanism to limit the photodiode voltage. The double sampling technique provides a self-reference for the data resolution and a feed-forward path dynamic offset modulation removes the input dependent variation in the double-sampled voltage. An efficient adaptation algorithm for adjustment of the feed-forward gain is proposed and investigated. The application of the baud-rate clock recovery to the receiver is also analyzed.

The proposed receiver discussed was implemented in 65nm CMOS that supports up to 24Gb/s of data rate. The low-voltage RC front-end receiver uses mostly digital building blocks and avoids the use of linear high-gain analog elements. The proposed receiver employs double-sampling and dynamic offset modulation to resolve arbitrary patterns. The receiver consumes less than 0.36pJ/b power at 20Gb/s, and operates up to 24Gb/s with -4.7dBm optical sensitivity ( $\text{BER} < 10^{-12}$ ). Since a large percentage of power consumption is due to the clock buffers and digital blocks, the overall power consumption can greatly benefit from technology scaling. It is also shown that this design is highly suitable for hybrid integration with low-capacitance photodiodes to achieve high optical sensitivity and high data rate. Experimental results validate the feasibility of the receiver for ultra-low-power, high-data rate and highly parallel optical links.

# Chapter 6

## On-Chip Wires: Characteristics, Models, and scaling

The miniaturization trend for CMOS integrated circuits (IC) has resulted in a tremendous cost advantage and performance improvement. The cost and performance will continue to drive miniaturization along the lines of Moore's law [1] for as long as fundamental limits allow. Aggressive shrinking of devices combined with a continuous increase in the chip size results in more functionality within the chip. Scaling results in faster transistor delays (lower delays) due to smaller channel size and less parasitics.

Unfortunately, miniaturization is not applicable to all the components of an IC. In particular, the speed of wires, which connect the transistors, is rapidly becoming a performance bottleneck as their characteristics degrade by scaling [147], [148], [149]. In the early days of VLSI, wires were wide and thick with relatively low resistance. Therefore, wires could be treated as ideal equipotential nodes with lumped capacitance. As a result of scaling, wires are becoming narrower, which drives their resistance to the point that the wire  $RC$  delay exceeds gate delay. Moreover, miniaturization exacerbates interconnect performance as increasing transistor numbers demand a proportionate increase in connectivity. Figure 6.1 shows how the total on-chip interconnection length has scaled over different technology nodes [150]. One way to accommodate this increasing number of interconnects is to increase the chip area. However, chip area can only increase very slowly as it is constrained by a balance between performance, cost and the reliability. There are different factors driving




Figure 6.1: On-chip interconnect length trend.

the chip area. It is desired to minimize the chip area to shorten the length of the wires (this improves performance) and to enhance manufacturing yield, which helps reducing the chip cost. On the other hand, in order to meet power density constraints, the chip size has to be increased. Chip temperature is very important as it affects both reliability and speed of various components on an IC.

Because the increasing complexity of VLSI chips continuously demands more from interconnects, a systematic and a realistic study of limitations of the currently used electrical interconnects under scaling is of great importance. In the remainder of this chapter we investigate the characteristics of the on-chip wires and develop simple models to analyze their behavior. Finally we investigate the effect of scaling on on-chip wires and discuss the resulting challenges.

## 6.1 Wire Characteristics

Three generations of Intel process technologies, shown as cross-sectional photographs in Figure 6.2, reflect advances in on-chip wiring. A 130nm technology from 2000, using six layers of copper metal. A 65nm technology from 2004, using eight layers of copper metal [151],




Figure 6.2: On-chip metal stack in different technology nodes, 130nm (a), 65nm (b), and 32nm (c).

and a 32nm technology from 2008, employing nine copper metal layers [152]. Between these three technologies the wires' cross-sectional areas and spacings have dropped dramatically. The importance of cross-sectional area and spacing lies in their effects on the wire electrical characteristics, resistance and capacitance. The following sections describe geometric models for resistance and capacitance that are based on cross-sectional area and spacing. We will also discuss wire inductance and why we can ignore its role in the performance modeling.

### 6.1.1 Resistance

All wires have a finite conductance, representing the ability of the wire to carry a charge flow. Aluminum and copper wires, which have been used in most CMOS processes, have a resistivity of  $3.5\text{-}4.0\mu\Omega\cdot\text{cm}$  and  $2.6\mu\Omega\cdot\text{cm}$ , respectively. Resistance (per unit length) can be approximated by the material resistivity divided by the conductor's cross-sectional area, but several wire non-idealities affect this model.

Most processes prior to 180nm generation employed aluminum wires. Modern processes use copper to reduce the resistivity and also to obtain better electromigration characteristics.




Figure 6.3: The schematic profile of diffusion barrier layer (a). SEM cross section of the diffusion barrier layer [171] (b).

For copper wires, a thin barrier layer, which has much lower conductivity, is required to prevent copper from diffusing into the surrounding oxide, Figure 6.3. As a result of this diffusion barrier layer and reduced effective cross-sectional area, the resulting resistance of the copper wires decreases. Moreover, the barrier layer may not be deposited evenly, which further degrades the wire conductivity. Another wire non-ideality is the surface scattering effect. When traveling along a wire, electrons inelastically scatter off lattice bonds at the edges of wires due to the surface roughness. As the wire dimensions grow smaller, the mean free path of electrons will reduce, effectively increasing the material resistivity [153].

Using advance fabrication techniques such as atomic layer deposition (ALD) an almost constant thickness barrier layer can be created [147], which results in the wire resistance being equal to

$$R_{wire} = \alpha_{scatter} \times \frac{\rho_{effective}}{(t - \delta)(w - 2\delta)}, \quad (6.1)$$

where  $\alpha_{scatter}$  is the resistance degradation factor due to surface scattering,  $\rho$  is the resistivity of the the thin film copper,  $\delta$  is the barrier layer thickness,  $w$  is the wire width, and  $t$  is the wire thickness.




Figure 6.4: A simple capacitance model for on-chip wires.

### 6.1.2 Capacitance

All wires have capacitance, modeling the ability of the wire to store electrical charge. To accurately model the capacitance of an on-chip wire both bottom-plate and fringing capacitances have to be taken into account [154]. The fringe capacitance models the field lines emerging from the edge and top of the wire. On-chip wires suffer from large aspect ratio (the ratio between the height and width) which makes the fringing portion of the capacitance significant. Therefore, capacitance is best modeled by four parallel-plate capacitors for the top, bottom, right, and left sides, as shown in Figure 6.4, plus a constant. This extra term lumps all the fringing field terms together and approximates their sum as a constant.

The total wire capacitance can be approximated as [155]

$$C_{wire} = \epsilon_0 \left( 2M \epsilon_{horiz} \frac{t}{s} + 2\epsilon_{vert} \frac{w}{h} \right) + C_{fringe}, \quad (6.2)$$

where  $\epsilon_{horiz}$  and  $\epsilon_{vert}$  represent the horizontal and vertical relative permittivity. This difference occurs in technologies that leverage low-k materials. The top and bottom plates are typically modeled as being grounded as they typically constitute a collection of orthogonally-routed conductors that, averaged over the length of the wire, maintain a constant voltage.

Capacitors to the left and right, on the other hand, have data-dependent effective capacitances that can vary: if the left and right neighbors switch in the opposite direction to the wire, the effective sidewall capacitances double, and if they switch with the wire, the effective sidewall capacitances approach zero. We model this multiplication effect by varying the  $M$  parameter in Equation 6.2 between 0 and 2; this is known as Miller factor ( $M$ ). These left and right neighbors are also the worst offenders for noise injection. The fringe term depends only weakly on geometry.

### 6.1.3 Inductance

Inductance is beginning to be important for accurately modeling on-chip wires. Unlike resistance or capacitance, inductance has no simple closed-form models. For advance CMOS technologies, the wire inductance can be ignored and the wire can be treated  $RC$  with a good accuracy [156,157]. According to [158], if the length of the wire falls within a certain range, Equation 6.3, the inductance effects are significant. This range depends on the parasitic impedances of the interconnect per unit length as well as the rise time of the signal at the input of the CMOS circuit driving the interconnect.

$$\frac{t_r}{2\sqrt{LC}} < l < \frac{2}{R}\sqrt{\frac{L}{C}}. \quad (6.3)$$

The upper bound of the range represents the case in which the inductance is not important because of the large transition time of the input signal. On the other hand, the lower bound shows the case where the wire attenuation is so large that the inductance becomes unimportant. Figure 6.5 shows the region in which the wire inductance becomes important for 28nm CMOS technology.

## 6.2 Wire Performance Metrics

The discussion of wire characteristics above provides the groundwork for an examination of wire performance. In this section we will go over the performance metrics of on-chip wires.






Figure 6.5: Transition time ( $t_r$ ) versus the length of the interconnect line ( $l$ ). The crosshatched area denotes the region where inductance is important.

Interconnect metrics discussed here are mostly applicable to any physical type of interconnect. However, some properties are more relevant to metal-based interconnect systems. In this section we will discuss three important figure of metrics, delay, power, and crosstalk.

### 6.2.1 Delay

The delay of the wires can be well approximated by the product of resistance ( $R$ ) and the capacitance of the wire ( $C$ ), if inductive effects are not important. The delay of longer wires is more accurately modeled by the  $RC$  product because this constitutive component is proportional to square of the length, whereas, for very short lengths, the wire delay is better approximated by purely the loss-less inductive delay formula (length divided by the speed of light in the medium) [159] and only increases linearly with length. Because the wire resistance is increasing very rapidly compared to the inductance, wire delays are becoming more and more  $RC$  limited even at short distances. To model the  $RC$  delay of the wires it is necessary to accurately model both the resistance and the capacitance. The analysis in the last section provides a powerful tool to investigate on-chip wire delay.




Figure 6.6: Simple lumped model for an RC dominated wire.

Interconnects increase circuit delay for two reasons. First, the wire capacitance adds loading to each driving gate. Second, long wires have significant resistance that contributes distributed  $RC$  delay or flight time. The distributed resistance and capacitance of a wire can be approximated with a number of lumped elements, as shown in Figure 6.6. Three standard approximations are the L-model,  $\pi$ -model, and T-model. The L-model is a very poor choice as a large number of segments are required for accurate results. On the other hand, the T- and  $\pi$ -models provide much better accuracy with a small number of segments [160].

Using the above model and applying the Elmore delay approximation [161], one can calculate the total delay of the wire. According to the Elmore model the delay of an  $RC$  ladder, Figure 6.6, is the sum over each node in the ladder of the resistance  $R_{n-i}$  between that node and the signal source multiplied by the capacitance on the node.

$$t_d = \sum_{i=1}^n R_{n-i} C_i, \quad (6.4)$$

where

$$R_{n-i} = \sum_{j=1}^i R_j, \quad (6.5)$$

is the resistance between the source and  $i^{\text{th}}$  node in the ladder, and  $t_d$  is the propagation delay. Considering this model for a wire with length  $l$ , capacitance per length of  $C_w$ , and resistance per length of  $R_w$ , we can calculate the total propagation delay as

$$t_d = R_w C_w l^2 \frac{N+1}{2N} \approx \frac{1}{2} R_w C_w l^2, \quad (6.6)$$

	Copper wire delay (FO-4/mm)
Local wire	136
Semi-global wire	24
Global wire	1.3

Table 6.1: PERFORMANCE OF ON-CHIP WIRES IN 28nm TECHNOLOGY.

where  $N$  is the number of the L-model segments. As  $N$  becomes large the wire delay approaches to half of its time constant ( $RC$ ). Another point that can be noticed from Equation 6.6 is that the delay of the wire grows quadratically with its length. This causes a huge latency in long wires. To remedy this problem inverter-based repeaters are employed, which will be explained in more detail in the next section. Table 6.1 shows the delay of wires per unit length (mm) relative to FO-4 gate delay for different metal layers in a 28nm CMOS technology.

### 6.2.2 Crosstalk

Coupling noise can greatly affect signal integrity in on-chip interconnects, as both mutual capacitance and inductance terms for wires can be large. Capacitive noise coupling usually has a larger effect, therefore we will investigate it first. The large aspect ratios of modern wires cause a significant coupling capacitance between neighboring wires. In particular, for minimum pitch spaced wires, the sideways cap can exceed 70% of the total wire capacitance. Many recent papers have modeled this noise carefully, and have shown that the noise voltage depends on the coupling capacitance to total capacitance ratio as well as on the ratio of the strengths of the gates driving the two wires [162–164].

Noise from inductive coupling can also present problems for VLSI wires. The current flowing in the aggressor wire generates a magnetic field which causes a return current to flow in the victim wire. Inductive coupling pushes the victim in the opposite direction to capacitive coupling: a rising edge on the aggressor wire drives the the victim up through capacitive coupling, while the same edge causes a negative glitch on the victim through inductive coupling. Capacitive coupling usually affects the nearest neighbor, while inductive coupling has a much larger range. Inductive noise becomes a problem only when a large

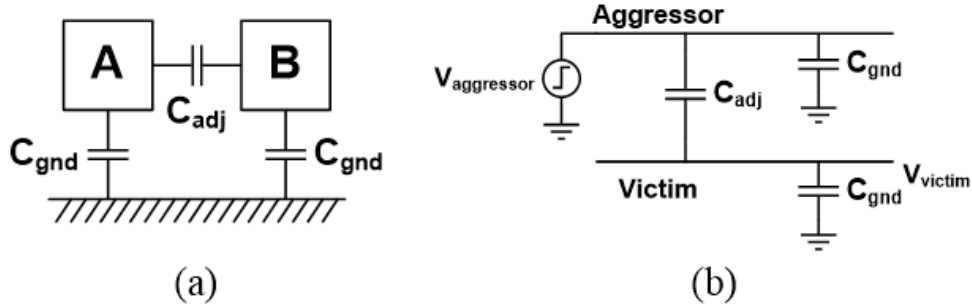


Figure 6.7: Simple model for evaluating capacitive coupling.

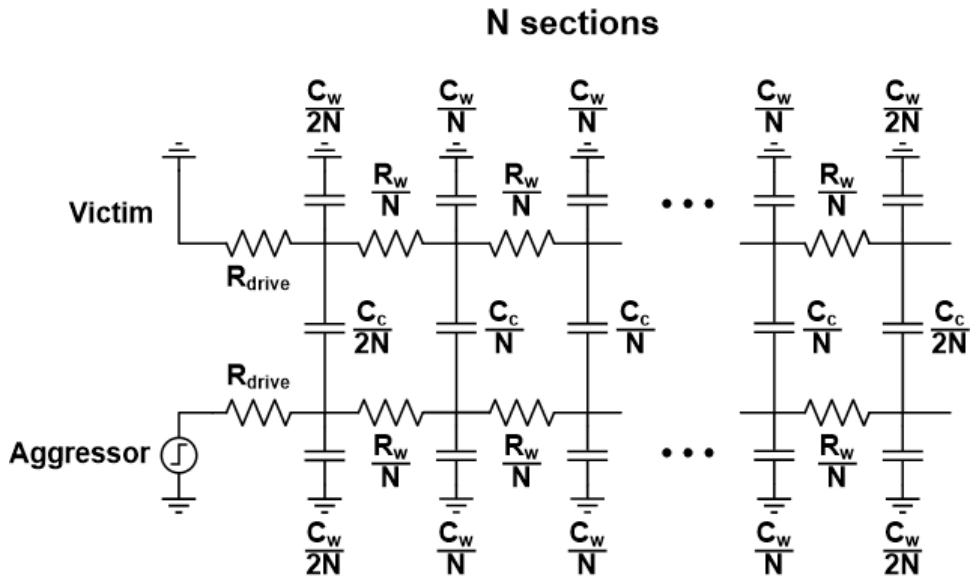


Figure 6.8: Capacitive coupling model in a wire with considerable resistive loss.

number of wires switch at the same time in bus-like situations [165, 166]. In the worst case multiple wires are switching, with near neighbors switching in one direction and far neighbors switching in the opposite direction, which results in constructive capacitive and inductive coupling noise from near and far neighbors, respectively. Therefore, the capacitive and inductive noises add, and the accumulated noise can be enough to cause signal detection failures.

There are different ways to reduce capacitive coupling. The simplest way is to increase the spacing between wires and hence reduce coupling capacitance. As will be explained later in this chapter, usually designers employ buffers (also known as repeaters) along the wire to improve its performance. Moving the repeaters in a bus such that each bit's repeaters




Figure 6.9: Techniques to combat capacitive coupling in on-chip wires.

are staggered from its neighbors forces capacitive noise to cancel itself. The structure of such a bus is illustrated in Figure 6.9(a). Because half of the injected noise must propagate down the RC wire to negate the other half, this cancellation is not perfect, but still effective [155]. Another technique to cancel capacitive coupling is charge compensation, in which a physical capacitor is introduced between the coupled lines to inject reverse noise to counteract parasitic coupling [167]. This requires extra area and power but can minimize noise as well as reduce data-dependent delay variation. Moreover, choosing the proper size for the capacitor requires careful modeling and simulation and could be susceptible to process and temperature variation. This technique is illustrated in Figure 6.9(b).

One of the most effective methods of reducing capacitive coupling noise is to employ differential signaling. In a differential system the difference is sensed at the receiver and hence any noise that affects the two signals in a same way appears as a common-mode interference and gets canceled. In order to effectively remove capacitive coupling in a differential scheme, the wires should be also twisted periodically, as shown in Figure 6.10. As a result, injected noise affects both wires equally and hence the differential voltage is unchanged. In addition, these systems have minimal inductive coupling to the rest of the system, because each wire

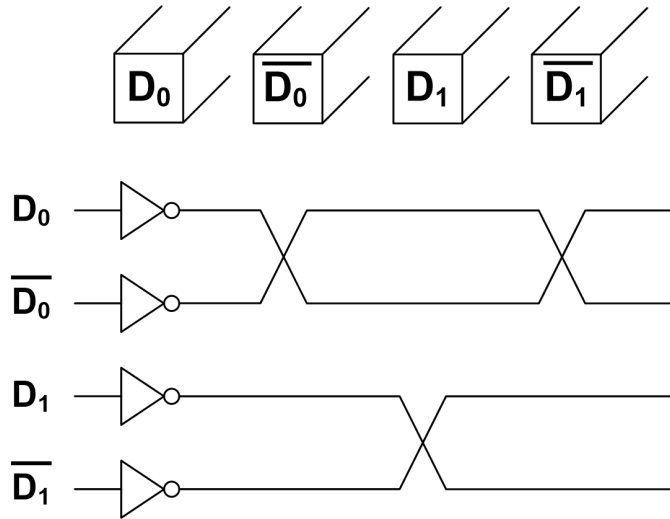


Figure 6.10: Differential signaling along with wire twisting to remove crosstalk.

acts as the return path for the other, creating the smallest possible current loops. The main problem associated with this technique is the extra dynamic power. As the driver has to drive twice the capacitance as in a single-ended version, the total dynamic power is increased by at least a factor of two. In addition, driving the wires differentially causes a Miller factor of two which further increases the power consumption. The twisting scheme also requires jumping to other metal layers which imposes area over-head. Nevertheless, with differential and twisted bits, wires can easily reject noise even if the coupling ratio approaches 90 to 100% [155].

Another effective technique in reducing the capacitive coupling is to insert ground shields between the data wires [168], as shown in Figure 6.11. The advantage of this technique compared with the differential signaling is the lower dynamic power, which is associated with the lower total driving capacitance (this includes a lower Miller factor of 1). It also offers a better area efficiency as no twisting is required. The need for a custom differential receiver is rectified as a simple inverter can perform the data resolution. With proper shielding a maximum coupling noise of less than 5% can be achieved for a minimum-pitch set of wires [169].

Designers cope with inductive coupling by adding power planes or densely gridded power supplies to reduce the number of wires that can couple into a victim. Power planes, or dense




Figure 6.11: Ground shield insertion to avoid croostalk.

power grids, effectively reduce both self and mutual inductances for wires in the direction of the grid, because they provide very nice return paths within a few microns of the wire itself and thus limit the extent of the magnetic coupling [170]. Most companies have design rules for buses, such as requiring every fifth wire to be a power supply wire, which makes inductive noise much less than capacitive noise and under 5% of the power supply.

### 6.2.3 Power

The second metric of importance is the power dissipation due to interconnects. This is a result of charging and discharging of the wire capacitance and is given by the dynamic power dissipation formula

$$P_{wire} = S_w C_w V^2 f, \quad (6.7)$$

where  $P_{wire}$  is the wire power per unit length,  $S_w$  is the switching activity factor representing the probability of a particular interconnect switching during a clock cycle,  $C_w$  is the wire capacitance per unit length,  $V$  is the voltage to which the interconnect charges and  $f$  is the clock frequency. Thus, at a given technology node, the interconnect power is heavily dependent on its total capacitance.

## 6.3 Repeaters

As mentioned earlier, the delay of an uninterrupted wire grows quadratically with wire length. This makes it difficult to communicate globally between different parts on a chip without encountering timing issues. Designers have employed different techniques to resolve this problem. Multi-core processors help to reduce the number global interconnects and




Figure 6.12: Inserting repeaters to improve on-chip wire delay.

hence partly mitigate the wire latency issue. However, the cores still need to communicate in order to fully exploit the parallel processing capability offered by these architectures.

The most popular design approach to reducing the propagation delay of a long wire is to break it into smaller pieces and introduce intermediate buffers (also known as repeaters) along the wire, as shown in Figure 6.12. When added in a way to optimize delay, repeaters make the total wire delay equal to the geometric mean of the total wire delay and the individual repeater stage delay. Hence, the length-squared term in wire delay falls out of the square root, making the total delay linear with total wire length [172]. Now let us see how optimal repeater insertion can be performed. By breaking the wire into short pieces, with a high accuracy we can model each piece with a simple  $\pi$ -model, as shown in Figure 6.13. Let's consider the capacitance and resistance per unit length of the wire to be  $C_w$  and  $R_w$ , respectively. Repeaters can be modeled with their input,  $C_g$  and output capacitances per unit width,  $C_T$ , as well as their equivalent resistance per unit width,  $R_T$ , [173]. Assuming equal size repeaters, the total delay of a wire with length  $L$  can be written as

$$T_d = R_T(C_T + C_g)N + \frac{R_w C_w L}{W} + \frac{R_w C_w L^2}{2N} + R_w C_g W L, \quad (6.8)$$

where  $W$  is the width of the repeaters. Setting  $\partial T_d / \partial N$  and  $\partial T_d / \partial W$  yields optimal values for  $N$ ,  $W$ , and  $l$  to minimize the wire delay as shown in equations 6.9, 6.10, and 6.11, respectively.



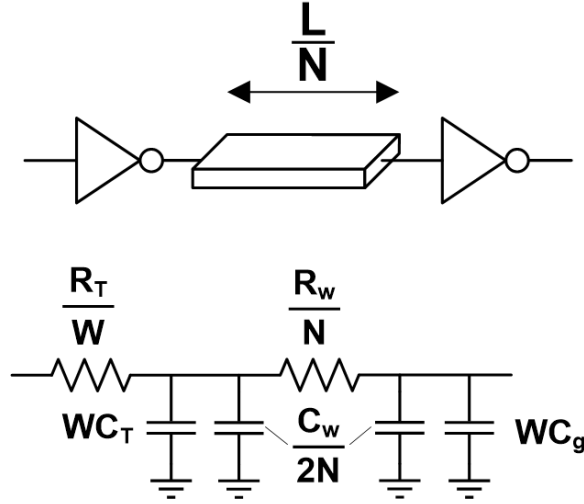


Figure 6.13: A segment of a repeated wire represented by a  $\pi$ -model.

$$N_{opt} = \sqrt{\frac{R_w C_w}{2R_T(C_g + C_T)}} \times L, \quad (6.9)$$

$$W_{opt} = \sqrt{\frac{R_T C_w}{R_w C_g}} \times L, \quad (6.10)$$

$$l_{opt} = \frac{L}{N_{opt}} = \sqrt{2R_T C_g (1 + \gamma) R_w C_w}, \quad (6.11)$$

where  $\gamma$  represents the ratio between the output and input capacitance of the repeater. This results in a total wire delay of

$$T_{d,opt} = 2\sqrt{R_w C_w R_T C_g (1 + \sqrt{0.5(1 + \gamma)})} \times L. \quad (6.12)$$

As a result, the delay of a repeated wire can be written in terms of  $l' = l/l_{opt}$  and  $w' = w/w_{opt}$  as

$$T_d = \sqrt{R_w C_w R_T C_g} \left( \left( l' + \frac{1}{l'} \right) \sqrt{0.5(1 + \gamma)} + \left( W' + \frac{1}{W'} \right) \right) \times L. \quad (6.13)$$

Although using repeaters is an attractive solution for the delay problem of long wires,

they add some design complexity. First, the simplest repeaters are inverting elements, so an even number of repeaters is necessary to maintain logic levels. Second, repeaters for global wires require many via cuts from the upper-layer wires all the way down to the substrate, potentially congesting routes on intervening layers. Third, designers are rarely afforded the luxury of placing repeaters in their optimal locations, because they require active area; designers usually floor-plan repeaters in pre-planned clusters. Finally, even with delay-power optimizations, repeaters are still large devices, and repeating an entire bus takes an impressive amount of silicon area. Fortunately for these last two complications, delay and capacitance curves for repeater insertion have fairly shallow optimizations, so that adding or removing a single repeater stage, moving repeaters back and forth, or resizing repeaters have fairly small costs.

As mentioned earlier, using repeaters directly adds to the total interconnect power. If we assume that dynamic charging and discharging of the capacitors is the only contributing factor to the interconnect power, the total power consumption for an optimally repeated link can be written as

$$P = K(N_{opt}(C_g + C_T)W_{opt} + C_w L)V_{dd}^2 f. \quad (6.14)$$

Replacing  $N_{opt}$  and  $W_{opt}$  from equations 6.9 and 6.10 in 6.14 results in

$$P = (1 + \sqrt{\frac{1 + \gamma}{2}})C_w K L V_{dd}^2 f, \quad (6.15)$$

where  $K$  is the logic activity factor. If we assume  $C_g = C_T$ , then the capacitance associated with the repeaters will be equal to the line capacitance and hence the total link power increases by 100% compared with an un-repeated wire. It should be noted that the factor of two increase in power happens in the case of an optimally repeated wire. Therefore it is worthwhile to investigate the effect of deviating from the optimal repeater design in favor of the power consumption. Figure 6.14(a) illustrates the variation of the repeated wire delay normalized to the optimal value. It can be seen that the graph is quite flat, which means that by slightly sacrificing the delay it is possible to obtain a large power saving. For instance,

reducing the number of repeaters to 80% and the repeater size to 70%, we can achieve about 44% reduction in the power associated with the repeaters while sacrificing the delay by 5%.


From the above discussion, it can be concluded that the delay-optimal solution is inefficient. Delay optimization naturally causes large power and energy costs in order to gain marginal delay benefits. In particular, at the optimal solution,  $N_{opt}$  and  $W_{opt}$ , the shallowness of the contours indicates only minor performance gains for fairly large changes in  $N_{opt}$  and  $W_{opt}$ . From a design perspective, the delay optimal design is wasteful. In fact by giving up small increments of delay performance a significant power saving can be achieved. In order to quantify gains in efficiency, we next construct models for repeater energy and optimize for energy-delay product.

While a thorough energy analysis would not only consider switched dynamic capacitive current but also device leakage and short circuit current associated with the repeaters, this discussion will focus primarily on switched capacitances and ignore leakage currents and short circuit current. A detailed analysis of the repeated wire delay-power performance can be found in [174]. Using equation 6.15 we can rewrite the total link power consumption as


$$P_{tot} = \left(\frac{W'}{l'} \sqrt{\frac{1+\gamma}{2}} + 1\right) KLC_w V_{dd}^2 f. \quad (6.16)$$

Combining this result with the delay expression in Equation 6.13 gives us the contours in Figure 6.15. The straight lines illustrate the power contours normalized to the optimal delay design. As can be seen, longer wire segments and smaller drivers achieve better power efficiency with minimal delay penalty.

A more intelligent approach to designing repeaters is to combine the expression for the wire delay and power consumption and try to minimize their product. The contours shown in Figure 6.16 illustrate the optimal  $l'$  and  $W'$  for product of power and delay. This graph illustrates the benefits of longer wire segments and smaller drivers to achieve low power consumption. The minimum energy-delay product occurs when  $l' = 1.333$  and  $W' = 0.667$ , which results in a delay that is 6% worse than that of the optimally repeated wire, and a power of 12.5% lower than that at the delay-optimal.



(a)



(b)

Figure 6.14: Delay of the repeated wire normalized to the optimal delay (a). The constant delay contours as a function of the repeater normalized width and wire segments relative to optimal values.




Figure 6.15: Constant power contours along with delay contours illustrating the trade-off between power and delay.




Figure 6.16: Constant power-delay product contours.




Figure 6.17: Eye diagram when 10-to-90% rise time equal to the bit time.

Another consideration in the design of a repeated wire is the rise and fall time of the signal in each repeated section. Signal integrity at the input of each repeater can be greatly degraded if the rise and fall time of the repeated wire section is larger than the bit time. As a result, in order to preserve signal integrity while traveling to the receiver side, the repeater strength and the wire segment length should be chosen such that the resulting rise and fall time at the end of each segment is equal to or smaller than the bit time. Using the dominant time-constant model, as in Equation 6.8, we can calculate the 10-to-90% rise time as

$$t_{rise} = 2.2 \times R_T C_g l' \left( \left( l' + \frac{1}{l'} \right) (1 + \gamma) + \left( W' + \frac{1}{W'} \right) \sqrt{2(1 + \gamma)} \right). \quad (6.17)$$

This equation quantifies the rise time for each segment of the wire. To calculate the overall rise time of the repeated wire, we need to take into account the effect of all the segments. Considering a system composed of  $n$  cascaded non-interacting blocks, each having a rise time  $t_{ri}$  and no overshoot in their step response, it can be shown that the output signal has a rise time equal to [175]

$$t_{rise} = \sqrt{t_{rs}^2 + t_{r1}^2 + \dots + t_{rn}^2}, \quad (6.18)$$

where  $t_{rs}$  is the rise time of the input signal. Combining equations 6.17 and 6.18 results in the total repeated wire rise time of




Figure 6.18: Constant rise time contours for a repeated wire.

$$t_{rise} = 2.2 \times R_T C_g l' \left( (l' + \frac{1}{l'}) (1 + \gamma) + (W' + \frac{1}{W'}) \sqrt{2(1 + \gamma)} \right) \sqrt{\frac{L}{l' \times l_{opt}}}. \quad (6.19)$$

$l_{opt}$  can be calculated using equation 6.11. In order to avoid ISI the rise time of the received signal should be almost equal to the bit time. Figure 6.17 shows the case in which the 10-to-90% rise time is equal to the bit-time. As can be seen in this figure, a rise time equal to the bit time results in a reasonable eye-opening at the end of each wire segment so that the following inverter can resolve the data. It should be noted that the constraint imposed on the repeater strength and wire segment length by the rise time might not necessarily agree with the optimal power-delay condition. Figure 6.18 illustrates the rise time contours with respect to the wire segment length and driver size normalized to the delay optimal values.

Figure 6.19 shows the power contours along with the rise time contours. It can be seen that, reducing power by increasing the wire segment length and decreasing driver size results in larger rise time and hence lower data rate.




Figure 6.19: Constant rise time along with power contours.

## 6.4 Wire Scaling

To properly discuss wire performance under technology scaling, we will first make a distinction between two kinds of wires, as shown in Figure 6.20. The first kind of wire connects gates locally within cores. For instance, a wire that runs across a block of 1000 gates. As technology scales by  $\alpha$ , these gates become smaller in area by  $\alpha^2$ , and hence this wire becomes shorter by  $\alpha$ . In general, local wires become shorter due to technology scaling.

The second kind of wire spans between cores and communicates data between them. Under scaling, die size does not typically decrease, we simply add more functionality to the chip, so these global wires that span across the chip, as shown in Figure 6.20, have the characteristic of remaining fixed in length.

Figure 6.21 shows unrepeated wire delays for both local and global wires. In this graph, local and global wires span 1mm. Along the x-axis are CMOS technologies for each year, according to ITRS 2011 data [4], plotted on log-scale to make them linear in time, and on the y-axis are wire delays, normalized to FO-4 inverter delays. In both graphs we show two pairs of trends, for intermediate layer (semi-global wires) and upper layer metals (global wires). Over this range of technology assumptions the conclusions are still consistent.



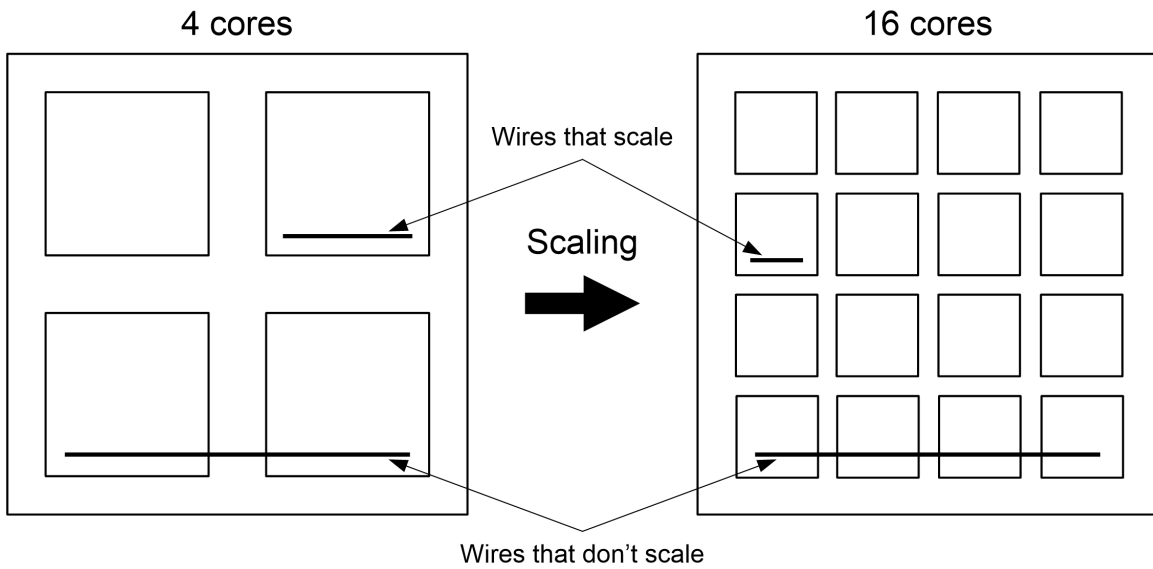


Figure 6.20: Two kinds of wire on a chip: local and global




Figure 6.21: Unrepeated global and semi-global wires delay normalized to FO-4 inverter delay for different technologies.




Figure 6.22: Repeated global and semi-global wires delay normalized to FO-4 inverter delay for different technologies.

Wire resistance grows quickly under scaling, leading to wire delays that increasingly lag gate delays. As can be seen in Figure 6.21, local wire delays degrade by more than 100X over the 15 year ITRS projection. Global wire delays follow almost the same trend, degrading over the same time span by more than 200X. Although these trends of wire performance paint a dismal picture of scaled designs, designers can use repeaters to improve wires. When applied to global wires, repeaters require a proportionally small area and design overhead and make these fixed-length wire delays, relative to gates, grow only by a factor of 20X over the 15 year time span, as shown in Figure 6.22.

## 6.5 Summary

This chapter examined how wire performance will scale as technologies advance. We first discussed a gate delay model, because wire delays are important only if they change relative to gate delays. We used the delay of a fanout-of-four inverter and modeled its speedup over time as linear with technology. Next, we discussed the two wire characteristics that are important to delay, resistance and capacitance, and discussed why inductance is not important. Using very simple yet sufficiently sophisticated geometric models for resistance

and capacitance, we constructed delay metrics and saw how they scaled with technology. On one hand, scaled-length, or local, wires keep up with gate delays when repeated; on the other hand, fixed-length, or global, wires cannot. Local wires get worse relative to gates by less than 2-3X over nine generations of technology scaling, while global wires degrade by 40-50X. This bifurcated view of wires leads to a number of broad implications across VLSI design, and we shall consider a few of these effects in the next chapter.

# Chapter 7

## On-Chip Interconnects

As VLSI technologies and multi-core processor chips continue to scale, long on-chip wires will present increasing performance limitations. While transistors benefit from technology scaling, the shrinking cross-sectional area of the on-chip wires increases electrical resistance and hence their latency, which has a quadratic relation with the wire length. Simple inverter-based repeaters can partially mitigate the latency problem, where an optimal design makes the repeated wire delay linear with length instead of quadratic. However, the associated power and area become prohibitive as the technology scales due to the increased number of repeaters per unit length.

In this chapter, we briefly go over the shortcomings of on-chip wires and solutions presented in prior works. Next we introduce our proposed technique which is inspired from the RC front-end optical receiver explained in Chapter 5. In the proposed on-chip link we take advantage of the RC dominant behavior of the minimum-pitch wire and apply the double-sampling technique along with a capacitively-driven transmitter to achieve high data rate and power efficiency. The chapter concludes with circuit implementation and experimental results.

### 7.1 On-Chip Communication Power Trend

Power dissipation of high-performance microprocessors is becoming a limiting factor and hence design for efficient power consumption is becoming a major consideration. Dynamic




Figure 7.1: Repeated and un-repeated wire delay variation trend with CMOS technology scaling.

power is currently the main component of the power dissipation. Under ideal scaling, all dimensions of the wires are shrunk 0.7X per generation [176]. As a result, the wire resistance per unit length doubles every process generation, while the wire capacitance remains almost constant, resulting in a wire delay degradation per scaled micron of 1.4X every generation. Figure 7.1 shows the wire delay scaling under ITRS technology projection [4].

As explained in the previous chapter, the RC-delay of an on-chip wire grows quadratically with wire length, therefore repeaters have traditionally been used to linearize the dependence of delay on interconnect length. In an optimally repeated interconnect, the delay of any given stage is approximately equally divided between the repeater and the wire. As the cross-sectional area of the wires shrink due to technology scaling, the repeated wires also illustrate an increasing trend in their latency, as shown in Figure 7.1. As expected, repeated wires have much smaller latency compared to un-repeated wires. This in fact comes at the cost of extra power and area associated with the repeaters. Recall from the previous chapter the number of repeaters in an optimally repeated wire can be expressed as

$$N_{opt} = \sqrt{\frac{R_w C_w}{2R_T(C_g + C_T)}} \times L. \quad (7.1)$$




Figure 7.2: The repeater distance and number of repeaters for an optimally repeated wire in different technology nodes.

As the technology scales,  $R_w$  increases while  $R_T$  shrinks, which results in a significant increase in the number of repeaters. Figure 7.2 illustrates how the number of repeaters in an optimally repeated wire increases as the technology scales. This imposes a large area overhead due to the increasing number of repeaters. As shown in previous chapter, the total power associated with the repeated wire can be approximated by the dynamic power required to drive a capacitor with twice the size of the wire capacitance. Wire capacitance remains almost constant with technology scaling; however, as the minimum wire pitch scales, the total number of wires within a fixed area increases accordingly. As a result, assuming that the clock frequency remains almost constant (which is the case in new processors), the total dynamic power dissipation associated with repeaters increase almost linearly with technology scaling, as shown in Figure 7.3. It should be noted that the short circuit and leakage power of the repeaters are not considered in this analysis. Therefore, the situation is even worse.

In microprocessors, the interconnect power analysis reveals that the interconnections are responsible for over 50% of the dynamic power consumption [178], as shown in Figure 7.4. Diffusion and gate power represent the total dynamic power dissipation in the diffusion and




Figure 7.3: Projection of the repeated wire power consumption in different technology nodes.

gate capacitances of the transistors. The interconnect power in this graph includes both data communication and clock distribution. With the advent of multiple core processors, this problem has been mitigated to some extent as integration of multiple cores on a chip allows for lower interconnect latency and therefore higher bandwidth between cores than their discrete counterparts [177]. Nevertheless, the interconnect power is a major portion of the microprocessors chips and as shown in Figure 7.3, it is growing with technology scaling.

## 7.2 Prior Art in Design of On-Chip Links

As explained in the previous chapter the most common approach to mitigate the high latency of the RC-limited on-chip wires is to employ repeaters. However, as the technology scales, the power and area overhead of the repeaters become prohibitive, as shown in the previous section. Designers have proposed different techniques to meet these challenges such as low-swing differential signaling (LVDS) [179, 182–185, 191], current-mode signaling [186, 187], equalization [185, 191] and transmission lines [179, 186, 190]. In this section we introduce these techniques in more detail and illustrate how they are becoming less adequate in meeting

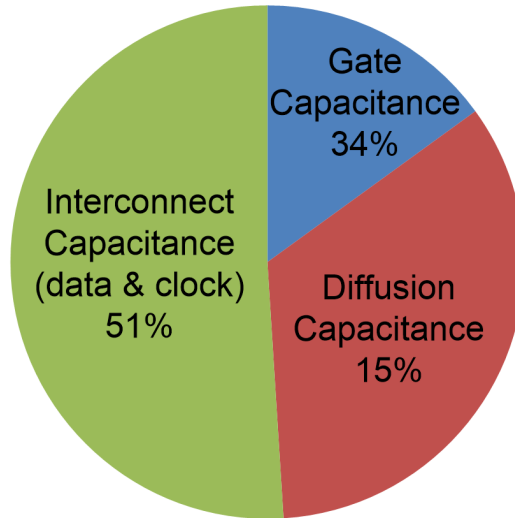


Figure 7.4: Dynamic power breakdown for a single core processor [178].

bandwidth density and power requirements.

### 7.2.1 Low Voltage Signaling

In RC-dominated on-chip links, the transmitter typically consumes most of the link power driving a large wire capacitance using a rail-to-rail swing buffer. As discussed in the previous chapter, this power is quadratically proportional to the supply voltage of the driver. As a result, a natural way to reduce the transmitter power consumption is to decrease the voltage swing of the driver. Zhang in [190] has extensively analyzed the effectiveness of low-swing signaling in power reduction of on-chip interconnects. In order to achieve this goal, a dedicated supply voltage can be employed for the transmitter. This comes at a performance cost, as driver resistance grows as supply voltage reduces and hence the latency is increased. In addition, it is not desirable to have different supply voltages on-chip as it adds complexity to the power distribution network. Designers have proposed different techniques to avoid an extra supply in low-swing signaling applications. In [183, 185] a novel capacitively coupled driver is proposed, which not only enables low-swing signaling, but it also enhances the overall bandwidth of the wire through high frequency pre-emphasis. The basic idea behind






Figure 7.5: Charge-recycling stacked transmitter employed to reduce effective supply voltage.

this technique is that driving a long wire of capacitance  $C_w$  through a capacitor,  $C_c$  reduces the signal swing on the wire through a capacitive voltage divider. The final load swing,  $V_{swing}$  in this configuration will be equal to  $C_c/(C_c + C_w + C_L)$ , where  $C_L$  represents the load capacitance. Usually,  $C_L$  is much smaller than the wire capacitance, therefore, in order to achieve 10X reduction in the output voltage swing the wire capacitance should be about 9X larger than the coupling capacitor.

Another advantage of using the coupling capacitor is that because it acts as a high-frequency short, the 3-dB bandwidth of the long wire increases. This allows for shorter cycle times and decreasing latency. Long wires have a low-pass frequency response that limits their operating speed due to excessive ISI. Capacitive coupling to a long wire creates a pole-zero pair that provides high-frequency emphasis that mitigates the low-pass wire response and increases performance [183, 185]. For instance, in the aforementioned case where 10X reduction was achieved, the bandwidth improves by more than a factor of three.

Another approach to implement reduced swing signaling is to stack up drivers with a single power supply so that the supply voltage is equally divided between the two drivers [191]. Figure 7.5 illustrates the implementation of this technique. A DC-DC converter is required to guarantee  $V_{dd}/2$  voltage at the mid-point. This DC-DC converter employs a large capacitance to provide enough output current without deteriorating from the nominal output voltage. In [191], this large capacitor is implemented using deep-trench capacitors in an SOI technology.


### 7.2.2 Current-Mode Drivers

In a current-mode driver, the data signal is transmitted using a current which is resolved at the receiver [187–189]. Current-mode drivers usually employ differential signaling. The benefit of the current-mode signaling technique is that the wire capacitance does not charge and discharge to the supply rails and hence potential power saving can be achieved. The key to current-mode signal transporting is the extension of signaling bandwidth and reduction of the system time constants that result from sensing signals with low impedance nodes. Figure 7.6(a) illustrates a simple current-mode driver, which employs a differential pair driven by the complementary data signals. The static power consumption in this configuration sacrifices the power saving achieved by the current-mode signaling. In order to avoid static power, the driver shown in Figure 7.6(b) can be adopted. In this configuration the complementary transistors are utilized to prevent a direct path between supply and ground and hence eliminate the static power dissipation.


One way to implement the current-mode receiver is to employ a TIA, which provides a large transimpedance and at the same time a low input resistance which helps improve the wire bandwidth [187]. Another approach is to use a current-mode sense amplifier at the receiver. This technique has been recently proposed in [189]. As shown in Figure 7.7, in this technique an open drain differential driver is employed at the transmitter along with a pre-emphasis equalization tap to remove ISI. The receiver utilizes a current-mode sense amplifier to convert the received current into a voltage which is consequently resolved by a comparator. In this technique the wire is essentially embedded into a sense amplifier which is composed of the transmitter differential pair and the receiver current-mode sense amplifier. This technique offers high power efficiency, however its operation is limited to 3-4Gb/s.

### 7.2.3 On-Chip Transmission Line

Electromagnetic wave propagation in on-chip transmission line structures is desirable because the peak phase velocity is determined by the speed of light in the dielectric surrounding the interconnect. Unlike distributed RC lines, where signals travel slowly by diffusion, the LC



(a)



(b)

Figure 7.6: Current-mode on-chip wire drivers.




Figure 7.7: Sense amplifier based current-mode on-chip wire driver.

nature of an optimized transmission line structure allows for faster signal transmission [180].

Even though this technique achieves the best latency, implementing an on-chip transmission line requires significant area, as large width is required to guarantee the operation of the transmission line in the LC regime [179, 180]. Furthermore, usually the transmission line is implemented with the top most metal as the signal line and the lowest metal as the ground line. This severely limits the usage of intermediate metal layers in the transmission line vicinity.

## 7.2.4 Equalization

Due to the RC nature of the on-chip wires, binary signals suffer from a long train of post-cursor inter-symbol interference (ISI). A common approach, which has been extensively employed by designers, is to employ transmit pre-emphasis through feed-forward equalization [181–183]. As explained in Chapter 2, in this technique, a delayed version of the signal is employed to reduce the low-frequency content of the transmitted signal and compensate for the channel loss.

To eliminate ISI, equalization techniques such as decision feedback equalization (DFE) can be utilized, but the long post-cursor tail necessitates many DFE taps, which results in significant power overhead. This problem is exacerbated as the technology scales. RC signal emulation in a DFE is also an attractive solution to eliminate many taps of post-cursors [185, 191]. The main limiting factor in this technique is to meet the timing requirement in the feedback loop, especially at high data rates.

In this chapter, we present an on-chip link using minimum-pitch wires for high-speed signaling to address the bandwidth requirement of future microprocessors. We propose a double-sampling technique with a feed-forward dynamic offset modulation (DOM) to achieve high data rates over minimum-pitch and long on-chip wires that suffer from excessive loss and latency. In order to further improve data rate and reduce power consumption, a capacitively-driven transmitter is employed [183]. The emphasis of the proposed design is low power consumption, high bandwidth density and scalability to future technology nodes.

## 7.3 Double-Sampling Link

It was shown in Chapter 5 that an RC dominated optical front-end can employ a double-sampling technique along with dynamic offset modulation to receive data without suffering from ISI. In this chapter we will demonstrate that a similar technique can be utilized to achieve a high data rate over minimum-pitch on-chip wires in a very power efficient manner. The proposed receiver and transmitter for on-chip signaling will be explained in the next two sections.

### 7.3.1 Receiver Design

Minimum-pitch wires can have a slow exponential response to a fast transition, with a time-constant ( $\tau$ ) much larger than the bit time ( $T$ ). Instead of conventional equalization techniques, in this work we propose to employ a mostly-digital double-sampling technique to break the trade-off between the data-rate and the on-chip wire time-constant [193]. Figure 7.8 shows the top-level architecture of the proposed receiver. As shown in Figure 7.8(b), the input voltage is sampled at the end of two consecutive bit times ( $V[n-1]$ ,  $V[n]$ ) and these samples are compared in order to resolve each bit:  $\Delta V[n] = V[n] - V[n-1] > 0$  results in 1, and  $\Delta V[n] < 0$  results in 0. Note that the overall sampling rate is still equal to the data rate. This double-sampled voltage ( $\Delta V[n]$ ) will be input-dependent due to the channel transfer function, Figure 7.9(b). To resolve this problem, the offset of the next stage comparator can be dynamically modulated to provide a constant voltage at its input regardless of the data sequence. In this receiver, an offset proportional to the previous sample,  $V[n-1]$ , is applied to the comparator. For instance, in case of a large  $\Delta V[n] = \Delta V_{max}$  (e.g. a one after many zeros), or a very small  $\Delta V[n]$  (e.g. a zero after many zeros), an offset equal to  $\frac{\Delta V_{max}}{2}$  is subtracted resulting in  $\Delta V'[n] = \frac{\Delta V_{max}}{2}$ , and  $-\frac{\Delta V_{max}}{2}$ , respectively as shown in Figure 7.9(b). The same scenario is true for the opposite case. Figure 7.9(a) shows the z-domain representation of the double-sampling and the dynamic offset modulation technique. Assuming a dominant pole of  $\omega_p = \frac{1}{\tau}$ , it can be shown that for an exponential signal, dynamic offset modulation can eliminate the input dependency of the double-sampled voltage if DOM gain,  $\beta$ , is equal

to

$$\beta = \alpha(1 - e^{-\frac{T}{\tau}}), \quad (7.2)$$

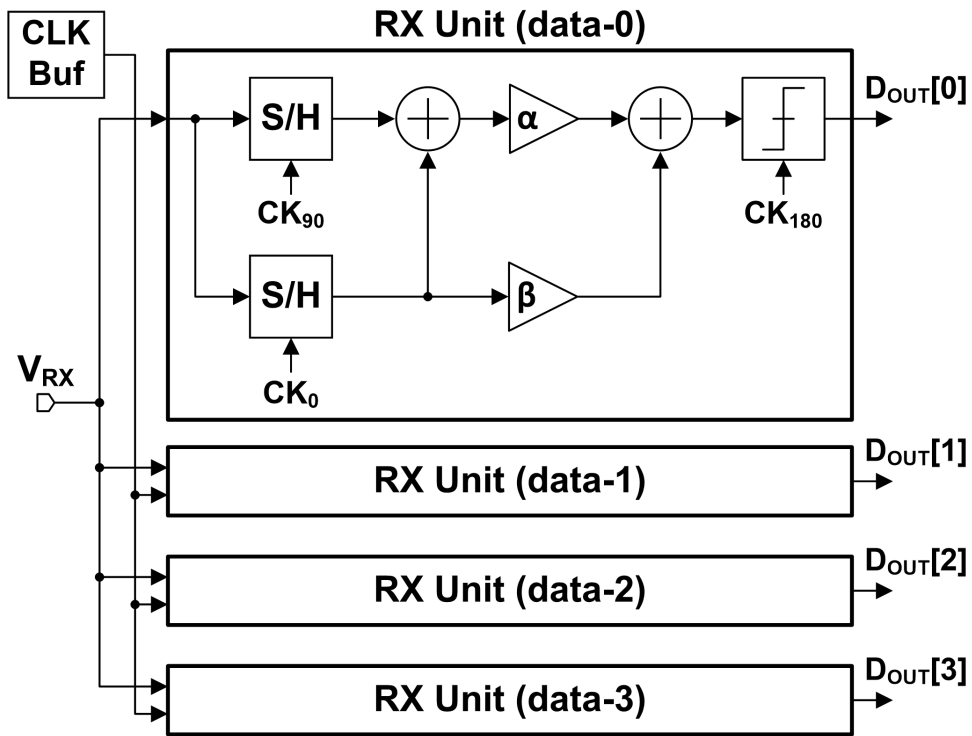
where  $\alpha$  is the main path gain. This results in a constant double-sampled voltage,  $\Delta V'[n]$ , equal to

$$\Delta V'[n] = \frac{\alpha}{2}(1 - e^{-\frac{T}{\tau}}) = \frac{\alpha}{2}\Delta V_{max}. \quad (7.3)$$

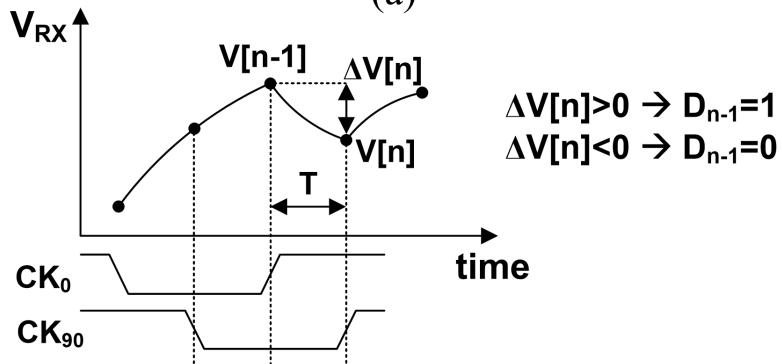
Another advantage of this technique is the capability to perform immediate demultiplexing at the front-end. A quarter-rate architecture (multiplexing factor of four) is employed in this design. As a result, the comparators will operate in a fraction of the data rate.

### 7.3.2 Transmitter Design

Utilizing low-swing signaling also reduces the power consumption in an on-chip interconnect, where most of the power is associated with the dynamic charging and discharging of the wire capacitance ( $C_w$ ). A separate supply can be employed for an inverter-based transmitter to reduce the signal swing and hence improve power efficiency. However, it is not desirable to have multiple supplies on chip, as it makes the power distribution complicated. An alternative approach to achieve low swing is to drive the wire through a capacitor,  $C_p$ . This helps reduce the signal swing on the wire through the use of a capacitive voltage divider. Ignoring the parasitic capacitance associated with the driver and the receiver, the resulting signal swing at the receiver side will be equal to  $\frac{C_w}{C_w + C_p} \times V_{dd}$ . The capacitor also pre-emphasizes transitions and reduces the driver's load. Because it acts as a high-pass filter, the capacitor increases the bandwidth of the wire by almost a factor of  $\frac{C_w}{C_p}$  and decreases latency. Figure 7.10 shows the capacitive driver and the resulting signals at the input and the output of the wire. As shown in Figure 7.10(b), the receiver input demonstrates an exponential behavior.



(a)



(b)

Figure 7.8: Receiver top-level architecture, double-sampling technique and DOM.

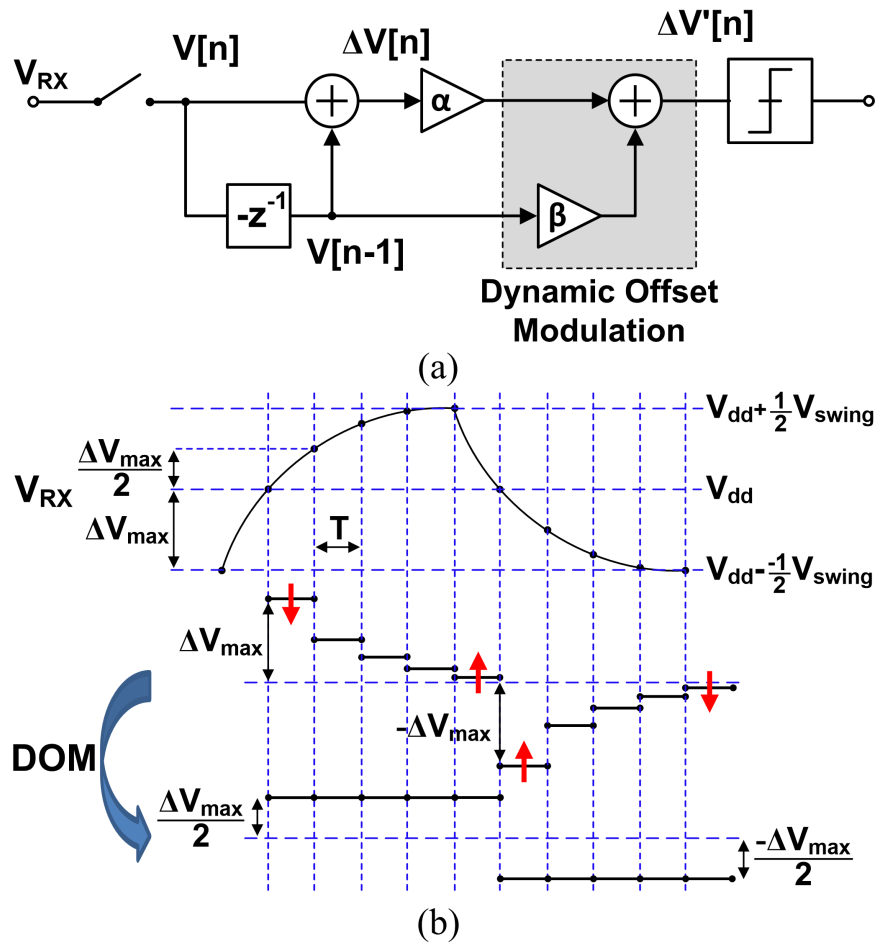


Figure 7.9: Z-domain representation of the double-sampler and the dynamic offset modulation (a). Operation of the dynamic offset modulation (b).




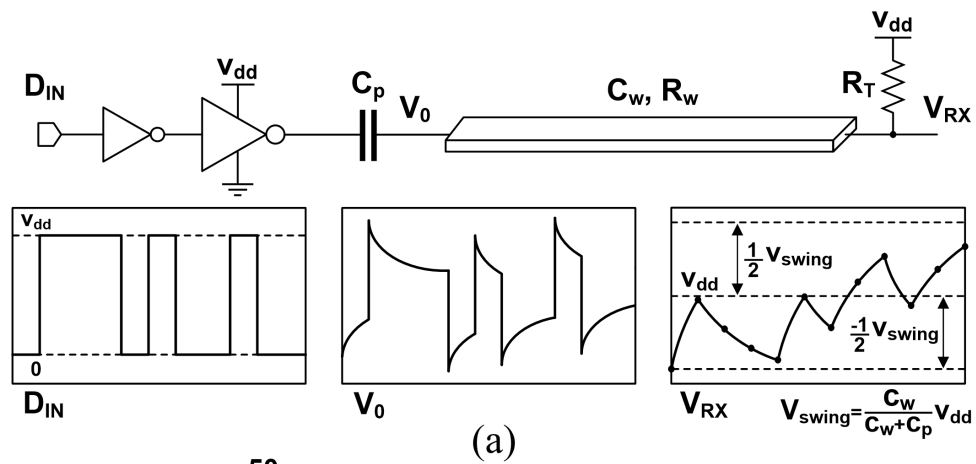


Figure 7.10: Capacitively-driven transmitter (a), double-sampling technique to resolve the received data (b).




Figure 7.11: Frequency characteristics of a minimum-pitch 7mm wire along with the power spectral density of the double-sampled pulse.

## 7.4 Link Latency

The latency of the repeated wires is proportional to their length. In a full-swing repeated wire, the threshold of the receiver inverter determines the effective arrival time of the transmitted signal. This threshold is usually set to  $V_{dd}/2$  to achieve maximum noise margin. As a result, the latency of the wire can be quite accurately approximated by the time it takes for the receiver input to rise to  $V_{dd}/2$ . As the wires become more resistive the rise time increases, which results in a long latency. In other words, the signal propagation velocity is a strong function of the frequency in on-chip wires [180]. For instance, Figure 7.11 illustrates the normalized velocity of a 7mm wire with respect to the speed of light in the same medium. As can be observed, at low frequencies the signal experiences much larger delay than in high frequencies where the signal propagation speed approaches the speed of light. Thus, on-chip wires impose a large latency to signals with low frequency content.

In a double-sampling receiver the incoming signal is resolved based on the difference between the two consecutive samples. Therefore, the low-frequency content of the received signal is filtered by the high-pass behavior of the double-sampler. In other words, this technique detects transitions in the transmitted signal. As transitions contain high frequency




Figure 7.12: Power spectral density of the transmitted pulse and the double-sampled pulse.

content, they propagate along the wire at near the speed of light. In addition, the capacitive-driven transmitter pre-emphasizes signal transitions to further enhance wire latency. To quantify these effects, Figure 7.12 shows the power spectral density of a transmitted pulse and a pulse after being processed by the coupling capacitor and the double-sampler. It can be seen that the double-sampled pulse has no power content at DC and the main portion of the signal power is concentrated around the baud-rate frequency where the signal propagation velocity is maximum.

## 7.5 Circuit Implementation

As mentioned in previous section, in the proposed design a capacitive driver is employed to achieve small voltage swing and reduce power consumption. It should be noted that the coupling capacitor limits the maximum number of consecutive ones/zeros due to the voltage drift associated with the high-pass behavior of the link. As a result, the coupling capacitor,  $C_p$  is optimized to reduce the time constant of the drift process while providing reasonable voltage swing and bandwidth enhancement. In this design a PMOS transistor realizes a 400fF capacitor for the driver. This results in about 140mV voltage swing over a 7mm wire and less than about 1mV drift in voltage after more than 40 consecutive ones/zeros.

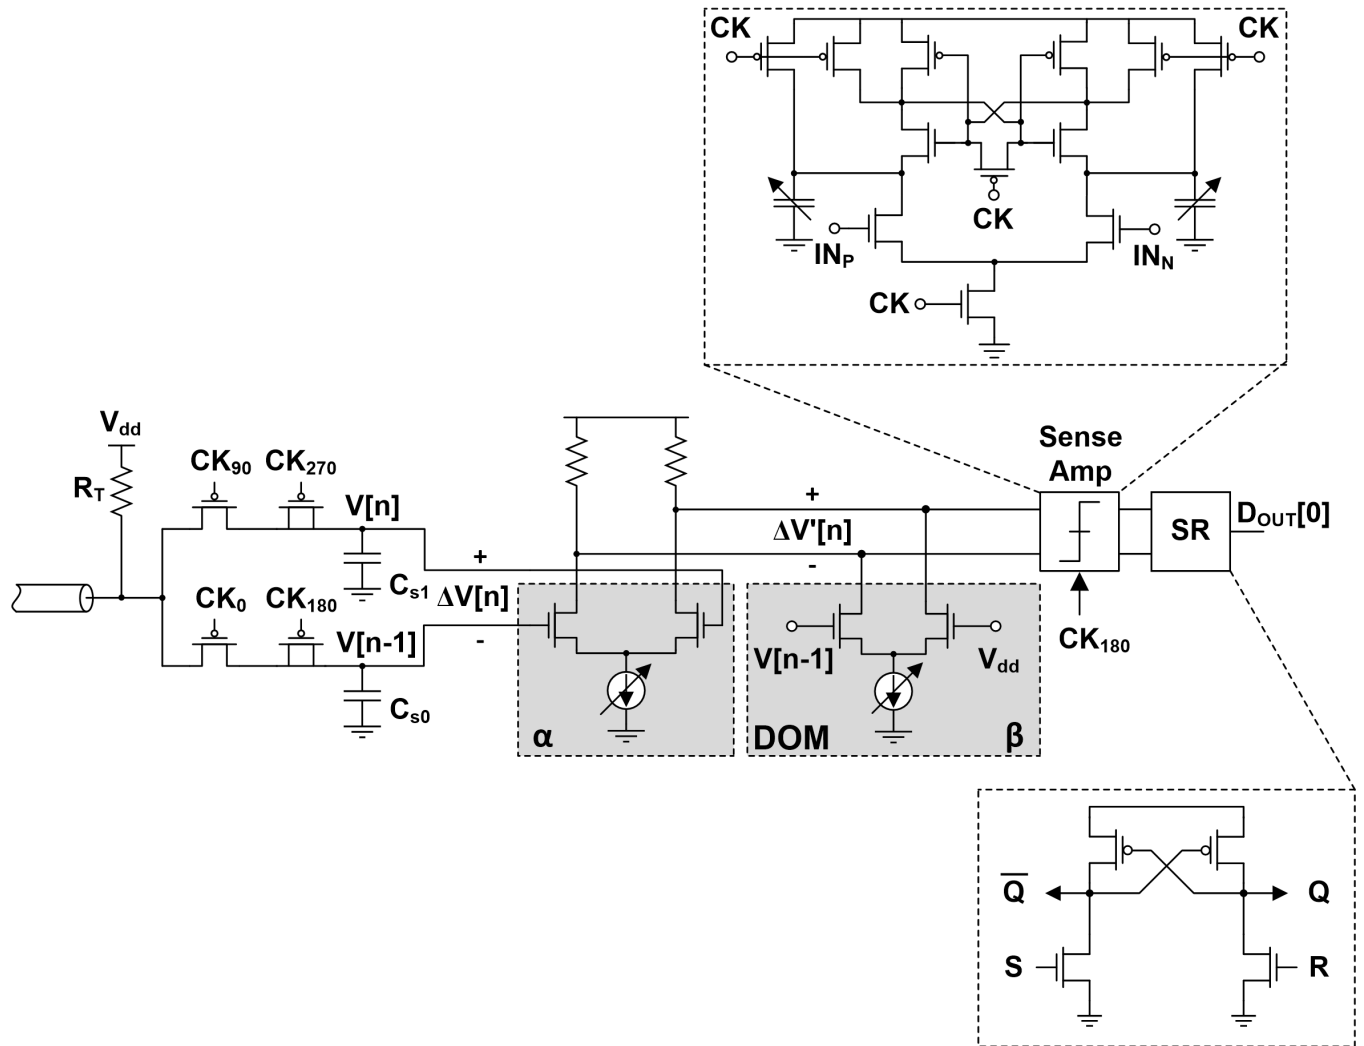


Figure 7.13: Transistor level schematic of the receiver front-end and the StrongArm sense amplifier with capacitive offset cancellation.

The termination resistor sets the receiver’s DC voltage to  $V_{dd}$ . The high DC voltage at the input of the receiver guarantees best operation of the PMOS samplers as shown in Figure 7.13. It also biases the PMOS coupling capacitor in the accumulation regime to ensure maximum capacitance and hence least area. The sampler utilizes dummy switches to reduce charge injection. The residual error due to charge sharing and clock feed-through is removed by the double sampling technique, which performs the single-ended to differential conversion immediately at the receiver input. The input voltage is sampled by a bank of four sample/holds (S/H) driven by quarter-rate clock phases. The sampling capacitor  $C_s$  is chosen such that enough SNR for  $\text{BER} < 10^{-12}$  is achieved. An amplifier with about 6dB gain provides isolation between the sensitive sampling node and the sense amplifier. It also creates a constant common-mode voltage and prevents input dependent offset. A StrongARM sense amplifier is employed to achieve high speed and low power. The sense amplifier has a separate offset cancellation for mismatch compensation through the variable capacitors shown in Figure 7.13. The sense amplifier generates a return-to-zero (RZ) output stream, which is converted to a non-return-to-zero (NRZ) signal through the following SR latch. The transistor level schematic of the latch is also highlighted in Figure 7.13. The wires are implemented using a minimum-pitch ( $0.36\mu\text{m}$  width,  $0.36\mu\text{m}$  spacing)  $M_7$  layer in the 9-metal process where  $M_6$  and  $M_8$  layers are densely populated to mimic orthogonal interconnects in a microprocessor chip. A shielded wiring structure is employed to minimize coupling noise from adjacent lines, as shown in Figure 7.14(a). This provides noise immunity while the double-sampling technique eliminates sensitivity to common-mode interferences at the receiver. Figure 7.14(b) shows the simulated and measured characteristics of the 5mm and 7mm on-chip wires.

## 7.6 Experimental Results

The link prototype was fabricated in 28nm LP CMOS technology with the receiver and transmitter occupying less than  $950\mu\text{m}^2$  and  $160\mu\text{m}^2$ , respectively, as shown in Figure 7.15. The functionality of the transceiver was validated using single-ended on-chip wires with




Figure 7.14: Shielded single-ended on-chip wire (a). Simulated and measured characteristics of the on-chip wires (b).




Figure 7.15: Die micrograph.




Figure 7.16: Total power consumption of the receiver and the transmitter for the 4mm, 5mm, and 7mm links.




Figure 7.17: Power breakdown for the 5mm, and 7mm links at different data rates.

different lengths (4-7mm). PRBS-7 to 31 data was generated off-chip and sent to the on-chip transmitters. Figure 7.17 shows how power consumption (including all clock buffers) changes with increasing the data rate. The 5mm link operates up to 18Gb/s while achieving better than 164fJ/b power efficiency. As the link length and thus wire capacitance increases, signal swing at the receiver degrades. This results in the maximum measured data rate of 15Gb/s for the 7mm link. At this data rate the power efficiency is about 180fJ/b. For the 1.6 $\mu$ m wire pitch, the bandwidth density for the 5mm and 7mm links is 11.25 and 9.375Gb/s/ $\mu$ m, respectively. As this design employs mainly digital blocks, the power consumption almost linearly scales with data rate, Figure 7.17(a).

An optimally repeated version of the link with the same geometry was also simulated for comparison purposes. The proposed scheme offers over 4X improvement in energy efficiency and about 40% lower latency compared to the repeated link. The receiver offers a peak energy efficiency of 136fJ/b at 10Gb/s data rate for 7mm wires. The transceiver was also tested using a 4mm wire. This link is composed of two adjacent wires to investigate the effect of crosstalk. Figure 7.18 shows the on-chip test setup to characterize the effect of adjacent wire crosstalk. Once the aggressor is activated while  $V_{dd}=0.95V$  the SNR drops due to the crosstalk noise, which causes an increase in the BER. By increasing the supply to 1V we could restore the BER, which translates into about 5% degradation in the SNR due to the crosstalk. This level of crosstalk noise is comparable to a twisted differential architecture. The power consumption of the receiver in presence of crosstalk is illustrated in Figure 7.19.






Figure 7.18: Crosstalk measurement setup.




Figure 7.19: Power consumption of the 4mm link in the presence of an aggressor at different data rates.

The advantage of the shielded structure is that it eliminates the Miller capacitance that exists in a differential pair and hence offers a better power efficiency as well as area efficiency. It also provides a return path for the signal and hence limits the extent of magnetic coupling. The immediate single-ended to differential conversion provided by the double-sampling technique also minimizes the sensitivity to common-mode noise. A maximum data rate of 20Gb/s with  $BER < 10^{-12}$  and 180fJ/b of energy efficiency was achieved over this link. Figure 7.20 illustrates the demultiplexed output of the receiver for 10Gb/s and 20Gb/s input data, which is sent off-chip through a 50-ohm output driver. Table 7.1 summarizes the performance of the proposed link and compares it with prior art.

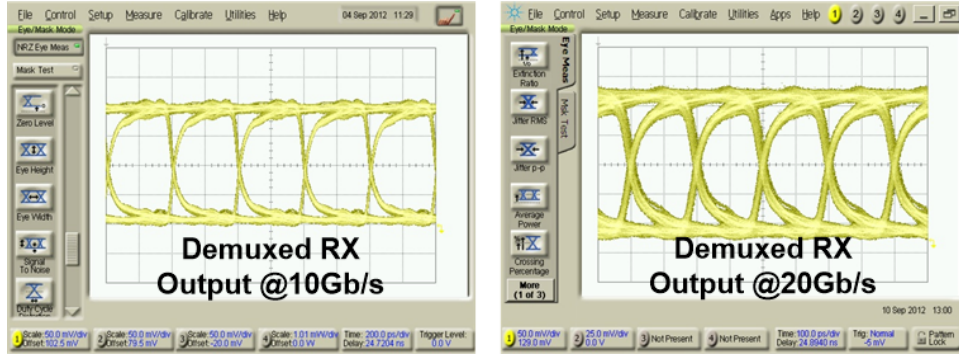


Figure 7.20: Receiver output eye diagram at 10Gb/s (a) and 20Gb/s (b) input data rate..

Technology	28nm Bulk CMOS
Supply	1.0V
Data Rate	20Gb/s
Power Consumption	0.11mW/Gb/s
Bandwidth Density	12.5Gb/s/ $\mu\text{m}$
Channel Length	4mm, 5mm, 7mm
Area	950 $\mu\text{m}^2$ (RX), 160 $\mu\text{m}^2$ (TX)

Table 7.1: OPTICAL RECEIVER PERFORMANCE SUMMARY.

## 7.7 Summary

On-chip wires present both latency and energy challenges to designers. We present a transceiver for repeater-less on-chip communication to address these issues. The proposed design demonstrates high bandwidth density, low latency, and low power consumption. It employs a capacitively-driven transmitter to pre-emphasize signal transitions and hence improve the overall bandwidth of the channel. In addition, it provides low-swing signaling to reduce transmitter power consumption. The receiver utilizes double-sampling in conjunction with dynamic offset modulation to resolve the reduce-swing received signal. A prototype was fabricated in 28nm CMOS to validate the functionality of the link. The transceiver offers up to 20Gb/s/ch data rate and 12.5Gb/s/ $\mu\text{m}$  bandwidth density with better than 180fJ/b energy efficiency. Silicon results show an energy savings of about 4X compared to full-swing CMOS repeaters. The shielded wiring reduces the crosstalk between adjacent channels and allows for minimum-pitch wires. Experimental results show only 5% degradation in SNR due to crosstalk. Since this technique employs mainly digital blocks, it is well-suited for

highly-scaled technologies.

# Chapter 8

## Conclusions

### 8.1 Electrical Interconnects

Most advanced electronic systems today require complex architectures that consist of interconnected integrated circuits (IC). The rapid scaling of CMOS technology continues to increase the processing power of microprocessors and the storage volume of memories which demands a corresponding growth in the communication bandwidth. This increase in the bandwidth can be achieved by employing large numbers of input and outputs (IOs) per chip as well as high data rates per IO. Electrical channels are a natural medium for chip-to-chip communication due to their efficient integration and low cost. Electrical channels are known to impose limited bandwidth mainly due to skin effect and dielectric losses, while CMOS scaling keeps increasing the switching speed of transistors. The dielectric and resistive losses of copper channels increase as the operation frequency increases. Such frequency dependent attenuation causes inter-symbol interference (ISI) and ultimately signal-to-noise-ratio (SNR) degradation. Such frequency dependent attenuations can create severe inter-symbol-interference and hence SNR degradation. Designers combat this problem using equalization. that can be removed by equalization techniques; nevertheless, the power overhead of equalization should be very low. One way to achieve high aggregate data rates is to employ a large number of channels to carry data in parallel. However, placing transmission channels in close proximity causes crosstalk and results in poor signal integrity. Therefore low-power crosstalk cancellation techniques are necessary to allow for a high level of parallelism.

In this thesis, we introduced a switched-capacitor equalization technique which allows the implementation of many taps of decision feedback equalization in an ultra low-power fashion. The key advantage of this technique is that the equalization tap summation is performed in the charge domain to enable low-power operation. Unlike conventional current-mode summers, the charge-mode summer does not require static current. It can also benefit from technology scaling due to switch performance improvement and reduced power consumption of the clock distribution network. Coupling noise is addressed in the proposed design through an efficient far-end crosstalk cancellation scheme. From electromagnetic theory we know that an aggressor signal experiences a differentiator operation to reach the far-end of the coupled line. Therefore, we have employed a simple RC as a differentiator to mimic the behavior of the FEXT noise. As a result, the crosstalk cancellation scheme imposes minimal power and area overhead.

To demonstrate the functionality of the proposed schemes a 4-tap architecture was implemented in which two taps were utilized for equalization and the other two for crosstalk cancellation. The prototype was fabricated in 45nm SOI CMOS. Half-rate clocking and loop-unrolling was employed to enable high data rates and perform data demultiplexing by a factor of two. To further reduce power consumption, analog multiplexers are used and combined with latches. This helps to reduce the number of digital latches. This receiver is suitable for channels with a considerable amount of ISI and crosstalk noise. The simple, low-power DFE can significantly enhance the data rate over lossy channels. In this design, high power efficiency (0.5mW/Gb/s) is achieved by using an SC summation technique, analog multiplexers, and half-rate clocking. The crosstalk cancellation removes more than 75% of crosstalk noise with less than 5% ( $33\mu\text{W}/\text{Gbps}$ ) extra power dissipation. Experimental results validate the feasibility of the DFE receiver for ultra-low-power, high-data rate and highly parallel I/O links.

Next, we introduced a high-speed low-power transmitter design. The proposed transmitter, in conjunction with the DFE receiver, enables highly parallel communication over bandwidth-limited copper channels required for future board-level chip-to-chip links. It employs passive equalization to compensate for the frequency-dependent loss of the channel.

Unlike conventional transmitter equalization techniques, the proposed design avoids active delay elements to reduce total power consumption. Instead, equalization is performed using a compact on-chip 3D inductor with adjustable size. Measurement results show that the transmitter supports data transmission over PCB channels with loss levels in excess of 20dB.

In summary, in this dissertation we have demonstrated a compact low-power electrical transceiver design suitable for implementation in highly advanced CMOS technologies. The proposed design addresses the increasing demand for chip-to-chip IO bandwidth scaling for future high-performance computing systems.

## 8.2 Optical Link Performance Summary

For practical optical chip IOs, small size, low-power CMOS interface circuits are required. In this thesis we proposed a novel RC front-end receiver to achieve these requirements. The receiver employs combined double-sampling and dynamic offset modulation to achieve low-power consumption and high data rate. It removes the need for high gain stages (employed in TIA-based receivers) operating at the bit rate while minimally compromising sensitivity, making it suitable for implementation in highly scaled CMOS technologies.

In this receiver, the optically generated current from the photodiode is converted into a voltage using a resistor at the front-end. The resulting voltage will be sampled at the end of two consecutive bit-times and compared for data recovery. In order to avoid data-dependent sensitivity degradation, the dynamic offset modulation technique is employed. In this technique, the offset of the receiver comparator is changed based on the received pattern to ensure constant double-sampled voltage. The data rate for this front-end is limited by the aperture time of the samplers. Thus by using de-multiplexing and parallelism it can support very high data rates while the timing constraints on the following comparators are relaxed. For this receiver we could achieve up to 24Gb/s in a 65nm CMOS test-chip.

The sensitivity of this receiver is determined by the total capacitance at the input node and noise and offset of the samplers and comparators. For a certain photodiode capacitance and de-multiplexing factor, we can optimize the sizing of the sampling capacitors for the best

sensitivity. The optimum sampling capacitance is a compromise between the  $kT/C$  noise and total capacitance of the input node. Further improvement of receiver sensitivity is possible only if the photodiode capacitance is reduced. The offset of the comparator also directly degrades the sensitivity and should be minimized. In our design an offset compensation technique is used to allow smaller transistors and lower power consumption.

The low power consumption of this receiver makes it an excellent choice for a dense array of receivers on-chip. The scaling behavior of this receiver was discussed in Chapter 5. The data rate per receiver increases with the same rate as feature sizes scale down. The advances in silicon photonics and integration technologies also promise low parasitic capacitances due to the photodiode. In Chapter 5 we also discussed how the data rate and sensitivity of the receiver improves with photodiode capacitance scaling. Since the sizings of the samplers and comparators are dictated by the noise, the power of the front-end scales slower than the digital and clocking circuits. However, the main portion of the receiver power consumption is associated with the digital blocks, thus the overall power consumption scales rapidly with the scaled technologies. For instance, simulations show a factor of two reduction in power by scaling from 65nm to 28nm CMOS technology. A major design challenge with advanced technologies is the reduced power supply voltage. The resistive front-end plays a key role in maintaining the performance of the receiver for reduced power supplies as it automatically limits the input node voltage and avoids non-linear operation of the sampler. In addition, a variable input resistance can guarantee linear operation of the receiver front-end for different input optical powers and hence rectifies the limited maximum optical power problem. It also eliminates the need for data encoding, which is necessary in integrating receivers to avoid long sequences of ones or zeros.

The proposed front-end needs a synchronous clock signal to perform the double sampling and comparison. In Chapter 4 we explored the possibility of baud-rate phase recovery technique for the proposed receiver. In this technique the samples that are separated by two bits are compared to extract phase information. Although this phase measurement is noisy and has low gain, it is effective in a low-bandwidth phase recovery loop. The principles and performance of this technique are also discussed in Chapter 4.

## 8.3 On-Chip Interconnects

This work described the limitations imposed by on-chip wires on the performance of integrated circuits. Wire delays are becoming larger while gate delays scale down. The shrinking cross-sectional area of the wires in advanced technologies makes them highly resistive, which in conjunction with the wire capacitance, creates severe latency in long global wires. To mitigate this problem, repeated wires have been commonly employed by designers to take advantage of the improved transistor quality and compensate for the wire shortcomings. The power and area overhead of the repeated wires is becoming prohibitive as the technology scales, therefore new on-chip signaling techniques are necessary to avoid wire-imposed performance limitations.

We introduced different techniques used by designers to achieve efficient on-chip interconnects such as low-swing signaling, current-mode drivers, transmission line-based interconnects, and equalization. We showed the inadequacy of these technique in meeting the stringent bandwidth requirements of future integrated circuits.

Next we proposed a novel technique inspired by the double-sampling RC receiver discussed in Chapter 5. Due to the RC-limited nature of the on-chip wires, sharp transmitted signals experience long exponential post-cursor tails, which can be removed using many taps of equalization at the cost of excessive power consumption. Instead, we employ double-sampling to measure the difference between the input voltage due to consecutive bits and resolve the incoming data. This technique offers low-power operation and is suitable for implementation in highly advanced CMOS technologies. To further reduce power consumption, a capacitive-driven transmitter is also employed, which enables low-swing signaling. The functionality of the proposed on-chip link was validated through silicon implementation. In summary, the proposed approach can provide the required bandwidth for future many-core processors.



# Bibliography

- [1] G. E. Moore, "Cramming More Components onto Integrated Circuits," *Electronics*, vol.38, no.8, April 1965.
- [2] W. Knight, "Two Heads Are Better Than One," *IEEE Review*, Sept. 2005.
- [3] J. Shalf et al., "Exascale Computing Technology Challenges," in Proceedings of *International Conference on High Performance Computing for Computational Science*, pp. 1-25, 2011.
- [4] International Technology Roadmap for Semiconductors (ITRS) 2011 Update. Semiconductor Industry Association (SIA), 2011.
- [5] Frank O. Mahony et al, "The Future of Electrical I/O for Microprocessors," *VLSI Design, International Symposium on VLSI Design, Automation and Test*, pp. 31-34, April 2009.
- [6] J. U. Knickerbocker et al, "3D Silicon Integration," *Electronic Components and Technology Conference*, pp. 538-543, May 2008.
- [7] C. Delong et al., "A Dual 23Gb/s CMOS Transmitter/Receiver Chipset for 40Gb/s RZ-DQPSK and CS-RZ-DQPSK Optical Transmission," *IEEE ISSCC Digest Technical Papers*, pp. 330-331, February 2012.
- [8] M. Harwood et al., "A 225mW 28Gb/s SerDes in 40nm CMOS with 13dB of Analog Equalization for 100GBASE-LR4 and Optical Transport Lane 4.4 Applications," *IEEE ISSCC Digest Technical Papers*, pp. 326-327, February 2012.
- [9] N. Kocaman et al., "11.3Gb/s CMOS SONET-Compliant Transceiver for Both RZ and NRZ Applications," *IEEE ISSCC Digest Technical Papers*, pp. 142-143, February 2011.
- [10] J. L. Zerbe et al., "Equalization and Clock Recovery for a 2.5-10-Gb/s 2-PAM/4-PAM backplane transceiver cell," *IEEE Journal of Solid State Circuits*, vol. 38, No. 12, pp. 2121-2130, Dec 2003.
- [11] R. Payne et al., "A 6.25-Gb/s Binary Transceiver in 0.13- $\mu$ m CMOS for Serial Data Transmission Across High Loss Legacy Backplane Channels," *IEEE Journal of Solid State Circuits*, vol. 40, No. 12, pp. 2646-2657, Dec 2005.

- [12] J. F. Bulzacchelli et al., "A 10-Gb/s 5-Tap DFE/4-Tap FFE Transceiver in 90-nm CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 41, No. 12, pp. 2885-2900, Dec 2006.
- [13] F. O'Mahony, J. Kennedy, J. E. Jaussi, G. Balamurugan, M. Mansuri, C. Roberts, S. Shekhar, R. Mooney, B. Casper, "A  $47 \times 10$  Gb/s 1.4mW/(Gb/s) Parallel Interface in 45nm CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 156-157, February 2009.
- [14] H. Lee et al., "A 16 Gb/s/Link, 64 GB/s Bidirectional Asymmetric Memory Interface," *IEEE Journal of Solid State Circuits*, vol. 44, No. 4, pp. 1235-1247, April 2009.
- [15] K. Kaviani et al., "A Tri-Modal 20-Gbps/Link Differential/DDR3/GDDR5 Memory Interface," *IEEE Journal of Solid State Circuits*, vol. 47, No. 4, pp. 926-937, April 2012.
- [16] S. J. Bae et al., "A 60nm 6Gb/s/pin GDDR5 Graphics DRAM with Multifaceted Clocking and ISI/SSN-Reduction Techniques," *IEEE ISSCC Digest Technical Papers*, pp. 278-279, February 2008.
- [17] M. Horowitz C. K. K. Yang, S. Sidiropoulos, "High-Speed Electrical Signaling: Overview and Limitations," *IEEE Micro*, vol. 18, no. 1, pp. 12-24, Feb 1998
- [18] C. K. K. Yang, "Design of High-Speed Serial Links in CMOS," Ph.D. Dissertation, Stanford University, Dec 1998.
- [19] J. G. Proakis, M. Salehi, "Digital Communications," McGraw-Hill, 2007.
- [20] R. Kollipara et al. "Design, Modeling and Characterization of High-Speed Back-Plane Interconnects," DesignCon, 2003.
- [21] M. Shafer, B Das, G. Patel, "Connector and Chip Vendors Unite to Produce a High-Performance 10 Gb/s NRZ-Capable Serial Backplane," DesignCon, 2003.
- [22] V. Stojanovic, M. Horowitz, "Modeling and Analysis of High-Speed Links," in Proceedings of *IEEE Custom Integrated Circuit Conference*, pp. 589-594, September 2003.
- [23] H. Johnson, M. Graham, "High-Speed Digital Design: A Handbook of Black Magic," Prentice Hall, 1993.
- [24] V. Stojanovic, "Channel-Limited High-Speed Links: Modeling, Analysis and Design," Ph.D. Dissertation, Stanford University, Sept 2004.
- [25] J. F. Buckwalter, "Deterministic Jitter in Broadband Communication," Ph.D. Dissertation, California Institute of Technology, Jan 2006.
- [26] D. M. Pozar, "Microwave Engineering," John Wiley & Sons, 2005.
- [27] S. Kao, S. Liu, "A 7.5-Gb/s One-Tap-FFE Transmitter with Adaptive Far-End Crosstalk Cancellation Using Duty Cycle Detection," *IEEE Journal of Solid State Circuits*, vol. 48, No. 2, pp. 889-896, Feb 2013.

- [28] C. Menolfi, T. Toifl, M. Rueegg, M. Braendli, P. Buchmann, M. Kossel, T. Morf, "A 14Gb/s High-Swing Thin-Oxide Device SST TX in 45nm CMOS SOI," *IEEE ISSCC Digest Technical Papers*, pp. 156-157, February 2011.
- [29] W. D. Dettloff, J. C. Eble, L. Lei, P. Kumar, F. Heaton, T. Stone, B. Daly, "A 32mW 7.4Gb/s Protocol-Agile Source-Series-Terminated Transmitter in 45nm CMOS SOI," *IEEE ISSCC Digest Technical Papers*, pp. 370-371, February 2010.
- [30] C. Menolfi, T. Toifl, P. Buchmann, M. Kossel, T. Morf, J. Weiss, M. Schmatz, "A 16Gb/s Source-Series Terminated Transmitter in 65nm CMOS SOI," *IEEE ISSCC Digest Technical Papers*, pp. 446-447, February 2007.
- [31] C. K. K. Yang, M. A. Horowitz, "A 0.8- $\mu$ m CMOS 2.5 Gb/s Oversampling Receiver and Transmitter for Serial Links," *IEEE Journal of Solid State Circuits*, vol. 31, No. 12, pp. 2015-2023, Dec 1996.
- [32] A. DeHon, T. Knight Jr., T. Simon, "Automatic Impedance Control," *IEEE ISSCC Digest Technical Papers*, pp. 164-165, Feb 1993.
- [33] W.J. Dally and J. Poulton. "Transmitter Equalization for 4-Gbps Signaling," *IEEE Micro*, vol. 17, no. 1, pp. 48-56, Jan 1997.
- [34] M. E. Lee, W. J. Dally, P. Chiang, "Low-Power Area-Efficient High-Speed I/O Circuit Techniques," *IEEE Journal of Solid State Circuits*, vol. 35, No. 11, pp. 1591-1599, Nov 2000.
- [35] Y. Hidaka, G. Weixin, T. Horie, J. H. Jiang, Y. Koyanagi, H. Osone, "A 4-Channel 1.25–10.3 Gb/s Backplane Transceiver Macro With 35 dB Equalizer and Sign-Based Zero-Forcing Adaptive Control," *IEEE Journal of Solid State Circuits*, vol. 44, No. 12, pp. 3547-3559, Dec 2009.
- [36] H. Cirit, M. J. Loinaz, "A 10Gb/s Half-UI IIR-Tap Transmitter in 40nm CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 448-450, Feb 2011.
- [37] M. Tomlinson, "New Automatic Equalizer Employing Modulo Arithmetic," *IEEE Electronics Letters*, vol. 7, no. 5, pp. 138-139, March 1971.
- [38] T. H. Lee, "The Design of CMOS Radio-Frequency Integrated Circuits," Cambridge University Press, 2003.
- [39] M. Austin, "Decision-Feedback Equalization for Digital Communication over Dispersive Channels," M.I.T. Res. Lab Electron., Tech. Rep. 461, Aug 1967.
- [40] S. Kasturia, J. H. Winters, "Techniques for High-Speed Implementation of Nonlinear Cancellation," *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 5, pp. 711-717, June 1991.
- [41] J. F. Bulzacchelli et al., "A 28Gb/s 4-Tap FFE/15-Tap DFE Serial Link Transceiver in 32nm SOI CMOS Technology," *IEEE ISSCC Digest Technical Papers*, pp. 324-326, Feb 2012.

- [42] T. Toifl, M. Ruegg, R. Inti, C. Menolfi, M. Brandli, M. Kossel, P. Buchmann, P. A. Francese, T. Morf, "A 3.1mW/Gbps 30Gbps Quarter-Rate Triple-Speculation 15-tap SC-DFE RX Data Path in 32nm CMOS," *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 102-103, June 2012.
- [43] H. Sugita, K. Sunaga, K. Yamaguchi, M. Mizuno, "A 16Gb/s 1st-Tap FFE and 3-Tap DFE in 90nm CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 162-163, February 2010.
- [44] A. Hazneci, S. P. Voinigescu, "A 49-Gb/s, 7-Tap Transversal Filter in 0.18- $\mu$ m SiGe BiCMOS for Backplane Equalization," in proceedings of *IEEE Compound Semiconductor Integrated Circuits Symposium*, pp. 101-104, Oct 2004.
- [45] J. Sewter, A. C. Carusone, "A 3-Tap FIR Filter With Cascaded Distributed Tap Amplifiers for Equalization up to 40 Gb/s in 0.18-  $\mu$ m CMOS," *IEEE Journal of Solid State Circuits*, vol. 41, no. 8, pp. 1919-1929, Aug 2006.
- [46] A. Momtaz, M. I. M. Green, "An 80 mW 40 Gb/s 7-tap T/2-Spaced Feed-Forward Equalizer in 65 nm CMOS," *IEEE Journal of Solid State Circuits*, vol. 45, no. 3, pp. 629-639, March 2010.
- [47] X. Lin, S. Saw, J. Liu, "A CMOS 0.25- $\mu$ m Continuous-Time FIR Filter With 125 ps Per Tap Delay as a Fractionally Spaced Receiver Equalizer for 1-Gb/s Data Transmission," *IEEE Journal of Solid State Circuits*, vol. 40, no. 3, pp. 593-602, March 2005.
- [48] T. Beukema et al., "A 6.4-Gb/s CMOS SerDes Core With Feed-Forward and Decision-Feedback Equalization," *IEEE Journal of Solid State Circuits*, vol. 40, no. 12, pp. 2633-2645, Dec 2005.
- [49] J. E. Proesel, T. O. Dickson, "A 20-Gb/s, 0.66-pJ/bit Serial Receiver with 2-Stage Continuous-Time Linear Equalizer and 1-Tap Decision Feedback Equalizer in 45nm SOI CMOS," in *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 206-207, June 2011.
- [50] Y. Kudoh, M. Fukaishi, M. Mizuno, "A 0.13- $\mu$ m CMOS 5-Gb/s 10-m 28AWG Cable Transceiver With No-Feedback-Loop Continuous-Time Post-Equalizer," *IEEE Journal of Solid State Circuits*, vol. 38, no. 5, pp. 741-746, May 2003.
- [51] C. Hogge, "A Self Correcting Clock Recovery Circuit," *IEEE Journal of Lightwave Technology*, vol. 3, no. 6, pp. 1312-1314, Dec 1985.
- [52] J. D. H. Alexander, "Clock Recovery from Random Binary Signals," *IEEE Electronics Letters*, vol. 11, pp. 541-542, 1975.
- [53] C.-K. K. Yang et al., "A 0.5 $\mu$ m CMOS 4Gb/s Serial Link Transceiver with Data Recovery Using Oversampling," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 5, pp. 713-722, May 1998.

- [54] J. Lee, K. C. Wu, "A 20-Gb/s Full-Rate Linear Clock and Data Recovery Circuit With Automatic Frequency Acquisition," *IEEE Journal of Solid State Circuits*, vol. 44, no. 12, pp. 3590-3602, Dec 2009.
- [55] Y. M. Greshichchev et al., "A Fully Integrated SiGe Receiver IC for 10-Gb/s Data Rate," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 12, pp. 1949-1957, Dec 2000.
- [56] R. Farjad-Rad, C. K. K. Yang, M. A. Horowitz, "A 0.3 $\mu$ m CMOS 8-Gb/s 4-PAM Serial Link Transceiver," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 5, pp. 757-764, May 2000.
- [57] J. T. Stonick, G. Wei, J. L. Sonntag, D. K. Weinlader, "An Adaptive PAM-4 5-Gb/s Backplane Transceiver in 0.25- $\mu$ m CMOS," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 3, pp. 436-443, March 2003.
- [58] G. Balamurugan, F. O'Mahony, M. Mansuri, J. E. Jaussi, J. T. Kennedy, B. Casper, "A 5-to-25Gb/s 1.6-to-3.8 mW/(Gb/s) Reconfigurable Transceiver in 45nm CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 372-373, February 2010.
- [59] J. E. Cunningham, D. Beckman, X. Zheng, D. Huang, T. Sze, A. V. Krishnamoorthy, "PAM-4 Signaling over VCSELs with 0.13 $\mu$ m CMOS Chip Technology," *Optics Express*, vol. 14, no. 25, pp. 12028-12038, 2006.
- [60] G. Cherubini et al., "Filter Bank Modulation Techniques for Very High-Speed Digital Subscriber Lines," *IEEE Communications Magazine*, vol. 38, no. 5, pp. 98-104, May 2000.
- [61] A. Amirkhany et al., "A 24Gb/s Software Programmable Multi-Channel Transmitter," *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 38-39, June 2007.
- [62] A. Amirkhany et al., "Analog Multi-Tone Signaling for High-Speed Backplane Electrical Links," *IEEE Global Telecommunications Conference*, pp. 1-6, Nov. 2006.
- [63] S. Palermo, "Design of High-Speed Optical Interconnect Transceivers," Ph.D. Dissertation, Stanford University, Sept 2004.
- [64] A. Emami-Neyestanak, A. Varzaghani, J. F. Bulzacchelli, R. Rylyakov, C. K. Yang, D. J. Friedman, "A 6.0 mW, 10.0 Gb/s Receiver with Switched-Capacitor Summation DFE," *IEEE Journal of Solid State Circuits*, vol. 42, No. 4, pp. 889-896, April 2007.
- [65] Y. Liu, B. Kim, T. O. Dickson, J. F. Bulzacchelli, D. J. Friedman, "A 10Gb/s Compact Low-Power Serial I/O with DFE-IIR Equalization in 65nm CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 182-183, February 2009.
- [66] T. O. Dickson, J. F. Bulzacchelli, D. J. Friedman, "A 12-Gb/s 11-mW Half-Rate Sampled 5-Tap Decision Feedback Equalizer with Current-Integrating Summers in 45-nm SOI CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 44, pp. 1298-1305, April 2009.

- [67] M. Park, J. F. Bulzacchelli, M. Beakes, D. J. Friedman, "A 7Gb/s 9.3mW 2-Tap Current-Integrating DFE Receiver," *IEEE ISSCC Digest Technical Papers*, pp. 230-231, February 2007.
- [68] T. Toifl, C. Menolfi, M. Ruegg, R. Reutemann, A. Prati, D. Gardellini, M. Brandli, M. Kossel, P. Buchmann, P. A. Francese, T. Morf, "A 2.6mW/Gbps 12.5Gbps RX with 8-tap Switched-Cap DFE in 32nm CMOS," *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 210-211, June 2011.
- [69] K. Fukuda et al., "A 12.3mW 12.5Gb/s Complete Transceiver in 65nm CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 368-369, February 2010.
- [70] M. Pozzoni et al., "A 12Gb/s 39dB Loss-Recovery Unclocked-DFE Receiver with Bi-dimensional Equalization," *IEEE ISSCC Digest Technical Papers*, pp. 164-165, February 2010.
- [71] S. A. Ibrahim, B. Razavi, "A 20 Gb/s 40 mW Equalizer in 90nm CMOS Technology," *IEEE ISSCC Digest Technical Papers*, pp. 170-171, February 2010.
- [72] K. Krishna et al., "A 0.6 to 9.6 Gb/s Binary Backplane Transceiver Core in 0.13 $\mu$ m CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 64-65, December 2005.
- [73] C. Pelard et al., "Realization of Multigigabit Channel Equalization and Crosstalk Cancellation Integrated Circuits," *IEEE Journal of Solid State Circuits*, vol. 39, pp. 1659-1670, October 2004.
- [74] H. K. Jung et al., "A 4 Gb/s 3-bit Parallel Transmitter with the Crosstalk-Induced Jitter Compensation Using TX Data Timing Control," *IEEE Journal of Solid State Circuits*, vol. 44, November 2009.
- [75] H. K. Jung, I.-M. Yi, S.-M. Lee, J.-Y. Sim, H.-J. Park, "A Transmitter to Compensate for Crosstalk-Induced Jitter by Subtracting a Rectangular Crosstalk Waveform from Data Signal During the Data Transition Time in Coupled Microstrip Lines," *IEEE Journal of Solid State Circuits*, vol. 47, no. 9, pp. 2068-2079, Sept 2012.
- [76] J. F. Buckwalter, A. Hajimiri, "Cancellation of Crosstalk-Induced Jitter," *IEEE Journal of Solid State Circuits*, vol. 41, no. 3, pp. 621-632, March 2006.
- [77] K.-I. Oh et al., "A 5-Gb/s/pin Transceiver for DDR Memory Interface with a Crosstalk Suppression Scheme," in proceedings of *IEEE Custom Integrated Circuit Conference*, pp. 639-642, September 2008.
- [78] K. J. Sham, R. Harjani, "I/O Staggering for Low-Power Jitter Reduction," in proceedings of *IEEE European Microwave Conference*, pp. 1226-1229, August 2008.
- [79] K. J. Sham, M. R. Ahmadi, S. B. G. Talbot, R. Harjani, "FEXT Crosstalk Cancellation for High-Speed Serial Link Design," in proceedings of *IEEE Custom Integrated Circuit Conference*, pp. 405-407, September 2006.

- [80] S.-Y. Kao, S.-I. Liu, "A 7.5-Gb/s One-Tap-FFE Transmitter with Adaptive Far-End Crosstalk Cancellation Using Duty Cycle Detection," *IEEE Journal of Solid State Circuits*, vol. 48, no. 2, pp. 391-404, Feb 2013.
- [81] T. Oh, R. Harjani, "4x12 Gb/s 0.96 pJ/b/lane Analog-IIR Crosstalk Cancellation and Signal Reutilization Receiver for Single-Ended I/Os in 65 nm CMOS," *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 140-141, June 2012.
- [82] M. Rau, T. Oberst, R. Lares, A. Rothermel, R. Schweer, N. Menoux, "Clock/Data Recovery PLL Using Half-Frequency Clock," *IEEE Journal of Solid State Circuits*, vol. 32, no. 7, pp. 1156-1159, July 1997.
- [83] M. H. Nazari, A. Emami-Neyestanak, "A 15Gb/s 0.5mW/Gbps 2-tap DFE Receiver with Far-End Crosstalk Cancellation," *IEEE ISSCC Digest Technical Papers*, pp. 446-448, February 2010.
- [84] M. H. Nazari, A. Emami-Neyestanak, "A 15Gb/s 0.5mW/Gbps 2-tap DFE Receiver with Far-End Crosstalk Cancellation," *IEEE Journal of Solid State Circuits*, vol. 47, no. 10, pp. 2420-2432, Oct 2012.
- [85] H. Wan, J. Lee, "A 21-Gb/s 87-mW Transceiver With FFE/DFE/ Analog Equalizer in 65-nm CMOS Technology," *IEEE Journal of Solid State Circuits*, vol. 45, no. 4, April 2010.
- [86] P. J. Lim, B. A. Wooley, "An 8-bit 200-MHz BiCMOS Comparator," *IEEE Journal of Solid State Circuits*, vol. 25, no. 1, pp. 192-199, Feb 1990.
- [87] B. Razavi, B. A. Wooley, "Design Techniques for High-Speed, High-Resolution Comparators," *IEEE Journal of Solid State Circuits*, vol. 27, no. 12, pp. 1916-1926, Dec 1992.
- [88] J. W. Jung, B. Razavi, "A 25-Gb/s 5-mWCMOS CDR/Deserialzer," *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 138-139, June 2012.
- [89] S. Palermo et al., "A 90 nm CMOS 16 Gb/s Transceiver for Optical Interconnect," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 5, pp. 1235-1246, May 2008.
- [90] S. U. H. Qureshi, "Adaptive Equalization," *Proceedings of IEEE*, vol. 73, no. 9, pp. 1349-1387, Sept 1985.
- [91] J. H. Winters, R. D. Gitlin, "Electrical Signal Processing Techniques in Long-Haul Fiber-Optic Systems," *IEEE Transactions on Communications*, vol. 38, no. 9, pp. 1439-1453, Sept 1990.
- [92] E.-H. Chen, J. Ren, B. Leibowitz, H.-C. Lee, Q. Lin, K. Oh, F. Lambrecht, V. Stojanovic, J. Zerbe, C.-K.K. Yang, "Near-Optimal Equalizer and Timing Adaptation for I/O Links Using a BER-Based Metric," *IEEE Journal of Solid State Circuits*, Vol. 43, No. 9, pp. 2144-2156, Sept 2008.

- [93] M. S. Lin, A. H. Engvik, J. S. Loos, "Measurements of Crosstalk Between Closely-Packed Lossy Microstrips on Silicon substrates," *IEEE Electronics Letters*, vol 26, pp. 714-716, May 1990.
- [94] J. H. R. Schrader, E. A. M. Klumperink, J. L. Visschers, B. Nauta, "Wireline Equalization Using Pulse-Width Modulation," *IEEE Custom Integrated Circuits Conference*, pp. 591-598, Sept. 2006.
- [95] M. H. Nazari, A. Emami-Neyestanak, "A Low-Power 20Gb/s Transmitter in 65nm CMOS Technology," in proceedings of *IEEE Radio Frequency Integrated Circuits Symposium*, pp. 149-152, June 2012.
- [96] S. Y. Kao, S. I. Liu, "A 20-Gb/s Transmitter with Adaptive Pre-emphasis in 65-nm CMOS Technology," *IEEE Trans. on Circuits and Systems-II*, Vol. 57, No. 5, May 2010.
- [97] Y. Hidaka, et al., "A 4-Channel 1.25–10.3 Gb/s Backplane Transceiver Macro With 35 dB Equalizer and Sign-Based Zero-Forcing Adaptive Control," *IEEE Journal of Solid State Circuits*, vol. 44, no. 12, Dec 2009.
- [98] M. S. Chen, Y. N. Shih, C. L. Lin, H. W. Hung, J. Lee, "A 40Gb/s TX and RX Chip Set in 65nm CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 146-147, Feb. 2011.
- [99] J. Kim, et al., "Circuit Techniques for a 40Gb/s Transmitter in 0.13 $\mu$ m CMOS," *IEEE ISSCC Digest Technical Papers*, pp. 150-151, Feb. 2005.
- [100] L. Kazovsky, S. Benedetto, A. Wilner, "Optical Fiber Communication Systems," Artech House, 1996.
- [101] M. Asghari, A. V. Krishnamoorthy, "Silicon Photonics: Energy-Efficient Communication," *Nature Photonics*, pp. 268-270, 2011.
- [102] K. W. Goossen, J. E. Cunningham, W. Y. Jan, "GaAs 850nm Modulators Solder-Bonded to Silicon," *IEEE Photonic Technology Letters*, vol. 5, no. 7, 776–778, July 1993.
- [103] S. Cheramy et al., "3D integration Process Flow for Set-Top Box Application: Description of Technology and Electrical Results," *IEEE European Microelectronics and Packaging Conference*, pp. 1-6, June 2009.
- [104] C. Gunn, "CMOS Photonics for High-Speed Interconnects," *IEEE Micro*, vol. 26, no. 2, pp. 58-66, March 2006.
- [105] D. M. Kuchta et al., "120-Gb/s VCSEL-Based Parallel-Optical Interconnect and Custom 120-Gb/s Testing Station," *IEEE Journal of Lightwave Technology*, vol. 22, no. 9, pp. 2200-2212, Sept 2004.
- [106] B. E. Lemoff et al., "MAUI: Enabling Fiber-to-the-Processor With Parallel Multiwavelength Optical Interconnects," *IEEE Journal of Lightwave Technology*, vol. 22, no. 9, pp. 2043-2054, Sept 2004.



- [107] F. E. Doany et al., "160 Gb/s Bidirectional Polymer-Waveguide Board-Level Optical Interconnects Using CMOS-Based Transceivers," *IEEE Transactions on Advanced Packaging*, vol. 32, no. 2, pp. 345-359, May 2009.
- [108] K. B. Yoon et al., "Optical Backplane System Using Waveguide-Embedded PCBs and Optical Slots," *IEEE Journal of Lightwave Technology*, vol. 22, no. 9, pp. 2119-2127, Sept 2004.
- [109] L. Schares et al., "Terabus: Terabit/Second-Class Card-Level Optical Interconnect Technologies," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 12, no. 5, pp. 1032-1044, Sept 2006.
- [110] A. A. Abidi, "On the Choice of Optimum FET Size in Wide-band Transimpedance Amplifiers," *IEEE Journal of Lightwave Technology*, vol. 6, no. 1, pp. 64-66, Jan 1988.
- [111] M. Ingles, M.S.J. Steyaert, "A 1 Gb/s, 0.7 $\mu$ m CMOS Optical Receiver with Full Rail-to-Rail Output Swing," *IEEE Journal of Solid-State Circuits*, vol. 34, no 7, pp. 971-977, July 1999.
- [112] B. Razavi, "Design of Integrated Circuits for Optical Communication Systems," McGraw-Hill, 2003.
- [113] S. M. Park, H. Yoo, "1.25-Gb/s Regulated Cascode CMOS Transimpedance Amplifier for Gigabit Ethernet Applications," *IEEE Journal of Solid State Circuits*, vol. 39, no. 1, pp. 112-121, Jan 2004.
- [114] B. Zand et al. "A Transimpedance Amplifier with DC-Coupled Differential Photodiode Current Sensing for Wireless Optical Communications," in proceedings of *IEEE Custom Integrated Circuits Conference*, pp. 455-458, May 2001.
- [115] A Bhatnagar et al., "Clocked-Sense-Amplifier-Based Smart-Pixel Optical Receivers," in *Proceedings of the SPIE - The International Society for Optical Engineering*, vol. 5359, no. 1, pp. 352-359, July 2004.
- [116] T. K. Woodward et al., "Receiverless Detection Schemes for Optical Clock Distribution," *IEEE Photonics Technology Letters*, vol. 8, pp. 1067-1069, Aug 1996.
- [117] A. Emami-Neyestanak et al., "A 1.6Gb/s, 3mW CMOS Receiver for Optical Communication," in *IEEE Symposium VLSI Circuits Digest of Technical Papers*, pp. 84-87, June 2002.
- [118] D. Bossert et al., "Production of High-Speed Oxide Confined VCSEL Arrays for Datacom Applications," *Proc. SPIE*, vol. 4649, pp. 142-151, June 2002.
- [119] D. Vez et al., "10 Gbit/s VCSELs for Datacom: Devices and Applications," proceedings of SPIE, vol. 4942, pp. 29-43, April 2003.
- [120] J. Jewell et al., "1310nm VCSELs in 1-10Gb/s Commercial Applications," proceedings of SPIE, vol. 6132, pp. 1-9, Feb 2006.

- [121] M. A. Wistey et al., "GaInNAsSb/GaAs Vertical Cavity Surface Emitting Lasers at 1534nm," *IEEE Electronics Letters*, vol. 42, no. 5, pp. 282-283, March 2006.
- [122] A. Kern, A. Chandrakasan, and I. Young, "18Gb/s Optical I/O: VCSEL Driver and TIA in 90nm CMOS," *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, June 2007.
- [123] A. V. Rylyakov, C. L. Schow, B. G. Lee, F. E. Doany, C. W. Baks, and J. A. Kash, "Transmitter Predistortion for Simultaneous Improvements in Bit Rate, Sensitivity, Jitter, and Power Efficiency in 20 Gb/s CMOS-Driven VCSEL Links," *Journal of Lightwave Technology*, vol. 30, no. 4, February 15, 2012.
- [124] K. Ohhata, H. Imamura, Y. Takeshita, K. Yamashita, H. Kanai, and N. Chujo, "Design of a 4x10 Gb/s VCSEL Driver Using Asymmetric Emphasis Technique in 90-nm CMOS for Optical Interconnection," *IEEE Transactions on Microwave Theory and Techniques*, vol. 58, no. 5, May 2010.
- [125] N. Chujo, T. Kawamata, K. Ohhata, T. Ohno, "A 25Gb/s Laser Diode Driver with Mutually Coupled Peaking Inductors for Optical Interconnects," in *Proceedings of Custom Integrated Circuits Conference*, pp. 1-4, September 2010.
- [126] P. Dong, S. Liao, D. Feng, H. Liang, D. Zheng, R. Shafiha, X. Zheng, G. Li, K. Raj, A. V. Krishnamoorthy, and M. Asghari, "High Speed Silicon Microring Modulator Based on Carrier Depletion," *IEEE Optical Fiber Communication Conference*, pp. 1-3, March 2010.
- [127] P. Dong, R. Shafiha, S. Liao, H. Liang, N. Feng, D. Feng, G. Li, X. Zheng, A. V. Krishnamoorthy, and M. Asghari, "Wavelength-tunable Silicon Microring Modulator," *Optics Express*, vol. 18, no. 11, May 2010.
- [128] G. Li, X. Zheng, J. Yao, H. Thacker, I. Shubin, Y. Luo, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy, "25Gb/s 1V-Driving CMOS Ring Modulator with Integrated Thermal Tuning," *Optics Express*, vol. 19, no. 21, May 2011.
- [129] F. Liu, et al., "10Gbps, 530fJ/b Optical Transceiver Circuits in 40nm CMOS," in *IEEE Symposium VLSI Circuits Digest of Technical Papers*, pp. 290-291, June 2011.
- [130] C. Li, R. Bai, A. Shafik, E. Z. Tabasy, G. Tang, C. Ma, C-H. Chen, Z. Peng, M. Fiorentino, P. Chiang, and S. Palermo, "A Ring-Resonator-Based Silicon Photonics Transceiver with Bias-Based Wavelength Stabilization and Adaptive Power-Sensitivity Receiver," *IEEE ISSCC Digest Technical Papers*, pp. 124-125, Feb 2013.
- [131] C. L. Schow, et al., "Low-Power 16x10 Gb/s Bi-Directional Single Chip CMOS Optical Transceivers Operating at < 5 mW/Gb/s/link," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 1, pp. 301-313, Jan 2009.
- [132] C. Kromer et al., "A 100-mW 4x10 Gb/s Transceiver in 80-nm CMOS for High-Density Optical Interconnects," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 12, pp. 2667-2679, Dec 2005.

- [133] T. Takemoto et al., “A Compact 4x25-Gb/s 3.0 mW/Gb/s CMOS-Based Optical Receiver for Board-to-Board Interconnects,” *IEEE Journal of Lightwave Technology*, vol. 28, no. 23, pp. 3343-3350, Dec 2010.
- [134] I. A. Young, et al., “Optical I/O Technology for Tera-Scale Computing,” *IEEE Journal of Solid-State Circuits*, vol. 45, no. 1, pp. 235-248, Jan 2010.
- [135] D. Kucharski et al., “10Gb/s 15mW Optical Receiver with Integrated Germanium Photodetector and Hybrid Inductor Peaking in 0.13 $\mu$ m SOI CMOS Technology,” *IEEE ISSCC Digest Technical Papers*, pp. 360-361, Feb 2010.
- [136] A. Narasimha et al., “A Fully Integrated 4x10-Gb/s DWDM Optoelectronic Transceiver Implemented in a Standard 0.13  $\mu$ m CMOS SOI Technology,” *IEEE Journal of Solid-State Circuits*, vol. 42, no. 12, pp. 2736-2744, Dec 2007.
- [137] J. Montanaro et al., “A 160-MHz, 32-b, 0.5-W CMOS RISC Microprocessor,” *IEEE Journal of Solid-State Circuits*, vol. 31, no. 11, pp.1703-1714, Nov 1996.
- [138] M. H. Nazari et al., “An 18.6Gb/s Double-Sampling Receiver in 65nm CMOS for Ultra Low-power Optical Communication,” *IEEE ISSCC Digest Technical Papers*, pp., Feb. 2012.
- [139] M. J. M. Pelgrom et al., “Matching Properties of MOS Transistors,” *IEEE Journal of Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1439, Oct 1989.
- [140] M. E. Lee et al., “Low-Power Area-Efficient High-Speed I/O Circuit Techniques,” *IEEE Journal of Solid-State Circuits*, vol. 35, no. 11, Nov 2000.
- [141] M. Georgas et al., “A Monolithically-Integrated Optical Receiver in Standard 45-nm SOI,” *IEEE Journal of Solid-State Circuits*, vol. 47, no. 7, pp. 1693-1702, July 2012.
- [142] L. Chen, M. Lipson, “Ultra-Low Capacitance and High Speed Germanium Photodetectors on Silicon,” *Optics Express*, vol. 15, no. 2, pp. 7901-7906, May 2009.
- [143] S. Sidiropoulos et al., “A Semidigital Dual Delay-Locked Loop,” *IEEE Journal of Solid-State Circuits*, vol. 32, no. 11, pp. 1683-1692, Nov. 1997.
- [144] A. Emami-Neyestanak et al., “CMOS Transceiver with Baud-Rate Clock Recovery for Optical Interconnects,” in *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 410–413, Jun. 2004.
- [145] M. H. Nazari et al., “Ultra Low-Power Receiver Design for Dense Optical Interconnects,” *IEEE Optical Interconnects Conference*, May 2012.
- [146] J. Proesel, A. Rylyakov, C. Schow, “Optical Receivers Using DFE-IIR Equalization,” *IEEE ISSCC Digest Technical Papers*, pp., Feb. 2013.
- [147] P. Kapur, “Scaling Induced Performance Challenges/Limitations of On-Chip Metal Interconnects and Comparisons with Optical Interconnects,” Ph.D. Dissertation, Stanford University, May 2002.

- [148] M. T. Bohr, "Interconnect Scaling-The Real Limiter to High Performance ULSI," IEDM Technical Digest, pp. 241-244, 1995.
- [149] K. C. Saraswat and F. Mohammadi, "Effect of Interconnection Scaling on Time Delay of VLSI Circuits," IEEE Transaction on Electron Devices, vol. ED-29, pp. 645-650, 1982.
- [150] G. Rangarajan, "Layer Aware Optimization," EETimes, 2012.
- [151] P. Bai et al., "A 65nm Logic Technology Featuring 35nm Gate Lengths, Enhanced Channel Strain, 8 Cu Interconnect Layers, Low-k ILD and  $0.57 \mu\text{m}^2$  SRAM Cell," in *IEEE International IEDM Technical Digest*, pp. 657-660, 2004.
- [152] S. Natarajan, et al., "A 32nm Logic Technology Featuring 2nd-Generation High-k + Metal-Gate Transistors, Enhanced Channel Strain and  $0.171 \mu\text{m}^2$  SRAM Cell Size in a 291Mb Array", in *IEEE International IEDM Technical Digest*, pp. 1-3, Dec 2008.
- [153] P. Kapur et al., "Technology and Reliability Constrained Future Copper Interconnects. I Resistance Modeling," *IEEE Transaction on Electron Devices*, pp. 590-597, April 2002.
- [154] R. Ho, K. W. Mai, M. A. Horowitz, "The Future of Wires," in *proceedings of IEEE*, vol 89, issue 4, pp. 490-504, April 2001.
- [155] R. Ho, "On-Chip Wires: Scaling and Efficiency," Ph.D. Dissertation, Stanford University, Aug 2003.
- [156] A. Naeemi, R. Venkatesan, and J. D. Meindl, "Optimal Global Interconnects for GSI," *IEEE Transactions on Electron Devices*, vol. 50, no. 4, April 2003.
- [157] K. Banerjee, A. Mehrotra, "Analysis of On-Chip Inductance Effects for Distributed RLC Interconnects," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 21, no. 8, pp. 904- 915, Aug 2002
- [158] Y. I. Ismail, E. G. Friedman, and J. L. Neves, "Figures of Merit to Characterize the Importance of On-Chip Inductance," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 7, no. 4, Dec 1999.
- [159] B. Kleveland, "CMOS Interconnections Beyond 10 GHz.," Ph.D. Dissertation, Stanford University, Nov. 1999.
- [160] T. Sakurai, K. Tamaru, "Simple Formulas for Two- and Three-Dimensional Capacitances," *IEEE Transactions on Electron Devices*, vol 30, issue 2, pp. 183-185, Feb 1983.
- [161] W. C. Elmore, "The Transient Response of Damped Linear Networks with Particular Regard to Wideband Amplifiers," *Journal of Applied Physics*, pp. 55-63, 1948.
- [162] T. Sato et al., "Accurate In-situ Measurement of Peak Noise and Signal Delay Induced by Interconnect Coupling," *IEEE ISSCC Digest Technical Papers*, pp. 226-7, Feb 2000.
- [163] K. Soumyanath et al., "Accurate On-Chip Interconnect Evaluation: A Time-Domain Technique," *IEEE Journal of Solid-State Circuits*, pp. 623-31, May 1999.

- [164] J. Lee, W. Lee, S. H. Cho, "A 2.5-Gb/s On-Chip Interconnect Transceiver With Crosstalk and ISI Equalizer in 130 nm CMOS," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol 59, issue 1, pp. 124-136, Jan 2012.
- [165] S. Morton, "On-chip Inductance Issues in Multiconductor Systems," in Proceedings of *IEEE Design Automation Conference*, pp. 921-6, June 1999.
- [166] L. He et al., "An Efficient Inductance Modeling for On-Chip Interconnects," in *Proceedings of Custom Integrated Circuits Conference*, pp. 457-60, May 1999.
- [167] S. Naffziger, "Design Methodologies for Interconnect in GHz+ ICs," Tutorial, *IEEE International Solid-State Circuits Conference*, Feb 1999.
- [168] K. Gala, V. Zolotov, R. Panda, B. Young, J. Wang, D. Blaauw, "On-chip Inductance Modeling and Analysis," in Proceedings of *Annual Design Automation Conference*, pp. 63-68, 2000.
- [169] J. Zhang, E. G. Friedman, "Effect of Shield Insertion on Reducing Crosstalk Noise Between Coupled Interconnects," in Proceedings of *International Symposium on Circuits and Systems*, pp. 529-532, 2004.
- [170] D. Priore, "Inductance on Silicon for Sub-Micron CMOS VLSI," *IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 17-18, June 1993.
- [171] T. Osaka, M. Yoshino, "New Formation Process of Plating Thin Films on Several Substrates by Means of Self-Assembled Monolayer (SAM) Process," *Electrochimica Acta*, vol 53, issue 2, pp. 271-277, Dec 2007.
- [172] H. Bakoglu, "Circuits, Interconnections and Packaging for VLSI," Addison-Wesley, 1990.
- [173] N. Weste, D. Harris, "CMOS VLSI Design: A Circuits and Systems Perspective," Addison-Wesley, 2010.
- [174] K. Banerjee, A. Mehrotra, "A Power-Optimal Repeater Insertion Methodology for Global Interconnects in Nanometer Designs," *IEEE Transactions on Electron Devices*, vol. 49, no. 11, pp. 2001-2007, Nov 2002.
- [175] N. S. Nise, "Control Systems Engineering," John Wiley & Sons, 2008.
- [176] P. Saxena, N. Menezes, P. Cocchini, D. A. Kirkpatrick, "The Scaling Challenge: Can Correct-by-Construction Design Help?," in proceedings of *IEEE International Symposium on Physical Design (ISPD)*, pp. 51-58, 2003.
- [177] J. Parkhurst, J. Darringer, B. Grundmann, "From Single Core to Multi-Core: Preparing for a New Exponential," in proceedings of *IEEE/ACM international conference on Computer-aided design*, pp. 67-72, 2006.
- [178] N. Magen, A. Kolodny, U. Weiser, N. Shamir, "Interconnect-Power Dissipation in a Microprocessor," in proceedings of *System Level Interconnect Prediction Conference*, pp. 7-13, 2004.

- [179] H. Ito et al., “A Low-Latency and High-Power-Efficient On-Chip LVDS Transmission Line Interconnect for an RC Interconnect Alternative,” *IEEE International Interconnect Technical Conference*, pp. 193-195, 2007.
- [180] R. Chang, N. Talwalkar, C. Yue, and S. Wong, “Near Speed-of-Light Signaling over On-Chip Electrical Interconnects,” *IEEE Journal Solid-State Circuits*, vol. 38, no. 5, pp. 834–838, May 2003.
- [181] D. Schinkel, E. Mensink, E. A. M. Klumperink, E. van Tuijl, B. Nauta, “A 3-Gb/s/ch Transceiver for 10-mm Uninterrupted RC-Limited Global On-Chip Interconnects,” *IEEE Journal Solid-State Circuits*, vol. 41, no. 1, pp. 297-336, Jan 2006.
- [182] J. Seo et al., “High-Bandwidth and Low-Energy On-Chip Signaling with Adaptive Pre-Emphasis in 90nm CMOS,” *IEEE ISSCC Digest of Technical Papers*, pp. 182-183, Feb. 2010.
- [183] R. Ho et al., “High Speed and Low Energy Capacitively Driven On-Chip Wires,” *IEEE Journal Solid-State Circuits*, vol. 43, No. 1, pp. 52-60, 2008.
- [184] D. Walter et al., “A Source-Synchronous 90Gb/s Capacitively Driven Serial On-Chip Link Over 6mm in 65nm CMOS,” *IEEE ISSCC Digest of Technical Papers*, pp. 180-181, Feb. 2012.
- [185] E. Mensink et al., “Power Efficient Gigabit Communication over Capacitively Driven RC-Limited On-Chip Interconnects,” *IEEE Journal of Solid-State Circuits*, vol. 45, No. 2, pp. 447-457, 2010.
- [186] A. P. Jose et al., “Distributed Loss-Compensation Techniques for Energy-Efficient Low-Latency On-Chip Communication,” *IEEE Journal Solid-State Circuits*, vol. 42, No. 6, pp. 1415-1424, 2007.
- [187] B. Kim et al., “A 4Gb/s/ch 356fJ/b 10mm Equalized On-Chip Interconnect with Non-linear Charge-Injecting Transmit Filter and Transimpedance Receiver in 90nm CMOS,” *IEEE ISSCC Digest of Technical Papers*, pp. 66-67, Feb. 2009.
- [188] N. Tzartzanis, W. W. Walker, “Differential Current-Mode Sensing for Efficient On-Chip Global Signaling,” *IEEE Journal Solid-State Circuits*, vol. 40, no. 11, pp. 2141-2147, Nov. 2005.
- [189] S.-K. Lee, S.-H. Lee, D. Sylvester, D. Blaauw, J.-Y. Sim, “A 95fJ/b Current-Mode Transceiver for 10mm On-Chip Interconnect,” *IEEE ISSCC Digest of Technical Papers*, pp. 66-67, Feb. 2013.
- [190] H. G. Rhew et al., “A 22Gb/s, 10mm On-Chip Serial Link over Lossy Transmission Line with Resistive Termination,” in proceedings of *European Solid-State Circuits Conference*, pp. 233-236, 2012.
- [191] Y. Liu, P-H. Hsieh, S. Kim<sup>1</sup>, J-S. Seo, R. Montoye, L. Chang, J. Tierno, and D. Friedman, “A 0.1pJ/b 5-to-10Gb/s Charge-Recycling Stacked Low-Power I/O for On-

- Chip Signaling in 45nm CMOS SOI,” *IEEE ISSCC Digest Technical Papers*, pp., Feb. 2013.
- [192] S-K. Lee, S-H. Lee, D. Sylvester, D. Blaauw, J-Y. Sim, “A 95fJ/b Current-Mode Transceiver for 10mm On-Chip Interconnect,” *IEEE ISSCC Digest Technical Papers*, pp., Feb. 2013.
- [193] M. H. Nazari, A. Emami-Neyestanak, “A 24-Gb/s Double-Sampling Receiver for Ultra-Low-Power Optical Communication,” *IEEE Journal Solid-State Circuits*, vol. 48, No. 2, pp. 344-357, Feb 2013.
- [194] M. H. Nazari, A. Emami-Neyestanak, “A 15-Gb/s 0.5-mW/Gbps Two-Tap DFE Receiver With Far-End Crosstalk Cancellation,” *IEEE Journal Solid-State Circuits*, vol. 47, No. 10, pp. 2420-2432, Oct 2012.